



Artificial

Intelligence and Food Safety: Hype vs. Reality

By Kristen M. Altenburger, A.M., and Daniel E. Ho, Ph.D.

To understand the promise and peril of artificial intelligence for food safety, consider the story of Larry Brilliant. Brilliant is a self-described “spiritual seeker,” “social change addict,” and “rock doc.” During his medical internship in 1969, he responded to a *San Francisco Chronicle* columnist’s call for medical help to Native Americans then occupying Alcatraz. Then came Warner Bros.’ call to have him join the cast of *Medicine Ball Caravan*, a sort-of sequel to *Woodstock Nation*. That caravan ultimately led to a detour to India, where Brilliant spent 2 years studying at the foot of the Himalayas in a monastery under guru Neem Karoli Baba. Toward the end of the stay, Karoli Baba informed Brilliant of his calling: join the World Health Organization (WHO) and eradicate smallpox. Brilliant did just that. He joined the WHO as a medical health officer, as a part of a team making over 1 billion house calls collectively. In 1977, he observed the last human with smallpox, leading WHO to declare the disease eradicated.

After a decade battling smallpox, Brilliant went on to establish and lead foundations and start-up companies, and serve as a professor of international health at the University of Michigan. As one corporate brand manager wrote, “There are stories that are so incredible that not even the creative minds that fuel Hollywood could write them with a straight face.”¹

In 2006, Larry Page and Sergei Brin hired Brilliant as the Executive Director—later, 'Chief Philanthropy Evangelist"—of Google's charitable arm, [Google.org](https://www.google.org), pledging 1 percent of Google's profits and equity to tackle complex social problems like climate change, poverty, and global pandemics. The same year, Brilliant received the TED Prize of \$100,000 for one "wish to change the world," announced in the accompanying TED Talk. The wish? "Early detection, early response." As Brilliant put it, "I envision a kid (in Africa) getting online and finding that there is an outbreak of cholera down the street. I envision someone in Cambodia finding out that there is leprosy across the street."

Answering his call for early detection, Google collaborated with the U.S. Centers for Disease Control and Prevention (CDC) and published a blockbuster paper in *Nature* in 2009, with Brilliant listed as last author. The team used over 1,000 search terms such as "thermometer," "flu," "symptoms," "muscle aches," and "chest congestion" to predict historical flu trends. With some 50 million search queries, Google Flu Trends (GFT), the team posited, could serve as an early warning signal of flu outbreaks. They reported being able to predict influenza 1–2 weeks faster and with higher accuracy than the CDC. The paper stated, "Harnessing the collective intelligence of millions of users...can provide one of the most timely, broad-reaching syndromic surveillance systems available today."²²

Viral prediction went viral. *Forbes* reported that "Google Flu Trends identifies outbreaks more effectively than doctors can."²³ According to an article in the *New York Times*,⁴ MIT professor Thomas W. Malone stated: "I think we are just scratching the surface of what's possible with collective intelligence." *The Times* also quoted Brilliant as stating that GFT "epitomize[s] the power of Google's vaunted engineering prowess to make the world a better place." GFT quickly became the poster child for [Google.org](https://www.google.org)'s style of philanthropy and was expanded to cover 30 countries. The paper has been cited over 3,500 times.

There was only one problem with GFT: It didn't work. Signs came early on. A few months after its release, GFT missed predicting the 2009 swine flu pandemic. Predicting the past is easier than predicting the future. GFT's aspirations to solve global pandemics were mismatched with digital coverage. As the *New York Times* put it, "One major shortcoming of Flu Trends is that in poor regions of the developing world, where devastating pandemics are most likely to start, computers are not widely available, so Google has little data to feed into the tool."²⁵ The death knell came in 2014, when a *Science* paper described the "Google Flu Parable" as the principal exemplar of "big data hubris," the mistaken belief that big data can substitute for traditional principles of data analysis.⁶ The *Science* team showed that the GFT algorithms almost surely engaged in massive overfitting, a no-no in machine learning. (We say "almost surely" because the search terms were never published, making it difficult to replicate the analysis directly, another no-no.) Such overfitting led the GFT team to initially select search terms like those for "high school basketball," meaning it was "part flu detector, part winter detector." Nor was the reliance on ordinary searches obvious. As one public health professional put it, "The problem is that most people don't know what 'the flu' is, and relying on Google searches by people who may be utterly ignorant about the flu does not produce useful information."²⁷

Denunciation came as swiftly as earlier praise. GFT has been called an "epic failure," and the University of Washington Information School features GFT as a case study in its course "Calling Bullshit."²⁸ [Google.org](https://www.google.org) shelved its poster child and ultimately reoriented its philanthropic strategy. The problem, as one former executive said, was that "we were creating solutions that were looking for problems rather than the other way around."

Artificial Intelligence and Food Safety

Much as GFT promised a new, faster, and more accurate surveillance system, can artificial intelligence (AI) and machine learning transform food safety?

Many think so. One trade publication describes AI as "playing a predominant role in the world of food safety and quality assurance."²⁹ Frank Yiannas, the deputy commissioner for food policy and response at the U.S. Food and Drug Administration (FDA), praised AI as one of the best tools for the future of

"...can artificial intelligence...and machine learning transform food safety?"

food safety at the 2019 International Association for Food Protection meeting. And earlier in the year, FDA announced a pilot program to use AI to help ensure food safety of imported foods.¹⁰ Just as AI promises to transform sectors across the economy, foodborne pathogen surveillance may be fundamentally altered by emerging technology.

At the same time, how can policy makers, industry, and consumers tell the difference between hype and reality?

To begin, we should be clear about what we mean by AI. The history of AI is one of boom and busts, with each generation of technology promising solutions with distinct technology, followed by "AI winters" when promises fall short. The changing technology also means that AI defies straightforward definition. What we focus on here is the use of AI in food safety, using emerging advances in *machine learning* to detect and predict risk based on large data sets. These techniques learn (potentially complex) patterns in big data to predict outcomes, often with significant improvements compared with human judgment. The set of techniques is large, but if you have heard of "deep learning," "neural networks," or "random forests," those are core techniques of machine learning.

While this is where much of the excitement resides, this scope of AI ex-

cludes two important categories that are often commingled in loose rhetoric about AI and food safety. First, this notion of AI excludes forms of process automation that may, for instance, improve the traceability of food products (e.g., track-and-trace systems using RFID tags, QR codes, cloud data storage, blockchain) or automate risk factor measurement (e.g., temperature logs). Second, it does not include improvements in testing technology (e.g., loop-mediated isothermal amplification). These developments are extremely important for modernizing the food safety system. Indeed, as we articulate below, these developments—which require basic investments in data infrastructure, sensing technology, and testing—may be more consequential for the future of food safety than AI technology that generates splashy news coverage.

Improved Surveillance through Big Data?

Underreporting, underdiagnosis, and uncertainty are endemic to foodborne disease surveillance.^{11,12} Only one of every 30 cases of nontyphoidal *Salmonella* spp. or *Campylobacter* spp. is captured by laboratory surveillance data.¹³ As a result, one of the biggest hopes is that machine learning with unconventional data can shore up foodborne surveillance systems. What has received a tremendous amount of attention is the use of AI to predict food safety risk from streams of untapped data sources like Twitter, Facebook, Yelp, and other social media. Just like GFT promised to use big data to overcome weaknesses of conventional disease surveillance, scores of ideas and pilots have emerged for conducting such foodborne disease surveillance using big data.¹⁴

Studies have used Yelp reviews in San Francisco and Seattle to predict health code violations^{15–17} and risky food items.¹⁸ The leading study, for instance, uses terms such as “sick,” “vomit,” or “food poisoning” to predict high risk. Others used Twitter to search for tweets indicating foodborne illness.¹⁹ One study used Amazon product reviews to predict FDA food recalls.¹⁷ Google researchers used Google search and location logs—much like GFT did—to predict health code violations.²⁰ And in another approach, the site waspoisoned.com attempts to crowd-source foodborne illness complaints.²¹ Nor are these limited to academic studies. Health departments in Boston, Chicago, Las Vegas, New York City, and St. Louis have each experimented with surveillance systems using Twitter, Yelp, and 311 call data. Earlier this year, Chick-fil-A built out a comparable system to mine social media posts to identify potential food safety issues. While the initial experience with AI in food safety has been in the use of social media for retail food safety, these pilots also have implications for AI in other parts of the food safety system, such as PulseNet, farm and food facility inspections, and food safety standards.

And just as with GFT, breathless media reporting declares how these tools might solve the problem of food safety. *The Atlantic* ran the headline: “Predictive Policing Comes to Restaurants.”²² *CBS News* described Yelp reviews as the “magic bullet for fighting foodborne illness.” *Popular Science* wrote, “Mining tweets for illness-related complaints can tell us what restaurants to avoid when.”²³ *Government Technology* magazine noted that “the algorithms have the potential to improve predictions in any city, based on any type of review.”²⁴ Yelp’s CEO baldly asserted that “online reviews could beat ‘gold standard’ healthcare measures.” Gizmodo ran with the headline “This AI Can Spot Unsafe Food on Amazon Faster Than the FDA.”²⁵ Echoing the claims of the *Nature* paper, one author of the Google study proclaimed, “Today, we can use online data to make epidemiological observations in near real-time, with the potential for significantly improving public health in a timely and cost-efficient manner.”²⁶

Do these claims exhibit big data hubris? We discuss three key areas where data va-

lidity, fairness, and predictive accuracy need to be rigorously examined before the large-scale adoption of AI in public policy occurs.

Validity and Goals in Using AI

When deploying AI, it is worth being clear about goals. There are several key considerations about the validity of social media as a data source to weigh before deploying an AI system in public health.

First, AI applications often gloss over the relationship between public health surveillance and social media. Online,

“Underreporting, underdiagnosis, and uncertainty are endemic to foodborne disease surveillance.”

crowd-sourced reporting tools like Yelp can provide independent, market-based incentives for restaurants and food providers to correct subpar food standards. The reason is that a report of “food poisoning” in a restaurant review can act as an independent sanction and incentive for correction. As one company puts it, “The simple accusation from a reviewer can severely damage any restaurant’s reputation.”²⁷ If so, it is less clear whether inspections should be considered complements to, as opposed to substitutes for, online reviews. Instead of sending health inspectors where social media already covers food safety, we may want to deploy them where Yelp is missing.

Second, it is well-known that complexities in the health code have led inspectors to carry out inspections in dramatically divergent ways.²⁸ *Who* does the inspection can matter as much as the conditions on the premises. In King County, Washington, for instance, one inspector was widely viewed as “the toughest health inspector,” citing violations at orders of magnitude higher than others. Sending more inspectors to locations where citations are predicted to be high may be exactly the wrong thing to do. In King County, it would mean de-

ploying more resources to zones where the health code is already the most stringently enforced. To the contrary, we might want to deploy resources to areas where the health code is underenforced due to inspector lenience.

Third, there is considerable debate about whether health code violations predict foodborne illness outbreaks. One study in King County found an association between violations and outbreaks,²⁹ but studies in Miami-Dade County and in Tennessee found none.^{30,31} Part of the explanation for such mixed findings is that inspectors may not conceive of their role as mindless checklisting, but rather as identifying the greatest health risks that can be remediated. As one scholar wrote, “If all meat-inspection regulations were enforced to the letter, no meat processor in America would be open for business...[T]he inspector is not expected to enforce strictly every rule, but rather to decide which rules are worth enforcing at all.”

Last, crafting enforcement based on social media can set up perverse incentives. Users and competitors, for instance, can submit suspicious terms to trigger health inspections, without the procedural and scientific protections of health department screenings of complaints. Health department reporting systems typically build in the systematic collection of symptoms and food histories to facilitate investigation, which is lost through unstructured speculation on social media.

Fairness Is Still a Challenge

One of the emerging areas of concern is that algorithms, rather than addressing human bias, can encode it. Google searches are systematically more likely to turn up advertisements of criminal records for black-sounding names (e.g., Latanya) than for white-sounding names (e.g., Allison).³² Gender prediction based on facial recognition technology is prone to error for darker skin tones.³³ And there are serious questions about bias with criminal risk assessment scores used for pretrial deten-

tion and sentencing.³⁴ In the food safety space, the leading paper used words such as “Vietnamese” and “Chinese” to infer that an establishment was dirty and words like “Belgian” and “European” to infer that an establishment was clean.¹⁵

Bias can be more subtle too, stemming from lack of representativeness in the data. Recall the shortcoming of GFT: The lack of Google search terms in the poor regions of the developing world made it useless where its need was most vital. Another example comes from Virginia Eubanks, who documents that the risk for child welfare determinations in Allegheny County may import bias because much of the data come from social welfare programs, such as Temporary Assistance for

SPONSORED CONTENT


AI in Security Systems: Innovations to Help Food Producers

The introduction of intelligent video surveillance with artificial intelligence (AI) applications has vastly improved the way security is managed and used. Breakthrough security technologies present a suite of features that can be an integral part of a good food defense plan.

Before now, food manufacturers deployed access control and video surveillance as a security measure only, without any consideration for food regulatory requirements or good working practices.

Food manufacturing companies can think differently about their security systems. It is now possible to deploy physical security that will also provide an additional tier of confidence for the Food Safety Modernization Act intentional adulteration and compliance mandates. Our solution offers:

- Seamless integration between security systems and the food production lines.
- Self-learning video analytics with a directory of event and warning notifications.
- Appearance search to indicate unknown persons at internal and external areas.
- Flexible configurations for workflow routines, access levels, and integrated video feeds.
- License plate recognition and improved security automation for transport and logistics
- New security projects to include specific design emphasis for the food industry.
- Upgraded systems to provide additional intelligent features that work for food producers.



The operation and functionality of a security system’s performance will improve with superior video quality and enterprise-level access control.

Find out more on how your systems can be working better for you. Contact ivelah, your security expert that understands food.

www.ivelah.com/food-safety-and-security-systems/

www.ivelah.com/ • info@ivelah.com • 1.800.216.0805



Needy Families, Supplemental Security Income, Supplemental Nutrition Assistance Program, and county medical assistance.³⁵ Child welfare determinations may hence reflect past poverty as much as future risk.

In food safety, one concern is that social media users are not a representative sample of society.³⁶ To illustrate this, Figure 1³⁷ displays a map of permitted establishments in King County on the left panel, with the number of Yelp reviews in the middle panel. We can see that Yelp reviews are concentrated disproportionately around the city of Seattle. The right panel plots the density of the Asian population living in the county, showing that there's a high negative correlation between Asian consumers and Yelp review penetration. Yelp reviewers are more likely to be white, affluent, and young. As a result, sending out inspectors based on Yelp reviews may risk creating different food safety systems across demographic groups.

Another source of bias lies in consumers themselves. Diagnosing and attributing foodborne illness is notoriously difficult given the incubation periods of foodborne pathogens. Relying on consumers who may be unaware of the etiology of foodborne illness may propagate bias. As *Slate* conjectured, stereotypes harbored by consumers against “ethnic foods” may affect the propensity by consumers to blame such restaurants for food poisoning.³⁸ In prior research, we tested for the presence of such bias by examining whether consumers are more likely to submit a complaint over a 311 call or suspicious terms in a Yelp review for Asian than for non-Asian establishments, holding constant the food safety inspection score.³⁹ We found significant evidence for such implicit bias in reviews: Asian restaurants are more likely to be accused of food poisoning than non-Asian restaurants, even though consumers observe none of the risk factors in the kitchen. The text of reviews corroborates differences in how “ethnic” establishments are described on Yelp.⁴⁰ The adoption of AI for food safety may thus perpetuate societal inequities imprinted in big data.

If these algorithms are adopted for enforcement, one unique risk is the creation of “runaway feedback loops.” A sobering finding is that when police are sent to ZIP codes based on a predictive model, and subsequent arrests are fed back into the predictive model, police

“...there is considerable debate about whether health code violations predict foodborne illness outbreaks.”

may be sent to the same area over and over again, even if the underlying crime rates were random. Risk can be a self-fulfilling prophecy.

Does a Gain in Accuracy Outweigh Bias?

Even given the above concerns, there may be real reasons to consider adoption of these AI tools. After all, the informational challenges to foodborne disease surveillance are so severe that some might deem the gain in accuracy worth the cost in bias. That said, performance can vary significantly: The reported accuracy for Chick-fil-A's

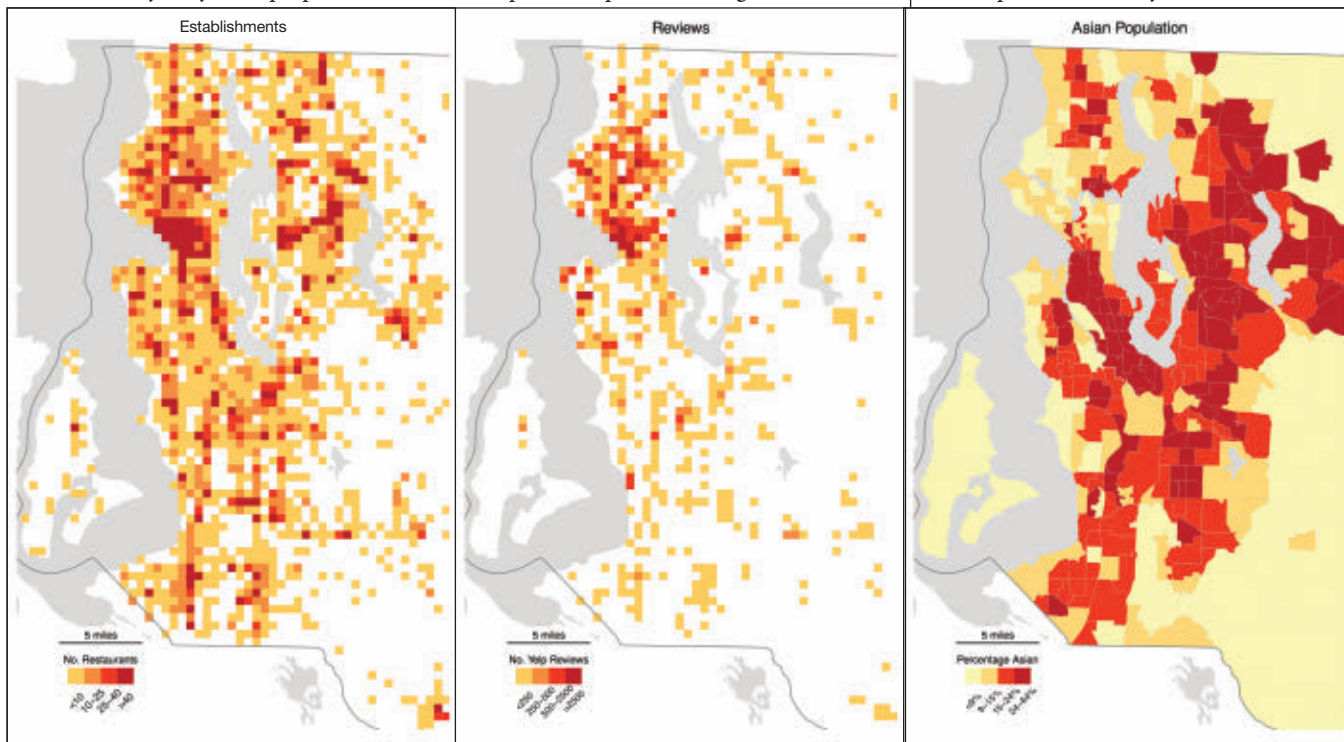


Figure 1. Density of permitted establishments (left panel), Yelp reviews (middle panel), and Asian population (right panel) across King County.³⁷

system is 78 percent, meaning that the prediction is wrong 22 percent of the time. (This assumes a balanced training sample. It is not possible to tell how exactly the accuracy rate was compiled, as the methodological details of Chick-fil-A's system are not documented.) The best model in the Amazon review system retrieves only 71 percent of unsafe food products. And in the Yelp example, using reviews yielded accuracy of 82 percent compared with 72 percent using inspection history alone.

But the story gets worse. One of us has worked for years to improve food safety inspections, evaluating interventions using randomized controlled trials.²⁸ We originally sought to test methods for de-biasing machine learning approaches. Our idea was that it may be possible to use state-of-the-art methods to purge consumer complaints of differences driven by ethnic bias. Yet in embarking on that project, we attempted for months to replicate the leading paper's results that generated the claims that "Yelp might clean up the restaurant industry."⁴¹ That paper merged roughly 152,000 Yelp reviews with roughly 13,000 subsequent inspection scores for 1,756 restaurants in King County from 2006 to 2013.¹⁴ Using the words in the Yelp review, cuisine, ZIP code, and inspection history, it developed a model for subsequent inspection scores. It claimed that words like "gross," "mess," "sticky," "smell," "restroom," and "dirty" predicted health code violations.

Our replication efforts showed that (a) the data set of 13,000 inspections was trimmed to only 612 inspections for analysis; and (b) these 612 observations included only establishments with either extremely high or extremely low inspection scores. By selecting the extreme ends of inspection performance, the paper made the prediction task trivially easier. Consider an analogy to predicting student performance. Instead of predicting the full grade range for students, the original study essentially predicted between straight-A students and dropout students, removing the vast majority of students in between.

Big data claims were based on little data support.

Still, we were so stymied by the replication that we asked the authors to share the original code. Generously, and much to their credit, they did. And after studying some 1,700 lines of code, we discovered one simple mistake. When choosing zero-violation establishments, the program chose systematically from a sorted list, so that the same restaurant was represented many times over multiple inspections. (The standard is random undersampling.) This line generated the artificially high predictive power of Yelp review words, because terms specific to the overrepresented restaurants (e.g., "Vietnamese food") occur in the review text. In a sense, this problem

NEED MEDIA? We're Prepared For You

Prepared Media

MediaBox™ Sterile Liquids
Fresh, ready-to-use broths and buffers in a convenient, stackable storage box.



EZ-Media Prepared Plates

Formulated with the highest quality components.

EZ-Media Bottles

Sterile broths and buffers; 500ml and 1L sizes.

EZ-Media Test Tubes

Prepared liquid solution in 9mL test tubes.



Now ISO 9001:2015 Certified

Microbiology International is pleased to announce it has been accredited to the ISO 9001:2015 International Quality Standard by the Performance Review Institute (PRI) for ready-to-use microbiology media, including MediaBox™





Microbiology INTERNATIONAL

Experts in Laboratory Instrumentation & Media Solutions
An ISO 9001:2015 Certified Company

800-EZMICRO (396-4276) x 123
• www.800EZMICRO.COM

is just another variant of the overfitting problem that plagued GFT. When this single mistake is fixed, Yelp reviews have less predictive power than inspection history. And a re-analysis of the full data set of 13,000 observations reveals that Yelp information does not add predictive power to a model based solely on information already available to health departments (e.g., inspection history, ZIP code). Notwithstanding the generous coverage of how Yelp can solve food safety challenges, social media actually adds little.

These findings are sobering. The replication failure is part of a broader replication crisis that is sweeping the biomedical sciences and the social sciences, as well as machine learning. The *Science* team struggled to precisely understand GFT's failure, as search terms were never published. Researcher discretion pervades the 1,700 lines of code, each one of which may—adventurously or inadvertently—change the findings. And a recent review of top machine learning conferences found that source code was available for only 6 percent of papers, making replication extremely difficult.⁴² News and journal coverage rarely exists to correct these kinds of errors, so that only splashy promises of silver bullets reach broader audiences.

Considerations for Cautious Use of AI

Much excitement exists around the potential for AI to reshape and improve food safety. Given the informational challenges of foodborne disease surveillance, we agree that the accumulation of more extensive data, the development of state-of-the-art machine learning methods, and technology to deploy and scale such methods at each step of the food production process are tremendously promising. Indeed, much of our research involves the development of machine learning.^{43,44}

Yet, we also sound a note of caution. The hype around AI should not drown out deliberate and responsible consideration of what AI can and cannot do for food safety. In other domains, hasty and ill-considered adoption has contributed to technology backlash, such as the calls for banning and regulating facial recognition technology.

We offer several recommendations for parties considering the adoption of AI. First, AI cannot substitute for conventional principles of data analysis and causal inference in public health and epidemiology. When data sources are not representative of the population, systems developed with such data may exacerbate population differences. Knowing those differences may also allow for adjustments to account for lack of representativeness.

Second, it is critical that domain experts and technologists work together to pilot, assess, and develop new solutions. Such exchange involves extensive communication and collaboration for technologists to understand the health inspection process and for domain experts to understand what machine learning can and cannot achieve. Rather than “creating solutions...looking for problems,” such collaborations can identify what problems can meaningfully be addressed by AI. (We speak from experience: Our lab's collaboration with Public Health—Seattle & King County has changed the research questions we ask and led to exceptionally fruitful exchange over the years.)

Third, testing performance in application is critical. GFT used only retrospective data, and a better test would have been a prospective test. Google's health inspection project, on the other hand, got this part right by evaluating violation scores for restaurants identified by the risk algorithm against a random sample of restaurants.

Last, while the literature on algorithmic bias is rapidly evolving, any machine learning system should be examined for its potential to create bias relative to the

status quo. While assessing bias can be a complex endeavor, the comparison is easier when the baseline for food safety inspections is clear: Health departments typically require full-service restaurants to be inspected at the same frequency.

In sum, while there is much promise in AI, the future depends on deeper engagement among industry, policy makers, public interest groups, and

“Much excitement exists around the potential for AI to reshape and improve food safety.”

technologists. An algorithm for food safety isn't something to be procured one time off the shelf. It is developed based on understanding the challenges of the domain of food safety, adapted to meet institutional constraints and dynamic changes, and evaluated (and re-evaluated) based on its ability to serve goals on a continuing basis.

While most of the hype has surrounded the use of social media data, there are many unexplored opportunities with more comprehensive existing data, such as source attribution using PulseNet data,⁴⁵ learning from combinations of supply chain inspection and audit data, environmental health enforcement using satellite imagery,^{44,46} and risk targeting based on existing inspection histories.⁴⁷ Using these data and techniques responsibly and reflectively might ultimately make good on the Brilliant idea for early detection and early response. ■

The authors thank Tim Lytton and Coby Simler for helpful comments.

Kristen M. Altenburger, A.M., is a research fellow (non-resident) with the Regulation, Evaluation, and Governance Lab at Stanford Law School and is finishing her Ph.D. in computational social science in the Management Science & Engineering Department at Stanford. **Daniel E. Ho, Ph.D.**, is the William Benjamin Scott and Luna M. Scott Professor of Law, professor of political science, senior fellow at the Stanford Institute for Economic Policy Research, and director of the

Regulation, Evaluation, and Governance Lab (RegLab) at Stanford University.

References

1. brandingforresults.com/larry-brilliant/.
2. Ginsberg, J, et al. 2009. "Detecting Influenza Epidemics Using Search Engine Query Data." *Nature* 457 (7232):1012–14.
3. www.forbes.com/sites/gregsatell/2014/02/19/why-experts-always-seem-to-get-it-wrong/#57b2ae6b3a36.
4. www.nytimes.com/2008/11/12/technology/internet/12flu.htm.
5. www.nytimes.com/2011/01/30/business/30charity.htm.
6. science.sciencemag.org/content/343/6176/1203.
7. www.forbes.com/sites/stevensalzberg/2014/03/23/why-google-flu-is-a-failure/#57a6060d5535.
8. callingbullshit.org/videos.htm.
9. www.foodqualityandsafety.com/article/artificial-intelligence-a-real-opportunity-in-food-industry/.
10. www.fda.gov/news-events/press-announcements/statement-acting-fda-commissioner-ned-sharpless-md-and-deputy-commissioner-frank-viannas-steps-usher.
11. Arendt, S, et al. 2013. "Reporting of Foodborne Illness by U.S. Consumers and Healthcare Professionals." *Int J Environ Res Public Health* 1(8):3684–3714, doi:10.3390/ijerph10083684.
12. Lytton, TD. *Outbreak: Foodborne Illness and the Struggle for Food Safety* (Chicago: University of Chicago Press, 2019).
13. Scallan, E, et al. 2011. "Foodborne Illness Acquired in the United States—Major Pathogens." *Emerg Infect Dis* 17(1):7–15.
14. Oldroyd, RA, MA Morris, and M Birkin. 2018. "Identifying Methods for Monitoring Foodborne Illness: Review of Existing Public Health Surveillance Techniques." *JMIR Public Health Surveill* 4(2):e57.
15. www.aclweb.org/anthology/D13-1150.
16. Schomberg, JP, et al. 2016. "Supplementing Public Health Inspection via Social Media." *PLOS ONE* 11(3):e0152117.
17. Effland, T, et al. 2018. "Discovering Foodborne Illness in Online Restaurant Reviews." *J Am Med Inform Assoc* 25(12):1586–1592.
18. doi.org/10.1093/jamiaopen/ooz030.
19. Sadilek, A, et al. 2017. "Deploying nEmesis: Preventing Foodborne Illness by Data Mining Social Media." *AI Magazine* 38(1):37.
20. Sadilek, A, et al. 2018. "Machine-Learned Epidemiology: Real-Time Detection of Foodborne Illness at Scale." *Npj Dig Med* 1(1):36.
21. Quade, P and EO Nsoesie. 2017. "A Platform for Crowdsourced Foodborne Illness Surveillance: Description of Users and Reports." *JMIR Public Health Surveill* 3(3):e42.
22. www.theatlantic.com/technology/archive/2016/01/predictive-policing-food-poisoning/423126/.
23. www.popsci.com/technology/article/2013-08/which-restaurants-will-give-you-food-poisoning-ask-twitter/.
24. www.govtech.com/dc/articles/What-Can-Boston-Restaurant-Inspectors-Learn-from-Yelp-Reviews.html.
25. gizmodo.com/researchers-say-this-ai-can-spot-unsafe-food-on-amazon-1836997785.
26. www.sciencedaily.com/releases/2018/11/181106073236.htm.
27. www.reviewtrackers.com/how-to-handle-food-complaints/.
28. Ho, DE. 2017. "Does Peer Review Work: An Experiment of Experimentalism." *Stanford Law Review* 69:1–119.
29. Irwin, K, et al. 1989. "Results of Routine Restaurant Inspections Can Predict Outbreaks of Foodborne Illness: The Seattle-King County Experience." *Am J Public Health* 79(5):586–590.
30. Cruz, MA, DJ Katz, and JA Suarez. 2001. "An Assessment of the Ability of Routine Restaurant Inspections to Predict Food-Borne Outbreaks in Miami-Dade County, Florida." *Am J Public Health* 91(5):821.
31. Jones, TF, et al. 2004. "Restaurant Inspection Scores and Foodborne Disease." *Emerg Infect Dis* 10(4):688–692.
32. Sweeney, L. 2013. "Discrimination in Online Ad Delivery." *Queue*.
33. proceedings.mlr.press/v81/buolamwini18a.html.
34. Angwin, J, et al. 2016. "Machine Bias." *Pro-Publica* May 23.
35. www.wired.com/story/excerpt-from-automating-inequality/.
36. www.aaai.org/ocs/index.php/ICWSM/ICWSM14/paper/view/8062.
37. Ho, DE. 2017. "Equity in the Bureaucracy Symposium Issue: Establishing Equity in Our Food System." *UC Irvine Law Review*, 2:401–452.
38. slate.com/human-interest/2014/06/ethnic-restaurants-and-food-poisoning-the-subtle-racism-of-saying-chinese-food-caused-your-stomachache.html.
39. Altenburger, KM and DE Ho. 2018. "When Algorithms Import Private Bias into Public Enforcement: The Promise and Limitations of Statistical Debiasing Solutions." *J Institution Theoret Econ* 175(1):98.
40. Zukin, S, S Lindeman, and L Hurson. 2015. "The Omnivore's Neighborhood? Online Restaurant Reviews, Race, and Gentrification." *J Consum Cult* 17(3):459–479.
41. Altenburger, KM and DE Ho. 2019. "Is Yelp Actually Cleaning Up the Restaurant Industry? A Re-Analysis on the Relative Usefulness of Consumer Reviews." In *The World Wide Web Conference*, 2543–2550. WWW '19. New York: ACM.
42. Gundersen, OE and S Kjensmo. 2018. "State of the Art: Reproducibility in Artificial Intelligence." In *Thirty-Second AAAI Conference on Artificial Intelligence*.
43. Altenburger, KM and J Ugander. 2018. "Monophily in Social Networks Introduces Similarity among Friends-of-Friends." *Nature Hum Behav* 2(4):284–90.
44. Handan-Nader, C and DE Ho. 2019. "Deep Learning to Map Concentrated Animal Feeding Operations." *Nature Sustainability* 2(4):298–306.
45. Zhang, S, et al. 2019. "Zoonotic Source Attribution of *Salmonella enterica* Serotype Typhimurium Using Genomic Surveillance Data, United States." *Emerg Infect Dis* 25(1):82.
46. Purdy, R. 2010. "Using Earth Observation Technologies for Better Regulatory Compliance and Enforcement of Environmental Laws." *J Environ Law* 22(1):59–87.
47. Hino, M, E Benami, and N Brooks. 2018. "Machine Learning for Environmental Monitoring." *Nature Sustainability* 1(10):583–588.