IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF MICHIGAN
SOUTHERN DIVISION

CONCERNED PASTORS FOR SOCIAL ACTION, et al.,

Plaintiffs,

v.

NICK A. KHOURI, et al.,

Defendants.
_______________________________________/

Case No. 16-10277

Hon. David M. Lawson

Mag. J. Stephanie Dawkins Davis

DECLARATION OF ERIC M. SCHWARTZ, Ph.D.

I, Eric M. Schwartz, declare as follows:

1. I am an Assistant Professor of Marketing and Arnold M. and Linda T. Jacob Faculty Fellow at the Ross School of Business of the University of Michigan. I hold a Ph.D. in Marketing from the Wharton School of the University of Pennsylvania, and a B.A. in Mathematics and Hispanic Studies from the University of Pennsylvania.

2. I am an applied statistician. My research interests include predicting customer behavior using statistical machine learning models, Bayesian data analysis and econometrics, adaptive sampling, and reinforcement learning.

3. My *curriculum vitae* is attached as Exhibit A to this Declaration.

4. In 2016, I built a statistical model to predict the locations of lead and galvanized steel water service lines in Flint, Michigan, in collaboration with my

colleague, Jacob D. Abernethy, currently an Assistant Professor of Computer Science at the Georgia Institute of Technology's School of Computing, who was at the time an Assistant Professor of Computer Science at University of Michigan. We built the model with the assistance of several undergraduate and graduate students in the Michigan Data Science Team of the University of Michigan.

5. In 2016 and 2017, Dr. Abernethy and I worked in collaboration with the Flint Action and Sustainability Team (FAST) Start Program, under the direction of Brig. Gen. (ret.) Michael McDaniel. During that time, we regularly provided the City with lists of addresses in Flint that our model predicted were most likely to have lead or galvanized steel service lines. It is my understanding that the FAST Start Program used our predictions to inform its decision-making about where to conduct hydro-excavations[1] and service line replacements in 2016 and 2017.

6. In 2017, I wrote a paper with Dr. Abernethy and our team of (now former) University Michigan students that describes our work developing the model for Flint. A version of this paper, titled "ActiveRemediation: The Search for Lead Pipes in Flint, Michigan," was published on June 10, 2018, in the

[1] A hydro-excavation involves "a high-pressure jet of water used to loosen soil and a powerful vacuum hose that sucks the loosened material into a holding tank," which allows workers to dig a small hole and observe the service line material underground. *See* Exhibit B at 5.

Proceedings of the 24th ACM SIGKDD Conference on Knowledge Discovery and Data, a peer-reviewed data science and machine learning publication outlet, and is attached as Exhibit B to this Declaration.[2]

Communications with the City of Flint in 2016-2017 and Original Model (Using 2016 Data)

7. In March 2016, in collaboration with University of Michigan-Flint, Dr. Abernethy and I started working on a project to develop a mobile app and website that would allow Flint residents to determine the risk of lead-contaminated water in their homes.

8. In June, after Dr. Abernethy and I learned that no one knew precisely how many homes in Flint had lead or galvanized steel service lines, or where those service lines were located, we started another project to help Flint residents learn the materials of their service lines, and to help the FAST Start Program find and remove all of Flint's lead and galvanized steel service lines.

9. As part of this project, we received data from the City on 55,893 homes in Flint, including addresses, parcel identifiers, property values, home age,

[2] An earlier version of the paper, titled "On the Search for Lead Pipes in Flint: An Algorithmic Approach to Remediation of Water Utilities," was posted online in November 2017, *see* https://goo.gl/qmkSLU, and http://faculty.chicagobooth.edu/workshops/marketing/pdf/Winter%202018/Schwartz_etal_flint_sl_main_20180216.pdf. The 2018 version of the paper is also available for free online at https://arxiv.org/abs/1806.10692.

and information concerning what the City's records said, if anything, about each home's service line material.

10. Dr. Abernethy and I also considered a digitized map of the City's service line records created by Professor Martin Kaufman and a group of students from the Geographic Information Systems (GIS) Center at the University of Michigan-Flint. It is my understanding that in early 2016, Professor Kaufman received from the City almost 200 pages of hand-drawn maps last updated in 1982 that had annotations indicating the service line material at certain homes. Professor Kaufman's team created the digitized map by looking at each of the hand-drawn maps, finding the annotations indicating service line material, and typing those annotations into a spreadsheet associated with the address. It is my understanding that the resulting map was filed in this case as Exhibit 14 to the June 21, 2018, Declaration of Dimple Chaudhary (ECF No. 166-2).[3]

11. In June 2016, Dr. Abernethy and I met with General McDaniel, who at the time managed the FAST Start Program. Dr. Abernethy and I agreed, on a pro

---

[3] Because this map is based on a small subset of historical data, the map's estimate of total lead service lines in Flint (4,000 to 8,000) grossly underestimates the total number of likely lead and galvanized steel service lines in the City. *See infra* Paragraph 15. Indeed, Professor Kaufman acknowledged that thousands of parcels in the hand-drawn maps had no information about service line composition. *See* UM-Flint News, *New UM-Flint Research Shows Location of Lead Pipes in Flint* (Feb. 22, 2016), https://news.umflint.edu/2016/02/22/new-um-flint-research-shows-location-of-lead-pipes-in-flint/.

bono basis, to build a predictive model to help identify the locations of lead and galvanized steel service lines in Flint. Our goals were to help the City estimate the total number and locations of lead and galvanized steel service lines in Flint, thereby helping the City minimize unnecessary excavation work and maximize replacement of lead and galvanized steel service lines in Flint.

12. We created our predictive model during the summer of 2016. The model works by generating a probability (between 0 and 1) that a lead or galvanized steel service line is present at each home in Flint based on information about that home, including its address, age, assessed value, the zoning of the property, the type of fire hydrant closest to the home, and old City records of the home's service line composition (if available). A list of this data appears in our paper (Exhibit B at 3, tbl.1 & Section 2.1). The model was designed to incorporate new information, including the results of excavations, on an ongoing basis, to improve the accuracy of its predictions and guide decision-making at future stages.

13. Beginning in August 2016, Dr. Abernethy and I started using our model to provide General McDaniel and the FAST Start Program with lists of addresses that the model predicted had the highest chance of having lead or galvanized steel service lines. It is my understanding that the City used these lists to make decisions about where to conduct traditional excavations during the initial phases of the FAST Start Program, which took place between September and

December 2016. Approximately 90% of the traditional excavations the City conducted during those months uncovered lead or galvanized steel service lines.

14. After we received data from the City's traditional excavations, service line replacements, and hydro-excavations through the fall of 2016, Dr. Abernethy and I updated our predictions and were able to make more confident statements about the locations and total number of lead and galvanized steel lines in the City.

15. Together with the City and the FAST Start Program, we released a short report on November 1, 2016, that estimated that between 20,600 and 37,100 homes in Flint were served by lead or galvanized steel lines out of the 55,893 total addresses in Flint. This report is attached as Exhibit C. At that time, we could not provide a more precise estimate due to the limited number of physically verified service lines, as well as uncertainty about the accuracy of the City's historical records. We also did not have reliable information about which homes in the City had active water accounts, and so our estimate was based on the total number of homes in the City, both with and without active water accounts (55,893).

16. We continued to assist General McDaniel and the FAST Start Program throughout most of 2017. For instance, we created interactive maps based on our model and excavation data that we received from General McDaniel. The maps provided predictions of the locations of lead and galvanized steel service lines throughout the City. We updated these maps when we received new

excavation data from the City and shared them with General McDaniel on a regular basis. We also provided General McDaniel with spreadsheets and maps of recommended homes for the City to excavate, based on our model. It is my understanding that the City used the information we provided to inform its decisions about where to conduct hydro-excavations during Phases III and IV of the FAST Start Program in 2017.

17. As part of our involvement with the FAST Start Program, in May 2017, the City provided us with scanned images of approximately 140,000 handwritten index cards containing information on the historical work records from the City's Water Department. Many (but not all) of these index cards contained information about service line composition at specific homes. Dr. Abernethy and I sent those index-card images to be digitally transcribed on a pro bono basis by a company called Captricity. In September 2017, Captricity provided us and the City with spreadsheets reflecting the digitized data for approximately 50,000 of the index cards.

Communications with the City in 2018 and Updated Model Results (Using 2016-2017 Data)

18. Our regular communications with the City largely ended when the City hired AECOM to oversee the FAST Start Program.

19. On December 4, 2017, Dr. Abernethy and I spoke with officials from

AECOM, the FAST Start Program, and MDEQ to explain our collaboration with the FAST Start Program and make plans for 2018. I had several follow-up communications with AECOM throughout that month. We provided AECOM staff with detailed information about our predictive model, including a link to our paper (*see supra* note 2) and the estimated cost savings from using our model and proposed approach. We also provided AECOM with an interactive map of up-to-date data (as of December 2017) on the composition of service lines discovered through hydro-excavations and service line replacements completed during 2016 and 2017.

20. We had expected to continue to collaborate with AECOM on the use of our model to help the City locate individual homes and groups of homes in neighborhoods likely to have lead or galvanized steel service lines. However, since our initial communications in December 2017, AECOM has largely not responded to our calls and emails offering assistance.

21. On January 2, 2018, I emailed Alan Wong, the FAST Start Program Manager, to schedule a phone call to discuss how we could continue to use our model to assist the FAST Start Program. Mr. Wong did not respond. I sent additional emails to Mr. Wong on January 3, February 12, March 1, and May 1. In the May 1 email, I stated that Dr. Abernethy and I would like to discuss strategies for selecting homes for hydro-excavation during Phase V of the FAST Start

Program. Mr. Wong did not respond to any of these emails.

22. On February 23, I wrote an email to Steven Branch at the City Administrator's Office, and the same day, Dr. Abernethy and I had a telephone conversation with Mr. Branch discussing our past involvement with the FAST Start Program and how to best continue that work with AECOM, including by providing them with our predictive model and updating the model's predictions.

23. On March 1 and March 9, I wrote emails to Mr. Branch to follow up on our discussion about the City's potential use of our model during Phase V of the FAST Start Program. Mr. Branch did not respond to these emails.

24. On May 7, 2018, Dr. Abernethy and I emailed Mr. Wong, and later that day, we had a phone conservation with Mr. Wong about AECOM and the City's potential use of our predictive model for Phase V of the FAST Start Program. Mr. Wong told us that he would put us in touch with Constantine Kontos, an AECOM staff member who specializes in data analysis and works on the FAST Start Program.

25. Dr. Abernethy and I connected with Mr. Kontos and, on May 24, 2018, we provided him with the results of our updated model that incorporated data from the excavations conducted by the City as of December 2017. This included a list of all 55,000 addresses in Flint, with the probability of finding a lead or galvanized steel service line at each address. He told us that he was not sure how

AECOM was going to use these updated results but stated that he would follow up with us regarding potential collaboration with AECOM. We did not receive a follow-up email from Mr. Kontos.

26. On May 31, 2018, I emailed Mr. Kontos and Mr. Wong requesting updated information on the hydro-excavations conducted so far in 2018. I informed AECOM that we could incorporate these updated results into our model to make the predictions more accurate than the predictions that we had previously provided. Mr. Kontos did not respond to my email. A few weeks later, on June 18, 2018, I emailed Mr. Kontos again and asked whether AECOM had found our updated model results useful and how AECOM had used them.

27. Also in the spring of 2018, we were coordinating phone calls between Capricity (the company that had volunteered to digitize the more than 100,000 index-card records), AECOM, and the City to discuss how AECOM could use Captricity's work. As of February 2018, Captricity had processed approximately 80,000 of the index-card records and made them available to the FAST Start Program. By June 22, Captricity had completed the digitization work for approximately 95,000 index cards and had sent all processed and transcribed records to us and AECOM's FAST Start staff. I do not know whether or how the City is using the data from the index cards it received from Captricity.

28. On July 10, 2018, Mr. Kontos informed Dr. Abernethy and me that he

thought it would be helpful to update our predictive model with excavation results from 2018 but stated: "[T]hat is a conversation to be had with the City in regards to sharing information and direction on future planning as I am currently sharing my opinion." I sent a follow-up email to Mr. Kontos seeking clarification, but I did not receive a response.

29. After I did not hear back from Mr. Kontos, I contacted the City and spoke with the City's press officer, Candice Mushatt, who requested that I email her information regarding our communications with AECOM. I emailed this information to Ms. Mushatt on July 13, 2018, and left her a voicemail on July 19, 2018. Ms. Mushatt did not respond to my email or voicemail.

30. Neither Dr. Abernethy nor I have had any further communications with AECOM or the City regarding our predictive model. I do not know if or how our model is currently being used by AECOM as part of the FAST Start Program. Because the City has not provided us with any updated data on the results of excavations conducted in 2018, we stopped updating our model in December 2017.

31. In August 2018, I learned from the press that AECOM had used our model results to create a map (based on excavation results as of December 2017, which we sent to AECOM in May 2018), which Mr. Wong attached to his declaration in support of the City's July 12, 2018 filing in this case (ECF No. 172-3).

32. Mr. Wong states in his declaration that he is "not aware of any algorithm that has been developed to reliably predict the composition of buried infrastructure that has been installed over 5 to 10 decades." However, as we explained to Mr. Wong in December 2017 and May 2018, our model was developed exactly for that purpose, and I had provided him and other AECOM staff with our paper via email, detailing how the predictive model worked and how it could be used for making the City's excavations more cost effective.

33. As described in our paper (Exhibit B), our model accurately estimates how likely it is for a given home in Flint to have a lead or galvanized steel service line. Our model has an Area Under the Receiver Operating Characteristic (AUROC) score of nearly 0.94, which means that, if you considered a home with a lead or galvanized steel service and a randomly chosen home and asked our model to tell you which one is which, our model would be correct 94% of the time. In contrast, a random guess would be correct only 50% of the time.

34. We found that the most informative home features for predicting the presence of a lead or galvanized steel service line are its age, value, and location. Homes that were built before about 1950 and those that are lower in value are more likely to contain lead service lines. The historical City records were also a useful predictor when combined with other factors. Information concerning the nearest fire-hydrant type was not a useful predictor of the presence of lead and galvanized

steel service lines. Further, while homes likely to contain lead and galvanized steel service lines are geographically scattered across the City, there were larger clusters of high-probability homes in specific neighborhoods.

Communications with Plaintiffs' Counsel

35. On August 29, 2018, Dr. Abernethy and I spoke with Plaintiffs' counsel about the filings recently submitted in this case and, in particular, the map attached as Exhibit 5 to Mr. Wong's declaration, which states that it depicts the results of our model.

36. Plaintiffs' counsel asked whether we could run our model with the updated results from the City's 2018 excavations. We stated we had asked the City for this updated information but had not received it. We told Plaintiffs' counsel that we could update the model predictions if we had that updated data.

37. Plaintiffs' counsel also informed us that the Settlement Agreement in this case requires the City to conduct excavations at homes that had active water accounts as of the date the Agreement was executed, as opposed to at all homes in Flint. They requested that, when running our model, we consider only homes with active water accounts.

38. Plaintiffs' counsel provided us with the documents listed in Exhibit D to this Declaration, which included information on the results of the City's excavations so far in 2018, and a list of active water accounts they had received

from the City.

Updated Model Results (Using 2016-2018 Data)

39. Dr. Abernethy and I ran our model with the new excavation data we received from Plaintiffs to estimate the total number of eligible homes[4] in Flint with lead and galvanized steel service lines as of March 28, 2017.

Projected total number of eligible homes in Flint with lead and galvanized steel service lines as of March 28, 2017

40. As discussed above in Paragraph 12, our predictive model works by assigning a probability (between 0 and 1) of having a lead or galvanized steel service line to every home in Flint. Using these probabilities, I applied two different statistical methods to project the likely number of lead and galvanized steel service lines at eligible homes in Flint.

41. **Method 1.** Under the first method, I calculated the projected number of lead and galvanized steel service lines at eligible homes in Flint by adding up the individual probabilities assigned to each eligible home in the City that had not been excavated as of August 15, 2018, adding this sum to the 6779 lead and

[4] I understand that the term "eligible household," for purposes of the Settlement Agreement, means a household that had an active water account as of March 28, 2017. *See* Settlement Agmt. ¶ 11, ECF No. 147-1. The earliest list of homes with active water accounts Plaintiffs received from the City was a list of such accounts as of July 6, 2017. Accordingly, I have used that list as the list of eligible households for purposes of this Declaration.

galvanized steel service lines the City had previously discovered using excavations, and then subtracting the 752 lead and galvanized steel service lines that the City replaced as of March 28, 2017.[5] Using this method, our model predicts that, as of March 28, 2017, there were **10,836** lead and galvanized steel service lines at eligible homes in Flint. Our model predicts that, as of August 15, 2018, there were **4809** remaining, unexcavated lead and galvanized steel service lines at eligible homes in Flint. The 4809 figure reflects the model's total predicted number of lead and galvanized steel service lines at eligible households as of March 28, 2017 (10,836), minus the number of service line replacements completed between that date and the "City of Flint's Paragraph 30 Evaluation" (Paragraph 30 Report) that was submitted by the City in February 2018 (6027).

42. I also calculated a 95% confidence interval of 10,716 to 10,969 for the predicted total number of lead and galvanized steel service lines in Flint eligible households as of March 28, 2017 (10,836). A confidence interval is an estimated

[5] Based on the 2016-2018 excavation data I received from the City and Plaintiffs, our data shows that the City replaced 6779 lead and galvanized steel service lines as of August 15, 2018. However, this number is slightly different from the 6839 service line replacements reported by the City in Mr. Wong's August 3, 2018, declaration. Aug. 3, 2018 Wong Decl. ¶ 9, ECF No. 181-1. This difference exists largely because our model requires a "parcel identification number" for each home, but not all of the addresses in the data provided to us have corresponding parcel identifiers. The 2016-2018 excavation data we received included 30,957 addresses. We were able to match 28,414 of these addresses to parcel identification numbers, but we were unable to incorporate the remaining 2543 addresses (about 8% of the total) into our analysis.

range of values likely to contain the true value for a given parameter. Thus, by reporting a 95% confidence interval around the predicted number of lead and galvanized steel service lines at eligible homes, I am stating that, after rerunning the model 100 times on random variations of the same dataset, my calculations indicate that in 95 out of the 100 variations, the predicted number of lead and galvanized steel service lines lies within this range.

43. **Method 2.** Under the second method, I calculated the projected number of lead and galvanized steel service lines at eligible homes by setting a specific probability "threshold" for considering a particular home's likelihood of having a lead or galvanized steel service line. Under this method, I added up the total number of eligible homes with yet-to-be-excavated service lines as of August 15, 2018, with a probability of having a lead or galvanized steel service line that equals or exceeds that threshold. I added this sum to the 6779 lead and galvanized steel service lines already discovered by the City as of August 15, 2018.

44. To project the number of lead and galvanized steel service lines at eligible homes in Flint, I first set the probability threshold at 0.3. In other words, I added up all the eligible homes that our model assigned at least a 30% probability of having a lead or galvanized steel service line. I then set the probability threshold at 0.1 and added up all the eligible homes that our model assigned at least a 10%

probability of having a lead or galvanized steel service line.[6]

45. Using the 0.3 threshold, our model predicts that, as of March 28, 2017, there were **10,991** lead and galvanized steel service lines at eligible homes in Flint, and as of August 15, 2018, there were **4964** remaining unexcavated lead and galvanized steel service lines at eligible homes in Flint.

46. Using the 0.1 threshold, our model predicts that, as of March 28, 2017, there were **12,146** lead and galvanized steel service lines at eligible homes in Flint, and as of August 15, 2018, there were **6119** remaining unexcavated lead and galvanized steel service lines at eligible homes in Flint.

47. To test the accuracy and reliability of the predictions described in Paragraphs 41 to 46 above, I conducted what is known as model validation. I split the data in two parts: I used a randomly selected subset of 75% of all homes already excavated to predict the materials of the other 25% of homes already excavated, and then I compared the predictions to the actual data. I repeated this process 100 times using different random splits. For the 25% of data not used to generate predictions, the model has an AUROC score of 0.94, on average across 100 different splits. This indicates that the model is accurate at predicting the

[6] We used thresholds of 0.3 and 0.1 to make our projections health-protective. If Flint's goal is to uncover and replace all of the City's hazardous service lines, then if there is a 30%—or even a 10%—chance that the service line is hazardous, it may be appropriate from a policy perspective to excavate that service line to determine whether it needs to be replaced.

materials of service lines in Flint.

Projected costs of completing 18,000 excavations and replacing all lead and galvanized steel service lines

48. As discussed in Paragraph 41, using Method 1 (which adds the individual probabilities for all homes), our model predicts that, as of February 2018, there were 4809 remaining lead and galvanized steel service lines at eligible homes in Flint.

49. Based on the City's Paragraph 30 Report, I understand that the City had completed 8843 excavations as of February 2018. *See* Paragraph 30 Report 2, ECF No. 172-4. Assuming the City is required to complete 18,000 total excavations under the Settlement Agreement, then as of February 2018, I understand that it was required to conduct an additional 9157 excavations (18,000 minus 8843 = 9157),[7] and complete service line replacements at those homes at which it uncovered a lead or galvanized steel service line. Settlement Agmt. ¶¶ 9-10. I calculated the costs of completing this remaining work, under three scenarios based on different cost-related assumptions.

50. For each scenario described below, I assumed that the City will prioritize its excavations at those addresses most likely to have lead or galvanized

[7] The Paragraph 30 Report states that the City must complete 9173 excavations to reach a total of 18,000. Paragraph 30 Report 3. This appears to be based on a subtraction error (18,000 – 8843 = 9157, *not* 9173).

steel service lines. I also assumed that, through the remaining 9157 excavations, the City will uncover and replace all remaining 4809 lead and galvanized steel service lines.

51. **Scenario 1.** In Scenario 1, I used the following cost projections provided in the City's Paragraph 30 Report:[8]

- Cost of one service line replacement: $4985
- Cost of one hydro-excavation: $285
- Cost of one traditional excavation: $2605
- Cost of one site restoration[9]: $1725
- Total cost of 2018-2019 administrative and management: $9,297,508

I also assumed, based on the cost calculations in the Paragraph 30 Report, that at an address where contractors conduct a traditional excavation *and* a service line replacement, there will be no cost efficiencies. In other words, I assumed that if a traditional excavation and a service line replacement are both required at a particular address, the cost will be $7590 ($4985 + $2605). Finally, I assumed in Scenario 1 that the City will use only traditional excavations, which I understand to be the City's current practice. *See* Oct. 1, 2018, Tallman Decl. Exs. D, E.

52. Based on the assumptions described in Paragraphs 49 through 51, I

[8] Where the Report projected different costs for a particular item in 2018 and 2019, I used the average of the two projections.

[9] Site restoration involves restoring the area disrupted by an excavation or service line replacement, including filling the hole dug during an excavation or service line replacement with soil, cement, or other material.

calculated that **in Scenario 1** the City will need **$72,920,183** to complete the remaining 9157 excavations and 4809 service line replacements. Adding this total to the $37.25 million the City spent of the Settlement Agreement funding prior to February 2018, *see* Paragraph 30 Report 3, yields a total cost of **$110,170,183** for all 18,000 excavations and service line replacements.

53. **Scenario 2.** I understand that in September 2018, the City provided Plaintiffs with revised, updated cost figures for its excavation and service line replacement work in 2018. *See* Tallman Decl. Exs. A, B, C. In Scenario 2, I calculated the costs of completing the remaining 9157 excavations and 4809 service line replacements using these revised cost figures. I assumed that:

- Cost of one traditional excavation and service line replacement: $4197.86
- Cost of one hydro-excavation: $285
- Cost of one traditional excavation: $1788.37
- Cost of one site restoration: $1725
- Total cost of 2018-2019 administrative and management: $9,297,508

In addition, based on the information the City provided to Plaintiffs, I assumed in Scenario 2 that there are cost efficiencies when the City conducts both a traditional excavation and a service line replacement at a particular address. That is, I assumed in Scenario 2 that, the cost of completing both a traditional excavation and service line replacement at a particular address is $4197.86.

54. Finally, I assumed in Scenario 2 that the City will use only traditional

excavations, which is what I understand the City's current practice to be. *See* Tallman Decl. Exs. D, E.

55. Based on the assumptions described in Paragraphs 49–50 and 53–54, I calculated that in **Scenario 2** the City will need **$53,056,675** to complete the remaining 9157 excavations and 4809 service line replacements. Adding this total to the $37.25 million the City spent of the Settlement Agreement funding as of February 2018, *see* Paragraph 30 Report 3, yields a total cost of **$90,306,675**.

56. **Scenario 3.** Although the City is currently using only traditional excavations, it could instead investigate the service line material at a particular address using a hydro-excavation. This was the method used by the City prior to June 2018. Hydro-excavations are significantly less expensive than traditional excavations.

57. To better understand the cost savings associated with using hydro-excavations, in Scenario 3 I calculated the cost of completing the remaining 9157 excavations and 4809 service line replacements using the assumptions in Scenario 2, but with the revised assumption that the City will use hydro-excavations where possible in lieu of a traditional excavation. In Scenario 3, I assumed that the City will conduct hydro-excavations at 81% of the addresses it investigates, and will conduct traditional excavations at the remaining 19% of addresses (the percentages

reported in the City's Paragraph 30 Report).[10]

58. In Scenario 3, based on the assumptions outlined in Paragraphs 49–50 and 53–54, my calculations show that it will cost the City **$42,796,900** to complete the remaining 9157 excavations and 4809 service line replacements. Comparing this figure to the projected cost in Scenario 2 shows that using hydro-excavations is associated with a cost savings of $10,259,775.

Analysis of locations of the City's excavations in 2018 as of August 15

59. I also analyzed the City's excavation data to understand trends in addresses it has selected for excavations in 2018. For each of Flint's nine wards, I calculated:

- the total number of eligible homes that were not yet excavated as of December 31, 2017;
- the predicted number of lead and galvanized steel service lines at eligible homes using our model (only using excavation data I had obtained as of December 2017);
- the predicted "hit rate" (considering only eligible households);
- the number of excavations the City conducted in 2018 as of August 15;
- the number of lead and galvanized steel service lines the City found in 2018 as of August 15; and
- the City's observed "hit rate" for excavations conducted in 2018 as of August 15.

[10] I applied this assumption (81% hydro-excavations, 19% traditional excavations) both to the excavations at the 4809 homes where a lead or galvanized steel service line will be found, and to the excavations at the remaining 4348 homes where a non-lead service line will be found.

A table showing these values for all nine wards is attached to this Declaration as Exhibit E.

60. The results of my analysis show that, at an aggregate level, the City is not conducting excavations at homes that are most likely to have a hazardous (lead or galvanized steel) service line.

61. For example, as of December 2017, there were 20,135 total remaining eligible homes at which the City had not yet conducted excavations. Of those homes, the model predicted that there were 6322[11] remaining lead and galvanized steel service lines (rounded to the nearest whole number). *See* Table 1 below.

62. If the City conducted excavations randomly at the remaining eligible homes in 2018, I would have expected its observed hit rate to be 31.4% (6322 divided by 20,135 = 0.314). However, as of August 15, the City's observed hit rate

[11] This prediction (6322 remaining lead and galvanized steel service lines) was the model's prediction as of the end of 2017, using only 2016 and 2017 data. For purposes of the model's most up-to-date predictions for the total remaining lead and galvanized steel service lines at eligible homes in Flint (see Paragraphs 39 to 41 above), I updated the model to reflect the City's excavation results through August 2018. This affected the overall prediction of the remaining lead and galvanized steel service lines in Flint both because (i) the City had uncovered and replaced nearly 750 hazardous service lines in 2018 as of August 15, and (ii) the City's additional 3774 excavations conducted in 2018 affected the model's predictions (i.e., the model "learned" from the results of the 2018 excavations). Nevertheless, the relative rank ordering of wards by number of expected hazardous service lines at eligible homes not yet excavated using the most up-to-date August 2018 predictions is largely unchanged from the December 2017 predictions.

in 2018 was 19.7%.[12]

63. Considering Ward 5 is illustrative. Ward 5 is the ward that our model predicted to have the largest number of remaining lead and galvanized steel service lines as of December 2017, but it is the ward where the City conducted the fewest number of excavations in 2018. As of August 15, the City had conducted a total of 3774 excavations throughout the City in 2018. Although the model predicted that there were more than 1100 lead and galvanized steel service lines remaining in Ward 5 at the beginning of 2018, the City conducted only 163 excavations in that ward (4.5% of the total excavations conducted in 2018). For those 163 excavations, the City's hit rate was 95.7%.

64. In contrast, based on excavations conducted as of December 2017, our model predicted that Ward 4 was the ward with the fewest remaining lead and galvanized steel service lines. But the city conducted 702 excavations in Ward 4 (18.6% of the total excavations conducted in 2018), making it the second most visited ward in 2018. For those excavations, the City's hit rate was 2.4%.

---

[12] Our model's analysis of the data I received from Plaintiffs shows that the City's 2018 observed hit rate (as of August 15) was 19.7%. This figure differs from the 2018 observed hit rate reported by Mr. Wong in his August 3, 2018 declaration submitted in this case (16.5%). Aug. 3, 2018 Wong Decl. ¶ 10. This discrepancy could be due to the cut-off date in August for including excavations. It could also have resulted from the issue relating to parcel identifiers described in note 5 above.

Table 1: Analysis of 2018 Excavations as of August 15

Ward	Total number of eligible homes not yet excavated, 12/31/2017	Predicted number of hazardous service lines among eligible homes not yet excavated, 12/31/2017	Predicted hit rate among homes not yet excavated, 12/31/2017	Total number of homes excavated, 1/1/2018 - 8/15/2018	Observed number of hazardous service lines identified, 1/1/18 - 8/15/18	Observed hit rate, 1/1/2018 - 8/15/2018
5	1454	1163	80%	163	156	95.7%
4	2433	411	17%	702	17	2.4%
All	20,135	6322	31.4%	3774	742	19.7%

65. If the City does not prioritize its excavations towards those homes where the likelihood of finding a lead or galvanized steel service line is high, it will likely complete 18,000 excavations without having located hundreds of lead and galvanized steel service lines at eligible homes in Flint. If the City continues with its current approach, it will also waste money by conducting thousands of excavations at homes served by full copper service lines.

I declare under penalty of perjury that the foregoing is true and correct to the best of my knowledge and belief.

Dated: October 1, 2018

Eric Schwartz

Eric M. Schwartz, Ph.D.

INDEX OF EXHIBITS

Exhibit	**Description**
A	*Curriculum vitae* of Eric M. Schwartz, Ph.D.
B	Abernethy, J., Chojnacki, A., Farahi, A., Schwartz, E., Webb, J., *ActiveRemediation: The Search for Lead Pipes in Flint, Michigan*, Proceedings of the 24th ACM SIGKDD Conference on Knowledge Discovery and Data (June 10, 2018)
C	Abernethy, J., Schwartz, E., Farahi, A., Webb, J., Anderson, N., Doyle R., Inventory of Service Lines in Flint (Nov. 1, 2016)
D	List of documents provided by Plaintiffs' counsel in August 2018
E	*Table*, Analysis of the City of Flint's Service Line Excavations in 2018 by Ward (as of Aug. 15, 2018)

EXHIBIT A

ERIC M. SCHWARTZ

Ross School of Business
University of Michigan
701 Tappan Street
Office R5472
Ann Arbor, MI 48109-1234

ericmsch@umich.edu
ericmichaelschwartz.com
ssrn.com/author=1192670
http://goo.gl/sAEQ8x
734-936-5042 (office)

Academic employment

Assistant Professor of Marketing, Ross School of Business, University of Michigan, July 2013-present
Arnold M. and Linda T. Jacob Faculty Fellow, July 2018-present

Education

Ph.D. Marketing, Wharton School, University of Pennsylvania, May 2013
B.A. Mathematics and Spanish, College of Arts and Sciences, University of Pennsylvania, May 2008

Research interests

Substantive: adaptive marketing experiments, digital advertising, dynamic pricing, customer acquisition and lifetime value, media consumption, public policy

Methodological: statistical machine learning, Bayesian data analysis and econometrics, adaptive sampling, multi-armed bandit, active learning, reinforcement learning, and dynamic programming

Published or forthcoming papers

Misra, Kanishka, **Eric M. Schwartz**, Jacob Abernethy* (2018), "Dynamic Online Pricing with Limited Information Using Multi-Armed Bandit Experiments," *Marketing Science, forthcoming.* https://goo.gl/mxJ8Rh. *First two authors listed alphabetically.

Schwartz, Eric M., Eric T. Bradlow, Peter S. Fader (2017), "Customer Acquisition via Display Advertising Using Multi-Armed Bandit Experiments," *Marketing Science*, 36 (4), 500-522. https://goo.gl/Ly7Bdv
- *Winner of the* ***2017 John D. C. Little Award*** *for best marketing paper in Marketing Science, Management Science, and other INFORMS journals*

Schwartz, Eric M., Eric T. Bradlow, Peter S. Fader (2014), "Model Selection Using Database Characteristics: Developing a Classification Tree for Longitudinal Incidence Data," *Marketing Science*, 33 (2), 188-205. https://goo.gl/Wsk2yL

Berger, Jonah, and **Eric M. Schwartz** (2011), "What Drives Immediate and Ongoing Word of Mouth?" *Journal of Marketing Research*, 48 (5), 869-880. https://goo.gl/WBt7te

Working papers

Under review or targeted at peer-reviewed publications

Aribarg, Anocha, **Eric M. Schwartz** (2018), "Consumer Responses to Native Advertising," Available on SSRN, https://goo.gl/Fx7nhg.
- Status: Under 2nd round Revision at *Journal of Marketing Research* *Authors listed alphabetically.

Schwartz, Eric M., Jacob Abernethy (2018) "Active Learning for Sequential Household-level Targeting: An Application to Find Lead Pipes in Water Infrastructure"
- Status: Preparing to submit to *Marketing Science.*

Schwartz, Eric M., Kenneth Fairchild, Bryan Orme, Alexander Zaitzeff (2018), "Active Learning for Ranking and Selecting Best Arms: Idea Screening with Bandit MaxDiff"
- Status: Preparing to submit to *Journal of Marketing Research*

Work in progress

Data analysis and/or writing has begun

"Binge Viewing: Ad-Supported Streaming Video" with Puneet Manchanda and Prashant Rajaram
- Status: Analysis in progress.

"Online Advertising to Generate Leads with Randomized Controlled Experiments : Recruiting for the Detroit Police Department," with Michael Braun and Hye Jin Yoon
- Status: Data collection in progress.

"Sequential Allocation for Customer Acquisition" with Liangbin Yang and S. Fader (2015) https://goo.gl/Z6xOC7.
- Status: Manuscript completed.

Peer Reviewed Conference Proceedings

Abernethy, Jacob, Alex Chojacki, Arya Farahi, Eric M. Schwartz, Jared Webb* (2018). ActiveRemediation: The Search for Lead Pipes in Flint, Michigan. *KDD 2018, Proceedings of SIGKDD Conference on Knowledge Discovery and Data Mining, London, England, U.K.* Available on Arxiv http://bit.ly/flint-lead-pipes *Alphabetical order.
- Winner of *Best Student Paper Award, Applied Data Science*, KDD 2018 (One of two awards out of 500 submissions)

Alex Chojnaki, Chengyu Dai, Arya Farahi, Guangsha Shi, Jared Webb, Daniel T. Zhang, Abernethy, Jacob, **Eric Schwartz*** (2017) "A Data Science Approach to Understanding Residential Water Contamination in Flint." *KDD 2017, Proceedings of SIGKDD Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada.* Available on Arxiv, https://goo.gl/EaGfmh.*Students first, then faculty; alphabetical order.

Published conference proceedings

Abernethy, Jacob, Cyrus Anderson, Chengyu Dai, Arya Farahi, Linh Nguyen, Adam Rauh, **Eric Schwartz**, Wenbo Shen, Guangsha Shi, Jonathan Stroud, Xinyu Tan, Jared Webb, Sheng Yang* (2016), "Flint Water Crisis: Data-Driven Risk Assessment Via Residential Water Testing" in proceedings of Bloomberg Conference Data for Good Exchange, NY, NY. https://goo.gl/rBHaIb. *alphabetical order

Abernethy, Jacob, Cyrus Anderson, Alex Chojnacki, Chengyu Dai, John Dryden, **Eric M. Schwartz**, Wenbo Shen, Jonathan Stroud, Laura Wendlandt, Sheng Yang, Daniel Zhang* (2016), "Data Science in Service of Performing Arts: Applying Machine Learning to

Predicting Audience Preferences," in proceedings of Bloomberg Conference Data for Good Exchange, NY, NY. https://goo.gl/GdM0DV. *alphabetical order
- In collaboration with University Musical Society, University of Michigan

Fairchild, Kenneth, Bryan Orme, **Eric M. Schwartz** (2015), "Bandit Adaptive MaxDiff Designs for Huge Number of Items," *Proceedings of 2015 Sawtooth Software Conference*, 105-117. https://goo.gl/5iql87.

Research seminars, invited talks, competitive conferences

Economics of Advertising Workshop, Columbia (July 2018)
Notre Dame Mendoza Marketing (April 2018)
UT Dallas Bass FORMS, Presenter (March 2018)
Emory Goizueta Marketing (March 2018)
UT Dallas Marketing (February 2018)
Chicago Booth Marketing (February 2018)
Michigan, School of Information Seminar (October 2017)
Carnegie Mellon Tepper (September 2017)
Erasmus University, Rotterdam School of Management (April 2017)
UT Dallas Bass FORMS, Discussant (March 2017)
Hosmer-Hall Seminar, Michigan Ross (January 2017)
Management Science Workshop, Chile (January 2017)
Quantitative Marketing and Economics (October 2016)
Michigan, School of Public Health (September 2016)
Michigan, Computer Science Engineering Faculty Seminar (September 2016)
Dartmouth, Tuck Marketing Camp (June 2016)
Texas A&M Marketing (April 2016)
Marketing in Israel 15 Conference (December 2015)
NYU Conference on Big Data and Marketing Analytics (October 2015)
Hosmer-Hall Seminar, Michigan Ross (March 2015)
Temple, Fox Global Center for Big Data and Mobile Analytics (November 2014)
Microsoft Research, Seattle (June 2014)
Cornell Johnson Marketing (February 2014)
Electronic Arts, Redwood City (February 2014)
Google Play, Mountain View (February 2014)
Stanford GSB Marketing (January 2014)
London Business School (November 2012)
INSEAD (November 2012)
University of Michigan (November 2012)
UCLA (October 2012)
NYU (October 2012)
Carnegie Mellon (October 2012)
Northwestern (October 2012)
Yale (October 2012)
Boston University (October 2012)
Emory (September 2012)
University of Pittsburgh (September 2012)
University of Washington (September 2012)
Rotterdam School of Management / Erasmus School of Economics (January 2012)
Tilburg University (January 2012)
Marketing in Israel 11 Conference (December 2011)

Other talks and conference presentations

LEAD Summer Institute, Michigan Ross (July 2018)
Marketing Science Conference, Philadelphia (June 2018)
- Machine Learning in Marketing, Special Track, Co-organizer
American Marketing Association Conference, New Orleans (February 2018)
Marketing Science Conference, Los Angeles (June 2017)
- Machine Learning in Marketing, Special Track, Co-organizer
Customer Analytics in Retail Marketing, Los Angeles (May 2017)
Michigan Student Symposium for Interdisciplinary Statistical Sciences, Keynote (March 2017)
Quicken Loans Data Science, MIDAS Collaboration (January 2017)
Water @ Michigan Conference (January 2017)
Artificial Intelligence Lab, Michigan (October 2016)
Customer Analytics in Retail Marketing, New York (October 2016)
SPARK Machine Learning Workshop, Ann Arbor (October 2016)
Kickstart Computer Science, Ann Arbor (September 2016)
Bloomberg Data for Good Exchange (Presenter and Panelist), New York (September 2016)
Big Data Summer Institute Symposium, Ann Arbor (July 2016)
SPARK Workshop for Startups, Ann Arbor (March 2016)
INFORMS Annual Meeting, Philadelphia (November 2015), session co-organizer
Marketing Science Conference, Baltimore (June 2015)
Sawtooth Software Conference, Orlando (March 2015)
American Marketing Association Conference, San Antonio (February 2015)
Joint Statistical Meetings, Montreal (August 2013), Session Organizer
ART Forum, Chicago (June 2013)
Capital One, Webinar (June 2013)
Marketing Science Conference, Boston (June 2012)
Wharton Customer Analytics Initiative, Webinar with Elea Feit (September 2012)
Marketing Science Conference, Houston (June 2011)
Marketing Science Conference, Cologne (June 2010)
Jay H. Baker Retailing Initiative Board Meeting (November 2009)
Marketing Science Conference, Ann Arbor (June 2009)

Awards, grants, and honors

Marketing Science Institute Young Scholar (2019)
Arnold M. and Linda T. Jacob Faculty Award for Junior Faculty Research, Michigan Ross (2018-19)
KDD Best Student Paper Award, Applied Data Science (2018)
John D. C. Little Award for Best Marketing Paper (2017)
MCubed Grant ($60,000) with Laura Balzano and Al Hero (2016)
20 in Their 20s, Crain's Business Detroit (2016)
Top 25 Reviewer for Marketing Science (2015)
Golden Apple Teaching Award Nominee, University of Michigan (2014, 2015)
MSI Clayton Dissertation Proposal Competition, Honorable Mention (2012)
ISMS Doctoral Dissertation Proposal Competition, Sheth Winner (2012)
AMA-Sheth Foundation Doctoral Consortium, Fellow (2011)
Workshop on Quantitative Marketing and Structural Econometrics, Fellow (2010)
Russell Ackoff Award for Doctoral Student Research, Recipient (2009-12)

Jay H. Baker Retailing Initiative Research Grant, Recipient (2009)
Lauder CIBER Grant, Recipient (2009)
University of Pennsylvania Class of 1939 Fellowship, Recipient (2008-2009)
INFORMS Marketing Science Doctoral Consortium, Fellow (2009,2010,2011,2012)
Wharton Doctoral Fellowship, Recipient (2008-12)
Summa Cum Laude, Dean's List, University of Pennsylvania, GPA: 3.9/4.0 (2004-08)
Benjamin Franklin Scholar, University of Pennsylvania (2004-08)

Media coverage

Data Science for Flint Water Crisis

December 2016 – Wrote report for City of Flint Mayor's Office, reported/cited in Detroit News, MLive, WNEM (TV), Wikipedia

September-October 2016 - Wrote article in The Conversation (with Jacob Abernethy), 8 September 2016 reported/reposted in Scientific American, Business Insider, Associated Press, USA Today, Government & Technology, Detroit Free Press (1, 2, 3, and 4) Huron Daily Tribune, RawStory, Ross Thought in Action, Civics Analytics on Medium,

May 2016 - Announcement for Google funding research and app development, joint with U-M Flint and Engineering, 3 May 2016 reported in Chicago Tribune, Tech Crunch, Gizmodo, The Hill, Detroit Free Press, MLive, Michigan Radio, The University Record, Michigan Engineering News

Quoted in "Is Detroit ready for more soccer?" Detroit Free Press, 5 July 2016.

Profile in 20 in their 20s, *Crain's Business Detroit, 23 May 2016*

Profile in *Dividend, Michigan Ross, Fall 2015*

Model Selection Using Database Characteristics: Developing a Classification Tree for Longitudinal Incidence Data
– Featured in *Ross Thought in Action* (November 2013)
– Featured in *Knowledge@Wharton* (August 2012)

What Drives Immediate and Ongoing Word of Mouth
– Featured in book *Contagious* (2013)
– Featured in *Insights from MSI* (Fall 2010), formerly *MSI Working Paper* [10-105]

Teaching

Instructor

Marketing Management (Ross MBA Core, MKT 300), Fall 2017, 2018 (upcoming)

Marketing Management (Ross BBA, MKT 300), Fall 2013. 2014, 2015, 2016

Student advising

PhD Dissertation

Committee Member
Prashant Rajaram, Marketing, current
Longxiu Liu, Marketing, current
Yanzhe (Murray) Lei, Technology and Operations, 2018
Evgeny Kagan, Technology and Operations, 2018
Eunsoo Kim, Marketing, 2017
Qi George Chen, Technology and Operations, 2017
Guy Benedict Wilkinson, Sports Management, 2017
Joseph Golden, Economics, 2015

Service

To the marketing field

Reviewing activity as ad hoc reviewer

Marketing Science
Journal of Marketing Research
Management Science
Quantitative Marketing and Economics
Journal of Consumer Research
Marketing Letters
Journal of the American Statistical Association
Journal of Applied Econometrics

Conference activity

Session co-organizer, "Machine Learning in Marketing," INFORMS Annual Meeting, Philadelphia (November 2015)
Track co-organizer, "Machine Learning in Marketing," Marketing Science Conference, Los Angeles (June 2017)
Track co-organizer, "Machine Learning in Marketing," Marketing Science Conference, Philadelphia (June 2018)

To Ross School of Business and the University of Michigan

Michigan Data Science Team, Faculty Co-Advisor (2015-16, 2016-17,2017-18)
Ross Marketing 300 Course Coordinator (2015 Fall, 2016 Fall, 2017 Fall)
Ross Marketing Student Awards Committee (2013-14, 2014-15, 2016-17,2017-18)

Ross MBA Data Insights and Analytics Club Co-Advisor (2015-16, 2016-17,2017-18)
Ross Undergraduate Marketing Club Advisor (2014-15, 2015-16, 2016-17)
School of Kinesiology, Sports Marketing Faculty Search Committee (2016 Winter)
Ross Marketing Internal Seminar Coordinator (2014-15)
Marketing Undergraduate Case Competition Coach (2014, 2015)
-First place at national L'Oreal Brandstorm Competition (2014)
Google Online Marketing Challenge Advisor (2015, 2016)

Professional affiliations

American Marketing Association
American Statistical Association
INFORMS Society for Marketing Science
Phi Beta Kappa

Computer and natural languages

Fluent: R, Spanish, SQL, Tidyverse
Proficient: Catalan, Matlab, Python, SAS

EXHIBIT B

ActiveRemediation: The Search for Lead Pipes in Flint, Michigan

Jacob Abernethy
Georgia Institute of Technology &
University of Michigan
prof@gatech.edu

Alex Chojnacki
University of Michigan
thealex@umich.edu

Arya Farahi
University of Michigan
aryaf@umich.edu

Eric Schwartz
University of Michigan
ericmsch@umich.edu

Jared Webb
Brigham Young University
webb@mathematics.byu.edu

ABSTRACT

We detail our ongoing work in Flint, Michigan to detect pipes made of lead and other hazardous metals. After elevated levels of lead were detected in residents' drinking water, followed by an increase in blood lead levels in area children, the state and federal governments directed over $125 million to replace water service lines, the pipes connecting each home to the water system. In the absence of accurate records, and with the high cost of determining buried pipe materials, we put forth a number of predictive and procedural tools to aid in the search and removal of lead infrastructure. Alongside these statistical and machine learning approaches, we describe our interactions with government officials in recommending homes for both inspection and replacement, with a focus on the statistical model that adapts to incoming information. Finally, in light of discussions about increased spending on infrastructure development by the federal government, we explore how our approach generalizes beyond Flint to other municipalities nationwide.

CCS CONCEPTS

• **Information systems** → **Data analytics**; • **Machine learning** → *Applied computing*;

KEYWORDS

Water Infrastructure; Flint Water Crisis; Risk Assessment; Machine Learning; Active Learning; Public Policy

ACM Reference Format:

Jacob Abernethy, Alex Chojnacki, Arya Farahi, Eric Schwartz, and Jared Webb. 2018. ActiveRemediation: The Search for Lead Pipes in Flint, Michigan. In *KDD '18: The 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, August 19–23, 2018, London, United Kingdom.* ACM, New York, NY, USA, Article 4, 10 pages. https://doi.org/10.1145/3219819.3219896

KDD '18, August 19–23, 2018, London, United Kingdom

ACM ISBN 978-1-4503-5552-0/18/08…$15.00
https://doi.org/10.1145/3219819.3219896

1 INTRODUCTION

The story of the Flint Water Crisis is long and has many facets, involving government failures, public health challenges, and social and economic justice. As Flint struggled financially after the 2008 housing crisis, the state of Michigan installed emergency managers to implement several cost saving measures. One of these actions was to switch Flint's drinking water source from the Detroit system to the local Flint river in April 2014. The new water had different chemical characteristics which were overlooked by water officials. Of course many water systems have lead pipes, but these pipes are typically coated with layers of deposits, and the water is treated appropriately in order to prevent corrosion and the leaching of heavy metals. City officials failed to follow such necessary procedures, the pipes began to corrode, Flint's drinking water started to give off a different color and smell [11], and Flint residents were exposed to elevated levels of lead for nearly two years before the problems received proper attention. In August 2015 environmental engineers raised alarm bells about contaminated water[1] [21], not long after a pediatrician observed a jump in the number of Flint children with high blood lead levels[2][14], and by January 2016 the Flint Water Crisis was international news.

As attention to the problem was growing, government officials at all levels got involved in managing the damage and pushing recovery efforts. In looking for the primary source of lead in Flint's water distribution, attention turned to Flint's *water service lines*, the pipes that connect homes to the city water system. These service lines are hypothesized to be the prime contributor to lead water contamination across the United States [20]. Service lines, therefore, became a top priority for the City of Flint in February 2016. The Michigan state legislature eventually appropriated $27M towards the expensive process of replacing these lines at large scale; later the U.S. Congress allocated another nearly $100M towards the recovery effort. The group directed to execute the replacement program was called Flint Fast Action and Sustainability program (FAST Start), and their task was to remove as many hazardous service lines as possible up to funding levels.

The primary obstacle that the FAST Start team has faced throughout their work is uncertainty about the locations of lead or galvanized pipes. Although the U.S. Environmental Protection Agency requires cities to maintain an active inventory of lead service line locations, Flint failed to do so. Service line materials are in theory

[1] Prior work by the authors involved estimation of water lead contamination [1].
[2] For further analysis of blood lead levels, see [19]

documented during original construction or renovation, but in practice these records are often incomplete or lost. Most importantly, because the information is buried underground, it is costly to determine the material composition of even a single pipe. Digging up an entire water service line pipe under a resident's yard costs thousands of dollars. City officials were uncertain about the total number of hazardous service lines in the city, with estimates ranging from a few thousand to tens of thousands. Uncertainty about the service line material for individual homes has dramatic cost implications, as construction crews will end up excavating pipes that do not need to be replaced. These questions—how many pipes need to be replaced and which home's pipes need remediation—are at the core of the work in this paper.

Beginning in 2016, our team began collaborating directly with Flint city officials, analyzing the available data to provide statistical and algorithmic support to guide decision making and data collection, focusing primarily on the work of the FAST Start pipe replacement efforts. By assembling a rich suite of datasets, including thousands of water samples, information on pipe materials, and city records, we have been able to accurately estimate the locations of homes needing service line replacement, as well as those with safe pipes, in order to target recovery resources more effectively. Specifically, we have combined statistical models with active learning methods that sequentially seek out homes with hazardous water infrastructure. Along the way we have developed web-based and mobile applications for coordination among government offices, contractors, and residents. Over time, the number of homes' service lines inspected and replaced has increased, as seen in Figure 1.

In the present paper, we detail the challenges faced by decision-makers in Flint, and describe our nearly two years of work to support their efforts. With the understanding that many municipalities across the US and the world will need to undertake similar steps, we propose a generic framework which we call ActiveRemediation, that lays out a data driven approach to efficiently replace hazardous water infrastructure at large scale. We describe our implementation of ActiveRemediation in Flint, and describe the empirical performance and potential for cost savings. To our knowledge, this is the first attempt to predict the pipe materials house-by-house throughout a water system using incomplete data and also the first to propose a statistical method for adaptively selecting homes for inspection to replace hazardous materials in the most cost effective manner. This work illustrates a holistic, data-driven approach which can be replicated in other cities, thereby enhancing water infrastructure renovation effort with data-driven approaches.

Key Results. Among our main results, we emphasize that our predictive model is empirically accurate for estimating whether a Flint home's pipes are safe/unsafe, with an AUROC score of nearly 0.92, and a true positive rate of 97%. Since our approach involves a sequential protocol that manages the selection of homes for inspection and replacement based on our statistical model, we are also able to compare the model's total remediation cost to that of the existing protocol of officials. ActiveRemediation reduces the costly error rate (fraction of unnecessary replacements) to 2%, lowering the effective cost of each replacement by 10% and yielding about $10M in potential savings.

Methodology. Let us now give a birds-eye view of our methodological template. ActiveRemediation manages the inspection and replacement of water service lines across a city, with the long-term objective of replacing the largest number of hazardous pipes in a city under a limited budget. The formal in-depth exposition of this framework will be given in Section 3.

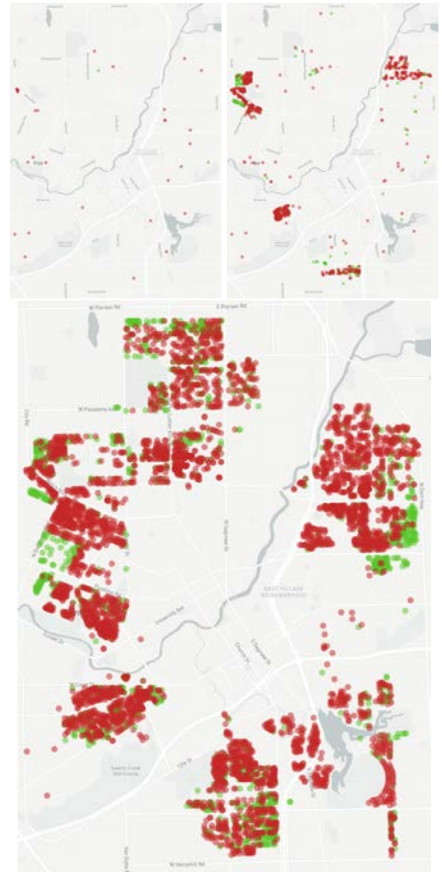

Figure 1: Progress of the replacement program. By March of 2016, only 36 homes had undergone replacement (top left); by December 2016, a total of 762 homes either been inspected or fully replaced (top right); as of September of 2017, this had grown to a total of 6,506 homes (bottom). Homes labeled green were selected for replacement but were deemed safe after copper lines were discovered by contractors.

Algorithm 1 ActiveRemediation

```
Input: parcel data, available labeled homes
for decision period t = 1, ..., T: do
    Predict hazardous/safe material via StatisticalModel
    if Budget remaining then querying StatisticalModel,
        Generate inspections via InspectionDecisionRule
        Generate replacements via ReplacementDecisionRule
        Input observed data to StatisticalModel
```

Since the process of identifying and replacing these lines around a city is naturally sequential, the decisions and observations made earlier in the process ought to guide decisions made at future stages. With this in mind, our framework continuously maintains three subroutines that are updated as data arrives. Following the outline

Date	Description
2016 Feb.	Attributes for all 55k parcels provided by the City of Flint
2016 Feb.	SL records digitized by M. Kaufman at UM Flint GIS
2016 March	Pilot Program, 36 homes visited, 33 SLs replaced
2016 June	Michigan DEQ provides SL private-portion inspections dataset
2016 Sept.	Phase One begins, contractors use our mobile data collection app
2016 Oct.	Fast Start begins hydrovac inspections to verify some home SLs
2016 Oct.	Congress appropriated $100M in WIIN Act.
2016 Dec.	Fast Start & authors release report: 20-30k replacements needed
2017 March	Federal court orders 18k homes to receive SL replacement by 2019
2017 Sept.	Fast Start replaced 4,419 hazardous service lines so far, identifying composition of a total of 6,506 homes.

Table 1: Timeline of service line data availability

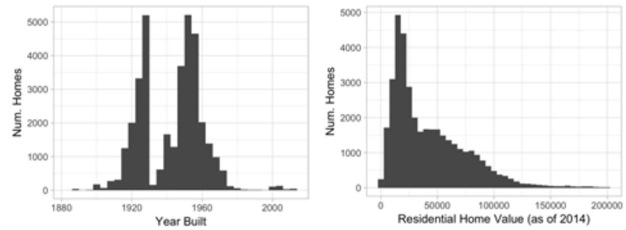

Figure 2: From city parcel data, distribution of home construction by year (left) and building value by dollar (right). The majority of the housing stock in Flint was built when it was a major automobile manufacturing hub, before current regulations about lead infrastructure were in place. Flint has experienced significant economic decline in recent years, leading to depressed real estate prices.

in Algorithm 1, the first of these is a StatisticalModel, that generates probabilistic estimates of the material type of both the public and private portion of each home's service lines. The input of this model is property data, water test results, historical records, and observed service line materials. The second subroutine is InspectionDecisionRule, the decision procedure that that generates a (randomized) set of homes for inspection. This should be viewed as an *active learning* protocol, with the goal of "focused exploration." The third routine, ReplacementDecisionRule, makes decisions as to which homes should receive line replacements; for reasons we discuss below, we typically assume that ReplacementDecisionRule is a *greedy* algorithm.

Roadmap. This paper is structured as follows. We begin in Section 2 by laying out the datasets available to us, with the story given chronologically to describe the shifting narrative as information emerged. We then explain the ActiveRemediation framework in greater detail in Section 3, and sketch out the statistical model mixed with the prediction, inspection, and decision-making framework. In Section 4 we employ ActiveRemediation on the data available in Flint, to show the empirical performance of our proposed methods in an actual environment, as well as in a simulated environment leveraged from Flint's data. We finish by detailing the potential for significant cost savings using our approach.

2 EMERGING DATA STORY OF FLINT'S PIPES

We now describe the various sources of data and the timeline during which these became available. This is summarized in Table 1 and more precise chronology is given throughout this section. More details will be available in the full version of this work.

2.1 Pre-crisis Information – Through mid-2015

In this section, we explain the relevant datasets that had been collected and maintained prior to the water crisis. This information, as we discovered later, was limited in both depth and quality.

2.1.1 Parcel Data. The city of Flint generously provided us with a dataset describing each of the 55,893 parcels in the city. These data include a unique identifier for each parcel and a set of columns describing City-recorded attributes of each home, such as the property owner, address, value, and building characteristics. A complete list of the parcel features is discussed in our previous work [7]. The distributions of the age of homes and their estimated values (Figure 2) tell an important story about the kinds of properties in Flint.

2.1.2 City Records of Service Lines. Initially, Flint struggled to produce any record of the materials in the city's service lines. Eventually, officials discovered a set of over 100,000 index cards in the basement of the water department[3] (see top of Figure 3). As part of a pro bono collaboration, the handwritten records have been digitized by Captricity.com and provided to the City of Flint.[4] Around the same time, a set of hand-annotated maps were discovered that contained markings for each parcel that specified a record of each home's service line (bottom of Figure 3). The map data was digitized by a group of students from the GIS Center at the University of Michigan-Flint lead by the director Prof. Martin Kaufman [10]. Many of the entries in the city's records list *two* materials for a given record, such as "Copper/Lead," but they do not specify the precise meaning of the multiple labels. However, our latest evidence suggests that, at least in the typical case, the double records were intended to specify that the second label ("Lead" in "Copper/Lead") indicates the public service line material (water main to curb stop), and the first label describes the private service line (curb stop to home), while an entry that is simply given as "Copper" may refer to both sections or only one. Lastly, there are a number of entries in the records that say "Copper/?" for the service line material, indicating missing information for the service line on the original handwritten records. Many other records are simply blank, recorded as "Unknown/Other."

2.2 Peak of Crisis & Replacement Pilot

In the wake of the crisis the State of Michigan began to discuss plans for lead abatement in Flint. It had become clear to lawmakers in Michigan that they would need to invest in a large-scale removal of lead pipes from the city. To begin, FAST Start initiated a pilot phase, with the goal of replacing the service lines of a small set of residences. Flint's Mayor and the FAST Start team awarded a contract to Rowe Engineering to replace pipes at 36 homes around the city. They selected these homes based on risk factors including the presence of high water lead levels, pregnant women, and children younger than 6 years old. Nearly all of the homes, 33 of 36, had some hazardous material (lead or galvanized) in one or both

[3] http://www.npr.org/2016/02/01/465150617/flint-begins-the-long-process-of-fixing-its-water-problem

[4] We would like to thank Captricity, especially their machine learning team, Michael Zamora, Michael Zamora, David Shewfelt, and Kayla Pak for making the data accessible.

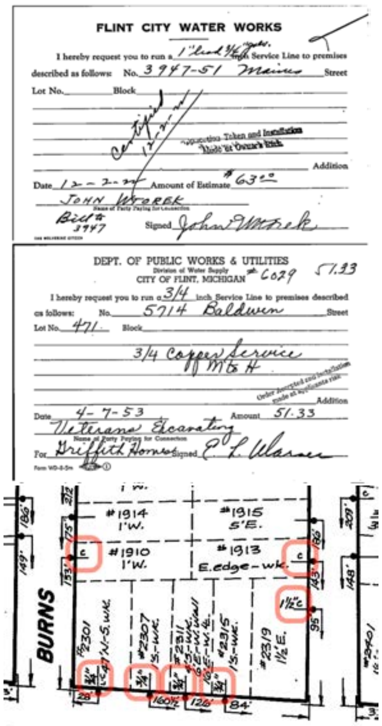

Figure 3: City officials located a set of over 100,000 handwritten index cards (top) with recorded work information dating back over 100 years, and annotated maps with data on home SLs (bottom). Red circles added to emphasize markings denoting material types.

portions of the service lines, while only 3 were safe. Therefore, the number of homes with physical verifications of both service line portions through September 2016 was only 36 out of over 55,000 homes. A map showing the progress of replacement in Flint can be found in Figure 1.

	Verified SL Materials (Public-Private)						
City Records	C-C	C-G	L-C	L-G	L-L	Other	All
Copper	1115	10	258	84	13	9	1489
Cop./Lead	109	20	816	91	15	25	1076
Galv./Other	113	18	565	1286	81	31	2094
Lead	24	2	29	14	12	3	84
Unknown	152	18	535	1169	118	42	2034

Table 2: Discrepancies between city records of service lines, and materials verified via inspection or replacement.

Meanwhile, in order to gather reliable information about private part of the service lines, the Michigan Department of Environmental Quality (DEQ) directed a team of officials and volunteers from the local plumbers union to personally inspect a sample of the homes of Flint residents. The public portion of the service line runs entirely under the street and sidewalk, while the private portion runs directly into the basement of the residents' home. Thus, the private portion can be inspected without any digging. The DEQ inspectors submitted their inspection results. As of June of 2016, the department had collected a data from over 3,000 home inspections. We consider this data to be reliable, since it was curated by DEQ officials who provided it to our team. This dataset allowed us to partially evaluate the reliability of the city records discussed in Section 2.1. It is important to note that the comparison is not "apples to apples," as the DEQ inspections were private-portion only whereas the labels in the city records did not specify which portion of the line was indicated. We report the confusion matrix between DEQ inspection data and city records in Table 2. The comparison suggests that, while the records were correlated with ground truth, the discrepancies were substantial.

2.3 Large-Scale Replacement, Mid-2016 to Now

Our group at the University of Michigan began engaging with the FAST Start team in the summer of 2016. One of the critical decisions the team needed to make was the selection of homes that would be recommended for service line replacement. According to the FAST Start payment agreements, contractors receive roughly half ($2500) the cost of a full replacement ($5,000) for excavated homes with copper on both public and private portions, due to removing concrete, refilling concrete, machine use, and labor. The choice of homes was deemed critically important, as the excavation of a home's service line that discovers a "safe" (e.g., copper) pipe is effectively wasted money, aside from the benefit of learning of the pipe's true material. Our work has focused on minimizing such unnecessary excavations, using the tools we describe below.

2.3.1 Early Replacement Activity and Findings (Fall 2016). By summer 2016, FAST Start had selected a set of 200 homes for replacement, scheduled to begin August, 31st. This selection is called Phase One. Like the Pilot Phase, their criteria included the presence of high water lead levels, pregnant women, children under six years old, as well as veterans and the elderly. In the present section, we describe how we helped facilitate data collection for Phase One, and how the results forced us to rethink our objectives and adjust our models.

By late September 2016, the early data from the service line replacement program began to arrive, and the rate of lead and other hazardous pipes discovered was alarming; 96% (165/171) of excavations revealed lead in the public portion of the line. These findings differed significantly from the city records, which had previously indicated that among those homes only 40% would contain lead in either portion. As data from Phase One arrived it was becomingly increasingly clear that *likely over 20,000 homes* have unsafe pipes serving their water – dramatically higher than earlier estimates. Critically, as these discoveries were being made, a debate was taking place in the U.S. Congress discussing the possibility of more than $100M in funding for the Flint's recovery efforts.

With the debate in the Congress ongoing, our team decided to put out an informal report to raise the alarm about the extent of the lead issue, and several news outlets reported on our findings [e.g. 5, 8]. This effort lead to a formal report in November of 2016 that provided a more precise estimate of the number of lead replacements likely to be needed [18], which was provided to the city's mayor, the DEQ, and the U.S. Environmental Protection Agency. Our report, based on comparing the city records and the data gathered from contractors, suggested that the number of needed replacements would be between 20,600 and 37,100. The large range accounts for the inherent uncertainty in data collection and model assumptions, as well as the question of *occupancy*. One challenge that is specific to Flint is the fact that around one third of the city's homes are not occupied, a rate that is the *highest in the country*[5].

2.3.2 Contractor Data Collection Application. With thousands of homes scheduled to have their water service lines excavated by multiple contractors, the collection and management of the data generated by this large-scale effort would prove to be a logistical challenge. While initially there was a plan in place to collect data via paper forms that would later get transferred to a spreadsheet, it was increasingly clear that digitally recording information, and storing it centrally, would be a more effective strategy and less prone to error.

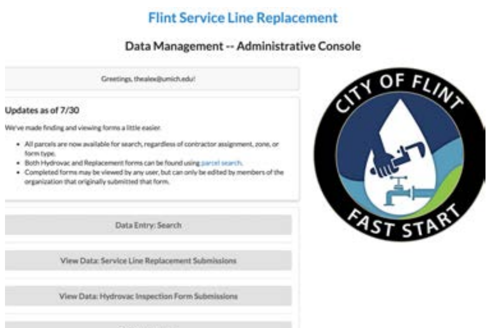

Figure 4: Mobile and web app, developed by the authors, to gather replacement data from contractors on-site.

Our team volunteered to facilitate the data collection efforts. In the fall of 2016, we developed a web and mobile application with various access levels. The latest version of this app is a custom-built web application using Python and the web framework using Flask. The users, on-site contractors as well as DEQ and Fast Start officials, are asked to select homes and to fill in essential information about service line work accomplished at each site. This information includes the excavated pipe materials, lengths, dates, and data on the home's residents. The output of the form appears in real-time in a live database with mapping capabilities. We adopted a tiered permissions structure with password-protected information to maintain the privacy of the data. The app continues to be used as of this writing for tracking progress for the public and for paying contractors for completed work.

2.3.3 Hydrovac Digging: Inspection without Replacement. The foremost challenge of a large-scale service line replacement program is the uncertainty about which homes possess safe service lines and which homes have lines made of hazardous materials. As of the summer of 2016, the only concrete verified data on pipe materials across the city consisted of the 36 data points provided by the Rowe engineering. By the end of Phase One, this number increased to about 250 homes. At this point, the excavation of pipes at a single home would cost anywhere from $2,500-5,000, a prohibitively high cost for data collection. At the same time, the available replacement data consisted of *cherry-picked homes*: houses were selected for line replacement if they were presumed to have an overwhelming likelihood of lead. These addresses and were highly concentrated in only three neighborhoods (see Figure 1) and provided nothing close to a representative sample of the broader city. We therefore realized, and emphasized to members of FAST Start, that the effort required a cheaper, quicker, and more statistically sound method to gather data.

Figure 5: Using a hydrovac truck for inspection, requires a large truck and crew (left) and exposes the pipe material underground (right).

After a lengthy discussion with water infrastructure experts and contractors, a new alternative emerged: *hydrovac inspections*. A hydro-vacuum truck, or simply a hydrovac (see Figure 5), has two main components: a high-pressure jet of water used to loosen soil and a powerful vacuum hose that sucks the loosened material into a holding tank. The hydrovac technique allows workers to dig a small hole quickly and then inspect whatever is observed underground. It is ideal for determining service line materials, as it can dig at the location of the home's curb box (connects the home's service line at the property line to the water main), and observe the pipe materials for both the public and private portions of the service line. The cost can be as low as $250 per inspection and often does not require prior approval from residents, as the digging site is mostly confined to city property. One limitation is that the hydrovac can only dig through the soil, and not through driveway or sidewalk pavement. This limitation led to unsuccessful excavations 20%-25% of the time, according to the hydrovac engineers.

The selection of homes for hydrovac inspection was one of the primary contributions of our team to FAST Start's efforts, and we were given wide discretion for "sampling" homes. This reflects the political and logistical challenges of service line replacement, as full excavation of service lines required a much longer process with oversight by the city council. We would emphasize that, in the following section where we describe our sequential decision protocols, our primary focus was on the model and inspection subroutines, and we assume the replacements are made using a simple greedy strategy.

3 PREDICTION & DECISION FRAMEWORK

In this section, we formally define the sequential decision-making problem for a city, in our case the city of Flint, seeking to remove all of the lead service lines from its homes under the following conditions: (i) for almost all homes, the service line materials of homes are unknown; (ii) there is a method of inspection to collect information; (iii) it is costly to excavate service lines that do not

[5]https://www.reuters.com/article/us-flint-vacancies-idUSKCN0VK08L

need to be replaced; and (iv) there is a fixed budget for replacement and inspection.

There are N total homes in the city, and it is unknown which homes need new service lines. We let the unknown label for home i be $y_i \in \{0, 1\}$, taking on the value 1 if the home needs a replacement and 0 otherwise. Note that a home needs replacement if either the public *or* private portion of the service line is hazardous. We also have information about each home, denoted by a vector x_i, with m features, that describe it (see Section 2). We want to learn the label y_i given x_i, for each $i = 1, \ldots, N$. We divide the procedure to find out these labels into two steps: first, a statistical model for prediction (StatisticalModel); and second, an algorithm that decides which homes to observe next (InspectionDecisionRule).

There is another decision rule, ReplacementDecisionRule, that determines which pipes to replace next. ReplacementDecisionRule is a *greedy* algorithm. That is this algorithm recommends that the replacement crew should go to the homes with the highest probabilities of having hazardous pipes. Given that, our InspectionDecisionRule is focused on learning, and ReplacementDecisionRule uses that learning to reduce costs.

3.1 StatisticalModel

In this section, we describe StatisticalModel, which assign a probability that a service line contains hazardous materials. StatisticalModel is a novel combination of predictive modeling using machine learning and Bayesian data analysis. First, a machine learning prediction model gives a prediction for the public and private portion of each home's service line using known features. These predictions then become the parameters to prior distributions in a hierarchical Bayesian model designed to correct some of the limitations to the machine learning model.

3.1.1 Machine Learning Layer. The machine learning layer of StatisticalModel outputs a probability of having a hazardous service line material for each home for which the material is unknown. Specifically, this layer gives a prediction, $\hat{y}_{i,k} = f_\theta(X_{i,k})$, the probability that service line portion k for home i is hazardous, and $X_{i,k}$ is a vector of features, described in Section 2.1. After examining several models empirically (see Section 4.1) we chose the machine learning layer, $f_\theta()$, to be XGBoost, a boosted ensemble of classification trees [6].

3.1.2 Hierarchical Bayesian Spatial Model Layer. One limitation of classification algorithms is how they handle unobserved variables, which may be correlated with the outcome. We address this limitation with a hierarchical Bayesian spatial model. This accounts for unobserved heterogeneity related to geographic location and similiarity of homes, which is used in hierarchical spatial models with conditional autoregressive structure [12, 13, 15, 16]. Empirically, each geographic region across the city (e.g., voting precincts) has a different number of observed service lines. While a city-level (pooled) model ignores precinct differences and a separate (unpooled) model for each precinct is limited by small sample sizes or even no observations, our full hierarchical (partially pooled) model strikes a balance with shrinkage. Precincts with little information will have their parameters pulled towards the city-wide distribution. Details of the Bayesian model, and how these are combined with the machine learning layer, are explained further in the full version of the paper.

Table 3: Summary of notation

Notation	Explanation
$\mathcal{X}$	observable feature space for each parcel/home
x_i, y_i	observable features for home i, label for home i
$\mathbf{h}_t/\mathbf{r}_t$	indicates "home i inspected/replaced at t?"
$y_t^{\mathbf{h}}/y_t^{\mathbf{r}}$	indicates "learned i's' label via inspect./replace?"
Q_{it}	indicates "learned i's label at t?"
q_{it}	indicates "already know i's label at t?"
$c^h, C^{\mathbf{r}+}, C^{\mathbf{r}-}$	cost of inspect., successful SLR, & failed SLR
U_t, L_t	set of labeled/unlabeled data at t

3.2 InspectionDecisionRule

Now we describe InspectionDecisionRule, which utilizes active learning [2, 3, 17] to efficiently allocate scarce resources to find and replace hazardous service lines. In general, a decision-maker may choose any active learning algorithm for inspection. In this work, we implement a version of Importance Weighted Active Learning (IWAL).

3.2.1 Active Learning Setup: Inspection and Replacement. We begin by describing the problem of efficiently locating and replacing hazardous pipes in a pool-based active learning framework (see Algorithm 2). Consider a budget of B total queries and a pool $\mathcal{P} = \{x_1, \ldots, x_n\}$ of unlabeled homes. Then at each time period t the algorithm will produce a probability vector $\phi_t = (\phi_{1,t}, \ldots, \phi_{n,t})$ that gives the probability that any home i is chosen at t.

Contractors can determine the material of a service line via either hydrovac inspection or service line replacement. When home i is chosen for hydrovac inspection at time t, we denote $\mathbf{h}_t = i$. When the service line for home i is replaced at time t, we denote $\mathbf{r}_t = i$. Once inspected or replaced, y_i is known for all subsequent rounds $t, t+1, \ldots$ and $p_{i,k}$ becomes 1 or 0, and we define $q_{i,t} = 1$ if home i has been observed through round t. $n_t^{\mathbf{h}}$ and $n_t^{\mathbf{r}}$ are the number of hydrovac and replacement visits, respectively. The number of successful replacements is denoted as $n_t^{\mathbf{r}+}$ (true positives) and the number of unnecessary replacements as $n_t^{\mathbf{r}-}$ (false positives).

We initially set $U_0 = \mathcal{P}$, and let $U_t = \{x_i | q_{i,t} = 0\}$ be the set of homes whose service line material is unknown at time t, and L_t be the set of homes with known service line materials. Finally, the budget also allows for a fixed number of inspections d for each period. The problem is how to select these d homes with unknown labels at each period t to maximize information gained.

3.2.2 Simple Active Learning Heuristics: Uniform and Greedy. We first propose several benchmark strategies for selecting homes for inspection. This family of algorithms randomly alternate between *random exploration* of the unobserved data and *greedy inspection* of the highest-predicted hazardous homes. As we see in Table 4, these decision rules differ in the costs they incur.

- **HVI uniform** *(egreedy(1.0))*: Select homes uniformly at random from the pool of those with unknown service lines.

Algorithm 2 ActiveRemediation for MultiEpochReplacement sequentially selects homes for both inspection and for replacement each epoch, incorporating ideas from both active learning and multi-armed bandits.

Input: parameters B, N, T
Input: initial observed data
for $t = 1, \ldots, T$: **do**
  Update: $\hat{p}_t(\theta) \leftarrow$ StatisticalModel $p_t(\theta, L_{t-1})$
  Inspect: $\mathbf{h}_t \leftarrow$ InspectionDecisionRule $\phi^{\mathbf{h}}(\theta)$
  Observe labels: $y_t^{\mathbf{h}}$
  Update: $\hat{p}_t(\theta) \leftarrow$ StatisticalModel $p_t(\theta, \{L_{t-1}, y_t^{\mathbf{h}}\})$
  Replace: $\mathbf{r}_t \leftarrow$ ReplacementDecisionRule $\phi^{\mathbf{r}}(\theta)$
  Observe labels: $y_t^{\mathbf{r}}$
  Update: $U_t \leftarrow \{U_t\} \setminus \{\mathbf{h}_t, \mathbf{r}_t\}$, $L_t \leftarrow \{L_t\} \cup \{\mathbf{h}_t, \mathbf{r}_t\}$
  $\text{TotalCosts}_t \leftarrow \text{TotalCosts}_{t-1} + (c^h + \mathbf{1}^{\mathbf{r}+}_{(c^{\mathbf{r}+})} + \mathbf{1}^{\mathbf{r}-}_{(c^{\mathbf{r}-})})$
  if $\text{TotalCosts}_t \leq B$ check budget **then** continue
  else stop
$\text{HitRate}^{\mathbf{r}}_T \leftarrow n^{\mathbf{r}+}_T / (n^{\mathbf{r}+}_T + n^{\mathbf{r}-}_T)$
$\text{EffectiveCost}_T = \text{TotalCosts}_T / n^{\mathbf{r}+}_T$

- **HVI greedy** *(egreedy(0.0))* Select the homes most likely to have hazardous service lines, based on current model estimates.
- **HVI ε-greedy** *(egreedy(ε))*: For a $1 - \varepsilon$ fraction of the inspections, select *greedily*, that is select homes for HVI based on the highest predicted likelihood of danger. For the remaining ε fraction, select homes uniformly at random for HVI. We experiment with values $\varepsilon = \{0.1, 0.3, 0.5\}$. Also, we note that **HVI uniform** and **HVI greedy** are special cases, with ε set to 1.0 and 0.0, respectively.

3.2.3 Importance Weighted Active Learning. We propose an algorithm that takes in the current beliefs about whether each home has hazardous pipe material, and outputs a decision of which homes should be inspected next period. This proposal is a variant of the Importance Weighted Active Learning (IWAL) algorithm [4]. The key idea behind IWAL is to sample unlabelled data from a *biased* distribution, with more weighted placed on examples with greater uncertainty, and then after obtaining the desired labels to incorporate the new date on the next iteration of model training. Our implementation of this approach takes the part of InspectionDecisionRule which is core to Algorithm 2. A full explanation of our IWAL implementation will be available in the full version of the paper.

3.2.4 Analyzing Costs. There are two categories of costs incurred in Algorithm 2: hydrovac inspections and replacement visits. Hydrovac inspections always cost the same amount and are denoted $c^{\mathbf{h}}$. Service line replacement costs, however, depend on what is actually in the ground. If contractors excavate a service line that does not need be replaced, we still incur a cost $c^{\mathbf{r}-}$ for labor and equipment, even though no replacement occurred. On the other hand, if contractors uncover a line that needs to be replaced then the direct cost of replacement is $c^{\mathbf{r}+}$.

But *effective cost per successful replacement* is greater than its direct cost, and we define formally it as TotalCosts$/n^{\mathbf{r}+}$, where

$$\text{TotalCosts} = c^{\mathbf{h}} n^{\mathbf{h}} + c^{\mathbf{r}+} n^{\mathbf{r}+} + c^{\mathbf{r}-} n^{\mathbf{r}-}$$

			Homes visited by rule		
Hydrovac Inspection	Replacement Visit	(Cost)	*Uniform*	*None*	*10%*
Finds Safety	→ not needed	($250)	230	0	23
Finds Danger	→ replaces Danger	($5,250)	770	0	77
None	→ finds Safety	($2,500)	0	230	207
None	→ replaces Danger	($5,000)	0	770	693
Effective Cost per Successful Replacement:			$5,325	$5,747	$5,705

Table 4: Average effective cost per successful replacement varies by simple InspectionDecisionRule, shown by 1,000 home visits.

(See Algorithm 2). In Flint, hydrovac inspection costs are summarized in Table 4. We note that the effective cost of a successful replacement is driven by two factors: the model accuracy (HitRate$^{\mathbf{r}}$) and the ratio of their costs, $c^{\mathbf{r}-}/c^{\mathbf{h}}$. Since unnecessary replacement visits can be avoided by prior inspection with a hydrovac, these two metrics, which naturally vary by city, will be critical guides to applying this approach to other cities.

4 AN EMPIRICAL ANALYSIS IN FLINT

In our empirical analysis, we use the data of the confirmed service line material from the 6,505 homes identified and replaced by Flint FAST Start, as of September 30, 2017 collected via our data collection app. This data is combined with our supplementary datasets describing homes (Section 2) and we train a suite of classification models to predict the presence of hazardous service line materials for a given home, and the predictive power of each model is measured on hold-out sets of homes (Section 4.1). After selecting a strong empirical model, we utilize the model predictions in our decision-making algorithms, which recommend those homes which will be most informative for inspection, and also those most likely containing hazardous service line materials for replacement (Section 4.2).

We emphasize that our methods and models were utilized by FAST Start officials for the management of the hydrovac process, and during the early days of the efforts we were given discretion over which homes would receive inspections. We used this freedom to select statistically representative samples, as well as targeted inspections on homes of interest. In practice, our modeling efforts had less impact on the choice of replacement homes, as these decisions carried greater political and logistical challenges.

4.1 Classification Algorithm Performance

Selecting a robust, precise, unbiased, and properly calibrated classification algorithm is key for our proposed active learning framework. Ultimately, the selected decision-making algorithm requires both accurate and well-calibrated probability estimates when selecting the next round of homes to investigate. To select such a classification model, we employ several machine learning model and compare them across various performance metrics. These metrics include the Area Under Receiver Operating Characteristic curve (AUROC), learning curves, and confusion matrices (including accuracy and precision). Using these scores, we find that tree-based methods are the most successful and robust category of models for this data. In particular, the model for gradient boosted trees implemented in the package `XGBoost` exhibits the strongest performance with a fewest data points.

4.1.1 ROC and Learning Curves. The overall accuracy of the best performing XGBoost model, based on a holdout set of 1,606 homes (25% of available data), is 91.6%, with a false-positive rate of 3% and false-negative rate of 27%. The homes falling in the top 81% of predicted probabilities are classified as having hazardous service lines. The ROC curves and AUROC scores show XGBoost's superior performance with an AUROC score of 0.939 on average in a range of [0.925, 0.951], Figure 6 and 7). While the ROC curves show a single run of each model, the AUROC scores are shown as distributions of 100 bootstrapped samples obtained using a stratified cross-validation strategy with 75%/25% of the data randomly selected for training/validation. We further examine AUROC scores using learning curves (Figure 8), using random subsets of data to illustrate diminishing returns of additional data on model performance using AUROC. We also introduce, *temporal learning curves.* These temporal learning curves reflect the exact order of data collection in 2016-17, and they show the AUROC as we re-estimate the model every two-week period to predict the danger for all remaining not-yet-visited homes. We finally ensure that the model's predicted probabilities, which we use to quantify our prediction uncertainty, are indeed well-calibrated probabilities. [6]

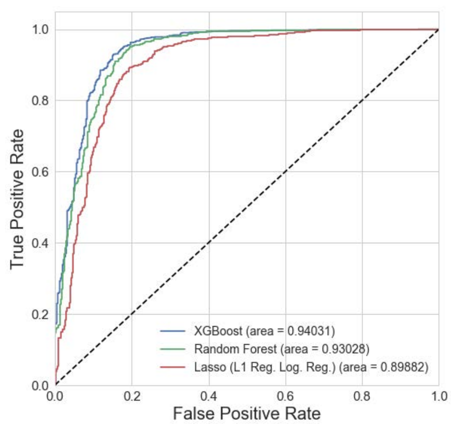

Figure 6: ROC curves measuring predictions of XGBoost, RandomForest, and lasso logistic regression on a random holdout set of all available data.

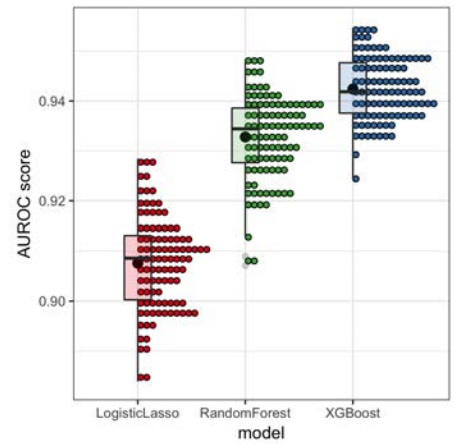

Figure 7: Empirical distributions of AUROC scores of classifiers over several runs on random holdout sets. Both XGBoost and RandomForest show marked performance improvement over lasso logistic regression, and XGBoost gives marginal improvement on RandomForest.

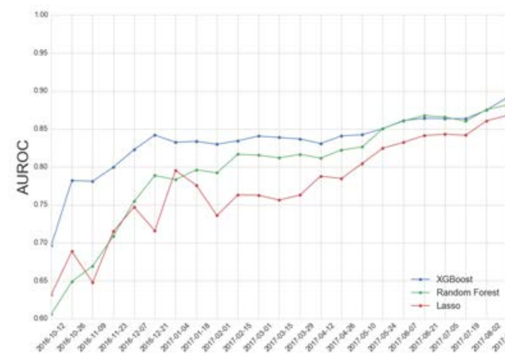

Figure 8: Temporal Learning curves for classification of hazardous service line materials. XGBoost consistently outperforms the other classifiers, especially at the beginning of the timeline when there is less data available.

4.1.2 Risk factors. Now that we have a robust predictive model, we can look at which features of a home and its surrounding neighborhood are the most predictive feature in identifying homes with hazardous service lines. But we are cautious to not make any causal claims from this analysis. We obtain the feature importance values[7] produced by each model by training with 20 bootstrapped samples of the data and reported the average feature importance values. The most informative home features relate to its *age*, *value*, and *location*, suggesting that the context (place and time) in which the home was built, as expected, is strongly correlated with service line material. For instance, homes built during and before World War II and those that are lower in value are more likely to contain lead in their public service line. Two additional features were the *city records* and the *DEQ private SL inspection reports*. Each was shown to be a noisy but useful predictor, as indicated earlier in Table 2.

4.2 ActiveRemediation: Evaluation

We now discuss our implementation of the ActiveRemediation framework applied to the particular case of Flint's large-scale pipe replacement program. With over $100M in investment, Flint is a perfect testbed to compare the performance of our proposed methods (developed in Section 3.2) with the actual empirical performance of the work of FAST Start thus far. Our goal is to show a high potential for savings by minimizing the number of unnecessary replacement visits, thus replacing more hazardous lines under the same budget.

4.2.1 Experimental testbed, and potential biases. Any experimental framework needs a quality dataset, with known labels for a large sample which we can evaluate our procedure. Fortunately for the City of Flint, where contractors have been working for over

[6] While not shown here, we also considered ExtraTrees, AdaBoost (with decision tree classifier), and Ridge Regression (regularized with L2 loss), but performance was lower than the three presented. Full details on hyperparameter optimization will be available in the full version.

[7] We calculate feature importance by weight, which is the normalized frequency with which a feature appears in a tree amongst the ensemble.

18 months, we have a total of 6,506 observations of service line materials. A natural choice for an experimental environment, which we call ActualFlint, is to use the set of observed homes in Flint as a template for the overall city, i.e. a municipality with precisely 6,506 homes whose service line material we can query as needed.

A major challenge of relying solely on observed data is that the actual home selection process is biased, in both the hydrovac inspections and the line replacements. While a certain fraction of the home selection was random, it was often reasonably arbitrary due to political and logistical constraints. For instance, many of the homes selected for service line replacement were chosen to maximize lead discovery. To assess the effect of sample bias, we developed an experimental environment, SimulatedFlint, in which we suppose Flint contains only those properties *not in* the observed dataset. For this dataset, labels are assigned based on the labeled hold-out data. With observed data as training, we used a K-Nearest-Neighbors (KNN) classifier to estimate a probability for each unknown home, and then sampled a Bernoulli random variable – "safe"/"unsafe" – to assign labels. This randomized dataset has lower potential selection bias concerns. In the reported results below, we focus on ActualFlint, but we note that results from SimulatedFlint were nearly equivalent.

4.2.2 Backtesting Simulation on ActualFlint. We quantify the cost savings from implementing our algorithm by comparing the sequential selection of homes from the proposed decision rules to what the Flint FAST Start initiative actually did in 2016-17. The goal is to stretch the allocated funds to remove hazardous pipes from as many homes as possible. One source of inefficiency in spending is unnecessary service line replacement (SLR) visits (the false-positive error rate). Therefore, our key performance metric is the SLR hit rate, i.e. the percentage of homes visited for replacement that required replacement.

The proposed approach greatly improves the hit rate. Our key finding from the simulation shows that we predict a reduced rate of costly unnecessary replacements visits from 18.8% (actual) to 2.0% (proposed). Figure 9 illustrates the direct comparison of hit rates for our proposed approach, IWAL(0.7), based on our ActualFlint simulation, compared to Flint FAST Start.

Second, the cost savings are substantial. The proposed algorithm, with a higher hit rate, increases the number of homes that receive service line replacements for the same number of visits. This, in turn, reduces the *effective cost* of a successful service line replacement. The effective cost includes both the direct costs of successful replacement visit and the average costs incurred by exploring homes from hydrovac inspections or unnecessary replacement visits. Having access to the exact same set of 6,505 homes actually observed, we find that the algorithm on average saves an additional 10.7% in funds per successful replacement (see Table 5). Across 18,000 total planned service line replacements, this would extend to an expected savings of about $11M out of current spending. In terms of the overall removal of lead pipes, this is approximately equivalent to 2,100 additional homes in the city that would receive safe water lines. These estimates are made using the current costs in Flint, where hydrovac inspection costs c^h = $250, unnecessary replacement costs c^{r-} = $2,500, and successful replacement costs c^{r+} = $5,000.

	Actual	Proposed Algorithm	
		Mean	*Range*
For every 1 successful replacement:			
Effective cost	$5,818	$5,196	($5,186 to $5,208)
Predicted savings ($)	–	$621.7	($610.4 to $632.4)
Predicted savings (%)	–	10.7%	(10.5% to 10.9%)
For every 1,000 successful replacements, the savings generate:			
Extra inspections	–	94	(92 to 96)
Extra replacements	–	120	(117 to 122)
For 18,000 successful replacement:			
Predicted savings ($ in millions)	–	$11.18m	($10.99m to $11.39m)

Table 5: Cost savings. The proposed method lowers the effective cost per successful service line replacement, saving $621.7 per home (10.7%), enough to remove lead from an additional 2,000 homes on the same budget.

Figure 9: Tracking hit rates over time, the proposed IWAL algorithm (blue; mean = 98.0%) outperform actual (green; mean = 81.2%; thick line is smoothed plot)

The proposed approach outperforms a competitive set of natural benchmark strategies. Instead of only comparing our proposed method to what actually occurred, we also consider a range of alternative methods. In particular, greedy (egreedy with 0% exploration) inspects the highest rate of hazardous homes inspected (HVI hitrate 91%), and uniform (egreedy with 100% exploration) inspects the lowest (63%). But IWAL does better with a more principled approach, selecting homes that are likely to be most informative, with risk probabilities near 70%. Figure 10 shows how IWAL and two greedy heuristics differ. Higher HVI hit rate is not better; instead, it is the choice of which homes to explore with inspection that matters. The uncertainty in performace of each algorithm comes from sampling variation from running 25 independent simulated experiments. We prefer IWAL to alternatives because it has greater savings and is less sensitive to tuning parameters.

We acknowledge some assumptions in our simulations. First, we only consider the cost of each job and not the time required for crews to move between homes, where there may be logistical issues with redirecting teams around the city. Second, in this analysis we have treated the ActualFlint as having only 6,506 homes of which all are visited. This creates an arbitrary finite end point, as the algorithm runs out of homes with unsafe service lines. To avoid this effect, the above calculations, figures, and tables are based on the first 4,500 replacement visits and 2,250 hydrovac inspections. Of course, to validate this, we would need access to a larger set, and thus we turn to our larger simulation using a full size of Flint. Finally, the results are robust to resource allocation schedule and batch size. We recognize that we used a schedule of SLR and HVI

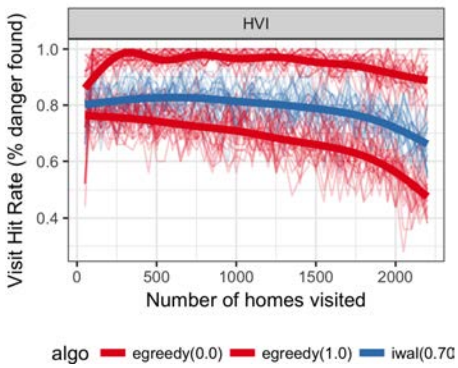

Figure 10: HVI Hit Rates. egreedy(0) tends to over-inspect whereas egreedy(1) is too conservative. IWAL more effectively optimizes HVI hit rate.

activities different than Flint FAST Start. To disentangle the confound between our choice of algorithms and the schedule, we ran an additional version of the ActualFlint backtest, with the schedule as closely aligned with Flint FAST Start in 2016-17 as possible. Across alternative scenarios tested the results differed only slightly.

4.2.3 Results from SimulatedFlint. In our second simulation, we demonstrate the potential value of deploying the algorithm at scale and characterize the long-term performance of the algorithms. Via SimulatedFlint we find that the proposed algorithms, with the aim of replacing hazardous lines from 18,000 homes out of a simulated city of 48,000 homes, can achieve 11.8% savings relative to the current rate of spending. The best algorithm using IWAL yields an average effective cost of $5,133 per successful replacement, better than $5,818 observed in Flint (Table 5). As a final note, the proposed algorithms' SLR hit rates are all above 98.0%.

ACKNOWLEDGMENTS

The authors would like to thank the FAST Start team for their phenomenal work and openness to collaboration. This includes Brigadier General (Ret.) Michael McDaniel, Ryan Doyle, Major Nicholas Anderson, and Kyle Baisden. Professors Lutgarde Raskin and Terese Olson, environmental engineering faculty at U-M, provided invaluable scientific support throughout. We are incredibly grateful to the work of, and communication with, Professor Martin Kaufman and Troy Rosencrantz at U-M Flint's GIS Center. We would like to thank Captricity, especially their machine learning team, Michael Zamora, Michael Zamora, David Shewfelt, and Kayla Pak for making the data accessible, and Kuang Chen for the generous support. We had major support from Mark Allison and his team of U-M Flint students. Rebecca Pettengill was enormously generous with her time and ability to help in the Flint community. We thank U-M Professors Marc Zimmerman and Rebecca Cunningham for their encouraging and helpful discussions. Among the many students involved in this work, we would like to recognize the roles of Jonathan Stroud and Chengyu Dai. And this work would not have happened without the expertise and enthusiasm of the students in the Michigan Data Science Team (MDST, http://midas.umich.edu/mdst/, [9]). The authors appreciate the many seminar and conference participants at U-M and elsewhere for their feedback on the academic work. The authors gratefully acknowledge the financial support of the Michigan Institute for Data Science (MIDAS), U-M's Ross School of Business, Google.org, and National Science Foundation CAREER grant IIS 1453304.

REFERENCES

[1] Jacob Abernethy, Cyrus Anderson, Chengyu Dai, Arya Farahi, Linh Nguyen, Adam Rauh, Eric Schwartz, Wenbo Shen, Guangsha Shi, Jonathan Stroud, et al. 2016. Flint Water Crisis: Data-Driven Risk Assessment Via Residential Water Testing. *arXiv preprint arXiv:1610.00580* (2016).

[2] Maria-Florina Balcan, Steve Hanneke, and Jennifer Wortman Vaughan. 2010. The true sample complexity of active learning. *Machine learning* 80, 2 (2010), 111–139.

[3] Maria-Florina F Balcan and Vitaly Feldman. 2013. Statistical active learning algorithms. In *Advances in Neural Information Processing Systems*. 1295–1303.

[4] Alina Beygelzimer, Sanjoy Dasgupta, and John Langford. 2009. Importance weighted active learning. In *Proceedings of the 26th annual international conference on machine learning*. ACM, 49–56.

[5] Steve Carmody and Mike Brush. 2016. Flint might have a bigger problem with lead pipes than previously thought. http://michiganradio.org/post/flint-might-have-bigger-problem-lead-pipes-previously-thought. (2016). (Accessed Feb, 16, 2017).

[6] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 785–794.

[7] Alex Chojnacki, Chengyu Dai, Arya Farahi, Guangsha Shi, Jared Webb, Daniel T. Zhang, Jacob Abernethy, and Eric Schwartz. 2017. A Data Science Approach to Understanding Residential Water Contamination in Flint. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17)*. ACM, New York, NY, USA, 1407–1416. https://doi.org/10.1145/3097983.3098078

[8] Matthew Dolan. 2016. Far more Flint homes have lead lines than expected, report shows. (2016). http://www.freep.com/story/news/local/michigan/flint-water-crisis/2016/09/28/more-than-half-flint-homes-could-have-lead-lines-report-shows/91225284/ (Accessed Feb, 16, 2017).

[9] Arya Farahi and Jonathan Stroud. 2018. The Michigan Data Science Team: A Data Science Education Program with Significant Social Impact.

[10] Ron Fonger. 2015. Flint data on lead water lines stored on 45,000 index cards. (2015). http://www.mlive.com/news/flint/index.ssf/2015/10/flint_official_says_data_on_lo.html (Accessed Feb, 16, 2017).

[11] Ron Fonger. 2015. Here's how that toxic lead gets into Flint water. http://www.mlive.com/news/flint/index.ssf/2015/10/see_step_by_step_how_lead_is_g.html. (2015). (Accessed Feb, 16, 2017).

[12] Alan E Gelfand and Penelope Vounatsou. 2003. Proper multivariate conditional autoregressive models for spatial data analysis. *Biostatistics* 4, 1 (2003), 11–15.

[13] Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. 2014. *Bayesian data analysis*. Vol. 2. Chapman & Hall/CRC Boca Raton, FL, USA.

[14] Mona Hanna-Attisha, Jenny LaChance, Richard Casey Sadler, and Allison Champney Schnepp. 2016. Elevated blood lead levels in children associated with the Flint drinking water crisis: a spatial analysis of risk and public health response. *American journal of public health* 106, 2 (2016), 283–290.

[15] Duncan Lee. 2011. A comparison of conditional autoregressive models used in Bayesian disease mapping. *Spatial and Spatio-temporal Epidemiology* 2, 2 (2011), 79–89.

[16] Duncan Lee. 2013. CARBayes: an R package for Bayesian spatial modeling with conditional autoregressive priors. *Journal of Statistical Software* 55, 13 (2013), 1–24.

[17] Alexander Liu, Goo Jun, and Joydeep Ghosh. 2008. Active learning with spatially sensitive labeling costs. In *NIPS Workshop on Cost-sensitive Learning*.

[18] Kristin Moore. 2016. Number of Service Lines that Need Replacing in Flint Rises to 29,100, According to Study. https://www.cityofflint.com/2016/12/01/number-of-service-lines-that-need-replacing-in-flint-rises-to-29100-according-to-study/. (2016). (Accessed Feb, 16, 2017).

[19] Eric Potash, Joe Brew, Alexander Loewi, Subhabrata Majumdar, Andrew Reece, Joe Walsh, Eric Rozier, Emile Jorgenson, Raed Mansour, and Rayid Ghani. 2015. Predictive modeling for public health: Preventing childhood lead poisoning. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2039–2047.

[20] Anne Sandvig, P Kwan, G Kirmeyer, B Maynard, D Mast, R Rhodes Trussell, S Trussell, A Cantor, and A Prescott. 2008. *Contribution of service line and plumbing fixtures to lead and copper rule compliance issues*. Environmental Protection Agency, Water Environment Research Foundation.

[21] Michael Torrice. 2016. How Lead Ended Up in Flint's Tap Water. *Chem. Eng. News* 94, 7 (2016), 26–29.

EXHIBIT C

Inventory of Service Lines in Flint

Jacob Abernethy and Eric Schwartz (Professors at University of Michigan, Ann Arbor)
Arya Farahi and Jared Webb (PhD students at U-M)
Nicholas Anderson (U.S. National Guard, FastStart), **Ryan Doyle** (MDOT, FastStart)

We acknowledge Martin Kaufman and Troy Rosencrants (U-M Flint GIS Center) who initially digitized Flint City Records.

What has been found?

There are over 55,000 unique parcels of land in Flint with roughly 51,000 residential properties. We have physical verification of the complete service line materials for 457 homes.

- 36 from replacements Rowe pilot;
- 262 from replacement phases 1 (complete) and 2 (in progress); and
- 159 from hydrovac program.

Using this data as of November 1, 2016, we summarize what has been found, compare it to what the city records suggested, and then estimate the total number of homes with lead requiring partial or full service line replacements.

	Truth in Public Service Line:			
Method of discovery	**Can't tell**	**Copper**	**Lead**	**All**
Hydrovac	72	65	94	231
Replacement Phase 1 (and 2)*	0	4	257	262
Rowe	0	6	30	36
All	72	75	381	529

*Note: One home in Replacement Phase had a Galvanized public service line.

Among the 457 homes where the service lines were physically verified, there were 381 public service service lines made of Lead, 75 Copper, and 1 Galvanized.

The hydrovac teams visited a total of 231 homes. They could not inspect the service lines at 72 homes for any one of a variety of reasons, such as, the curbstop was under a driveway, or they were unable to find curbstop. Rowe Pilot Phase was completed in March 2016. Replacement Phase 1 is complete (and excludes the Rowe Pilot Phase). Replacement Phase 2 is underway. The data we use throughout this document includes results from the replacements (from Phase 1, 2, and Rowe Pilot) and hydrovac program.

The key decision is whether the home needs Full Replacement (Public and Private portion), Partial Replacement (Public or Private portion), or No Replacement. If a service line portion is made of Lead or Galvanized (or Tubeloy) we say it requires a replacement. A copper portion does not require replacement.

Service Line Replacement Needed/Performed:		
No	Partial	Full
70	208	180

The replacement and hydrovac programs show that the vast majority of the concern (lead) is typically found in the public portion of the service line. We present the true public service line material grouped by each city record type.

	Truth in Public Service Line:				
City Records	**Can't tell**	**Copper**	**Lead**	**All**	**Replace Rate**
Copper	19	42	16	77	21%
Copper/Lead	16	9	145	170	85%
Galvanized/Other*	1	2	98	102	97%
Lead	8	2	22	32	69%
Unknown/Other	28	20	100	148	68%
All	72	75	381	529	72%

Note: For the one home in Replacement Phase which had a Galvanized public service line but its City Record was Galvanized/Other. The City Records come from UM Flint GIS Center data. The row for Lead represents the following labels in the City Records: "Lead," "Lead/Zinc," "Lead/Tubeloy," and "Tubeloy." Copper includes "Copper," "Copper/Zinc." The Replace Rate is the proportion out of homes that were successfully physically verified.

We report a "Replace Rate," which is the proportion of homes in any group requiring a Partial or Full replacement. Partial Replacement occurs when Lead or Galvanized appears in only the Private or Public portion. Full Replacement occurs when Lead or Galvanized appears in both portions of the service line.

Looking at the complete description of the Public and Private portions of the service lines suggest how many partial replacements and full replacements are needed. We also provide these rates in two ways -- using all data and using only data from the hydrovac program.

The following letters will be used in the tables below:

C-C = Copper-Copper	L-C: Lead-Copper
C-G = Copper-Galvanized	L-G: Lead-Galvanized
C-L = Copper-Lead	L-L: Lead-Lead
	L-O: Lead-Other

All homes where public and private lines verified

	Truth in Public-Private service line											
City Records	Can't tell	C-C	C-G	C-L	L-C	L-G	L-L	L-O	All	Full Replace Rate	Partial Replace Rate	Replace Rate
Copper	19	42	0	0	10	4	2	0	77	10%	17%	27%
Copper/Lead	16	7	0	2	134	9	2	0	170	7%	88%	95%
Galvanized/Other	1	2	1	0	24	72	2	0	102	73%	25%	98%
Lead	7	2	0	0	5	3	14	1	32	68%	20%	88%
Unknown/Other	28	17	2	1	29	48	21	2	148	57%	27%	84%
All	71	70	3	3	202	136	41	3	529	39%	45%	84%

Note: One home had a Galvanized/Other record, and actually had Galvanized-Copper, since that is the only case of Galavnized in public line, we count it with the Copper-Galvanized.

Homes where public and private lines were verified by hydrovac only

	Truth in Public-Private service line											
City Records	Can't tell	C-C	C-G	C-L	L-C	L-G	L-L	L-O	All	Full Replace Rate	Partial Replace Rate	Replace Rate
Copper	19	39	0	0	2	3	1	0	64	9%	4%	13%
Copper/Lead	16	7	0	0	20	1	0	0	44	4%	71%	75%
Galvanized/Other	1	0	0	0	1	2	0	0	4	67%	33%	100%
Lead	7	1	0	0	1	0	9	1	19	75%	8%	83%
Unknown/Other	28	16	1	1	17	19	17	1	100	50%	26%	76%
All	71	63	1	1	41	25	27	2	231	32%	27%	59%

How Much Lead Is There? Estimating the Number of Service Line Replacements Needed

Estimating the number of lead service lines in Flint is still not easy. The service line replacement program targets at-risk homes, so it does not give us a representative sample of the 55,000 parcels in Flint. The hydrovac program samples from a broader group of homes since its purpose is to gain information, but still does represent all of Flint.

So how many homes require a Full replacement (private and public portion) or Partial replacement?

We estimate this by first calculating the replacement rate by City Record. For example, a home with a City Record of "Copper/Lead" has chance of 75% needing a replacement (71% for Partial or 4% for Full). Then we take into account how common each City Record is throughout Flint. For instance, there are 4,161 Copper/Lead records (7% of Flint).

We use the estimated "Replacement Rate" described above. Using the hydrovac data only reflects a lower total number of replacements than using hydrovac and replacement data. We will use the rates from the hydrovac only since we know the hydrovac sample is a better reflection of all of Flint than the homes that have already received service line replacement.

Using the rates observed in the sample of homes so far, we estimate approximately 29,100 parcels (about 52% Flint parcels) would require some replacement. Among those, about 17,500 would be full replacements and 11,600 would be partial replacements.

We want to provide some uncertainty around this estimate and give the reader an illustration of how sensitive this estimate may be to changes in what we understand about the true rate of lead corresponding to each City Record. For instance, if you decided to be optimistic that only 60% of parcels with non-Copper records and 10% of Copper records some lead or galvanized service lines, then you would estimate 20,600 parcels required some replacement.

City Record	Percent of City	Total Parcels	Optimistic Assumption	Hydrovac Only	Replacement and Hydrovac	Pessimistic Assumption
Copper	46%	25843	10%	13%	27%	30%
Copper/Lead	7%	4161	60%	75%	95%	100%
Galvanized/Other	22%	12261	60%	100%	98%	100%
Lead	0%	111	60%	83%	88%	95%
Unknown/Other	24%	13517	60%	76%	84%	95%
Total Predicted Number of Replacements			20614.3	**29106.4**	**34398.3**	37121.5
As Percent of All Parcels			36%	**52%**	**61%**	66%

*Note: Lead represents records saying "Lead, Lead/Tubeloy Lead/Zinc Tubeloy Copper/Tubeloy." "Replacement Rate" is the probability of finding either lead or galvanized in either the private or public service line.

It is important to note that since the replacement program made up about 2/3 of our records and they were located in three at-risk neighborhoods, it still does not give us a representative sample of the 55,000 parcels in Flint. A large scale hydrovac excavation project is needed to determine the true percentage of the lines that are all or partially lead. It would also allow for a more sophisticated estimate using statistical algorithms using all factors that describe parcels such as city record of service line, year built, home value, land value, zoning, location, vacancy status, home condition, and others. In addition to determining a much more precise number of lead lines, the hydrovac excavation project would provide us with other useful information, such as breakdowns by city ward, occupied vs. vacant, age of home, material listed in city records, and any other parcel attribute. After receiving guidance on how to define occupied residential homes, we will be able to provide these numbers for occupied vs unoccupied residential properties.

We will continue to update these results and add additional helpful information as new data is submitted.

EXHIBIT D

DOCUMENTS PROVIDED BY PLAINTIFFS' COUNSEL IN AUGUST 2018

"Phase IV SLR Completed Addresses (2018).xlsx," a list of all Phase IV replacements completed in 2018 as of August 15, 2018

"Phase V SLR Completed Addresses.xlsx," a list of all Phase V replacements completed as of July 31, 2018

"Phase V SLR Issued Addresses 20180815.xlsx," a list of all addresses that have been issued to their contractors for excavation

"Copy of 2018.08.28 FAST Quarterly Report Data-forUofM.xlsx," a dataset regarding the results from excavations and service line replacements conducted between May 15, 2018, and August 14, 2018

"2017-09-21-Att-Active 7-6-17.csv," a list of active water accounts in Flint as of July 6, 2017

"FAST START ACCOUNTS 08-02.18.xlsx," a list of all active water accounts in Flint as of August 2, 2018

EXHIBIT E

Analysis of the City of Flint's Service Line Excavations in 2018 by Ward (as of Aug. 15, 2018)

Ward	Total number of eligible homes not yet excavated, 12/31/2017	Predicted number of hazardous service lines among eligible homes not yet excavated, 12/31/2017	Predicted hit rate among homes not yet excavated, 12/31/2017	Total number of homes excavated, 1/1/2018 - 8/15/2018	Observed number of hazardous service lines identified, 1/1/18 - 8/15/18	Observed hit rate, 1/1/2018 - 8/15/2018
5	1454	1162.9	80.0%	163	156	95.7%
8	3551	1134.9	32.0%	455	172	37.8%
7	2291	701.1	30.6%	216	114	52.8%
2	2404	690.1	28.7%	1220	46	3.8%
6	2017	684.7	33.9%	210	52	24.8%
9	1936	607.7	31.4%	230	42	18.3%
3	1418	473.7	33.4%	215	90	41.9%
1	2631	455.3	17.3%	363	53	14.6%
4	2433	411.4	16.9%	702	17	2.4%
All Wards	**20,135**	**6321.8**	**31.4%**	**3774**	**742**	**19.7%**