Data Analytics and Machine Learning for Environmental Protection: Targeting Air Inspections

Benjamin Laufer

Princeton University, Department of Operations Research and Financial Engineering
Volunteer Intern, EPA Region 2 DECA–CAPSB, Data Management Team

January, 2018

Introduction

Ensuring compliance with environmental standards is a challenging endeavor. The Environmental Protection Agency (EPA) and analogous state agencies are tasked with holding public and private groups accountable to abide by environmental laws. While routine inspections are carried out across the country, there are too many polluting facilities to check every site for violations. Instead, with limited time and resources, inspectors must strategically focus their attention on a handful of sites most likely to pose an environmental risk to Americans. As historical data becomes more widely available and trustworthy, inspectors can make more informed decisions about which sites to target.

This project aims to uses predictive modeling and machine learning techniques to assist the EPA with Clean Air Act (CAA) targeting. The algorithm put forward by this study uses the available data to find an optimal policy for inspecting facilities that are most likely to be violating air pollution standards. The policy function, implemented in R Programming Language, classifies inspections as either 'worth conducting' or 'not worth conducting,' based on the likelihood that an inspection leads to a formal legal enforcement action. The algorithm iterates through the data by date to update and refine its policy in order to maximize expected reward. The model seems to be a viable opportunity for EPA's Department of Enforcement and Compliance Assistance (DECA) to improve its inspection-targeting strategy. Additionally, the model has certain advantages over existing algorithms used in forecasting inspections and allocation problems.

The Issue

Environmental laws are only effective if they hold organizations accountable for non-compliant and unsafe pollution practices. The EPA and state governments must conduct inspections each year to ensure that facilities meet environmental standards established by the Clean Air Act (CAA), Clean Water Act, Safe Drinking Water Act, Resource Conservation and Recovery Act (RCRA), and others. With recent advances in the scope and availability of historical data, inspectors and enforcement branches have more tools at their disposal to strategically allocate time and resources. Still, this paper suggests that current EPA targeting strategies do not fully utilize the available data.

Currently, CAA inspectors choose targets based on a few key policies and facility characteristics.[1] Very large facilities – known as Title V Majors – are required by law to be inspected periodically. Smaller facilities do not have required inspections and are therefore especially important to target as potential non-compliant or even fraudulently registered polluters.

Important sources of data include self-reported emissions data from the Toxics Release Inventory (TRI) and National Emissions Inventory (NEI),[2] as well as stack EPA- and state-conducted stack tests of pollution levels.[3] Once a significant violation is found, the EPA often targets organizations that have common attributes with the non-compliant organization.

While these rules of thumb are likely to help EPA's targeting strategy, formal data analytics and mathematics may further help enforcement branches meet their goals. This is

[1] Observations of current EPA targeting operations come from personal experience as an intern in EPA Region II. Experiences include frequent meetings with CAA enforcement branch members and inspectors.
[2] Toxics Release Inventory data and information can be found at https://www.epa.gov/toxics-release-inventory-tri-program/learn-about-toxics-release-inventory. National Emissions Inventory data and information can be found at https://www.epa.gov/air-emissions-inventories/national-emissions-inventory-nei. Only TRI data, downloaded from EPA Envirofacts Database, was used for this particular study.
[3] Stack test data accessed via EPA's Enforcement and Compliance History Online (ECHO), retrieved at: https://echo.epa.gov/tools/data-downloads.

because, of course, human decision-making is imperfect. An internal EPA study by Stephanie Wilson revealed that there are certain 'hot-spots' and 'cold-spots' in Region II that are disproportionately over- or under-inspected.[4] While these particular locations are confidential to non-EPA personnel, more general and qualitative evidence suggests that EPA inspections biases do exist.

Figure (1) depicts inspections in Region 2 over the past 10 years. As to be expected, inspections occur more often in heavily populated areas. However, beyond that, it seems that very many of the inspections that are outside of cities occur near major roads between cities. One explanation is that there are simply more facilities in easy-to-access places, so we should expect more inspections along roads. But it seems very possible that part of the observed phenomenon is simply a function of convenience. Inspectors may not want to venture far for an individual facility, even if that facility may be just as likely to pose a threat to the environment as a facility in the middle of New York City. Of course, the EPA must consider the costs and resources that might be required to inspect a far-away, potentially hazardous facility. Still, it is important to acknowledge where biases might exist, and how these biases can impact successful enforcement.

Beyond the potential for bias, a data-driven approach may prove beneficial simply because of the scope of data available. Currently, inspectors consider only a handful of features when making their decisions.[5] But a data-driven approach is able to handle hundreds of variables that may be causally related to whether an inspection is worth conducting. Further, machine-learning algorithms can analyze trends and changes over time and accordingly refine the decision-making strategy for targeting.

[4] Stephanie Wilson is a colleague on the Region 2 Data Management Team. She used geographical data software to cluster inspections by location. The quantitative results of the study are confidential to non-EPA personnel.
[5] Information comes from meetings with Air branch inspectors in EPA's New York City office.

Improving the EPA's targeting strategy will enable the agency to find and fix hazardous, non-compliant pollution behavior so that fewer American citizens may be exposed to Hazardous Air Pollutants (HAPs). A data-driven approach will not only help increase the efficiency of EPA operations, but also protect citizens in the process.

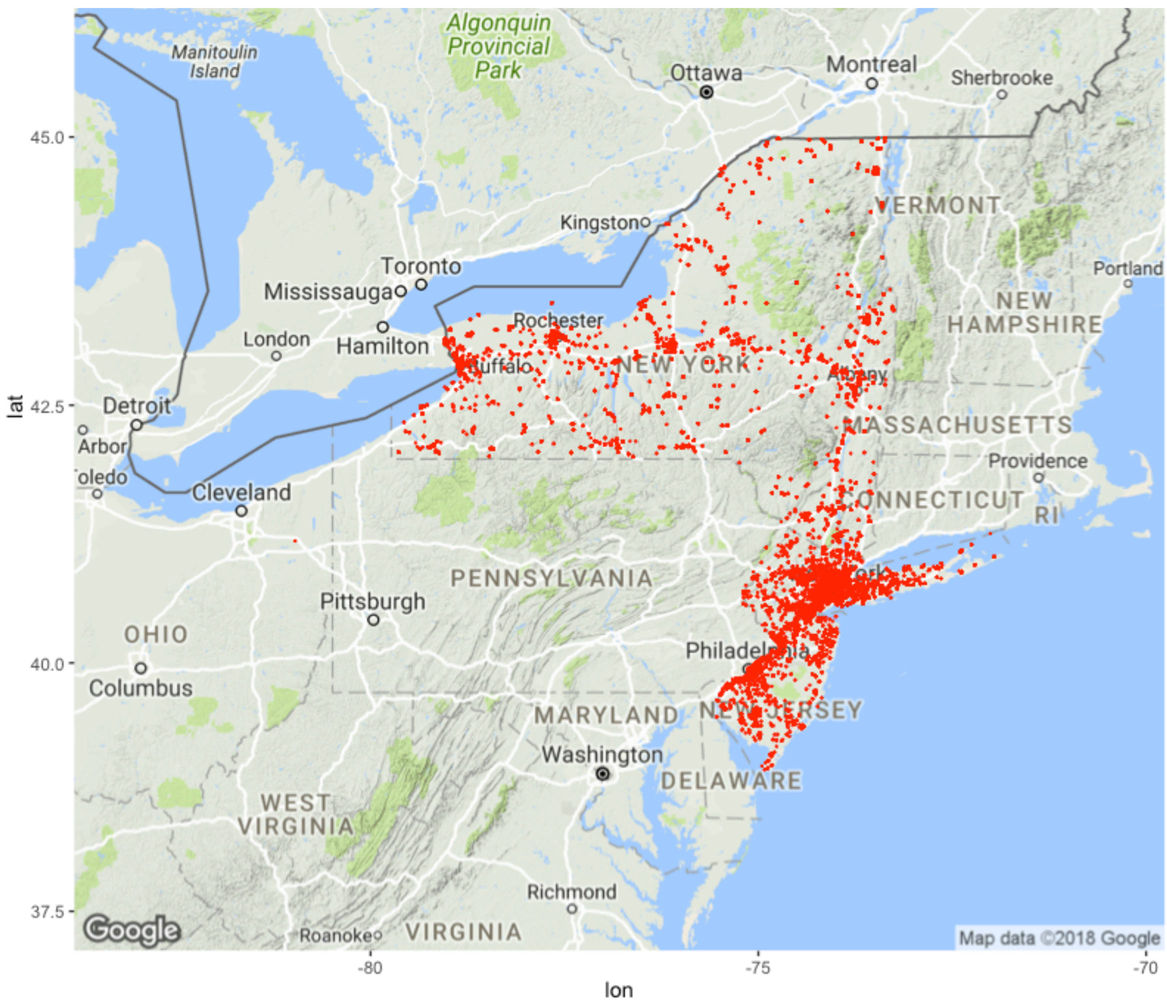

Figure (1)*: Map of state and EPA inspections in New York and New Jersey over the past 10 years Each red point represents one full- or partial-compliance evaluation. Many inspections occur in cities and near roads between cities.*

Finding the Proper Model

Using statistics to optimize inspections is a challenge for both enforcement agencies and researchers. A brief overview of related studies will help describe the motivations and reasons for the model used in this project.

Very similar research has been conducted on solid waste management inspections under the U.S. Resource Conservation and Recovery Act (RCRA). University of Chicago has led research initiatives to devise models that optimize RCRA inspections for solid waste management and pollution.[6] After meetings with both U. Chicago researchers and EPA employees, the initial approach that this project took was to implement a Random Forest algorithm - a supervised machine-learning algorithm that uses decision trees to classify inspections that are likely to uncover non-compliance.[7]

Another related project is the Food Inspection Forecasting project conducted by the City of Chicago. The open-source project on GitHub used data analytics to markedly improve the city government's ability to find non-compliant and hazardous health practices in restaurants.[8] Also using a supervised model-based algorithm, the project conducts regressions with regularization parameters to identify which inspections maximize the probability of non-compliance.

The models implemented in these studies are promising strides towards more effective regulation of non-compliant and hazardous behaviors by both companies and public facilities. However, they have their disadvantages. In particular, the Random Forest and traditional

---

[6] See Eric Potash, et. al. "Predictive Modeling for Environmental Protection: Hazardous Waste Management." *U. Chicago*, 2016. http://k2co3.net/assets/pdf/rcra_preprint.pdf.
Ongoing research is being conducted by Cynthia Gyles and Sarah Armstrong at Univerity of Chicago Urban Labs. See https://urbanlabs.uchicago.edu/labs/energy-environment.
[7] A full report and implementation of the Random Forest model can be provided by request. The report was written for EPA Region II in August 2017.
[8] See "Food Inspection Forecasting: Optimizing Inspections with Analytics." *City of Chicago* 2017. https://chicago.github.io/food-inspections-evaluation/.
For implemented code in R and more resources, see https://github.com/Chicago/food-inspections-evaluation.

regression models fail to account sufficiently for facilities that have never been inspected. This is a typical example of the 'multi-armed bandit' problem, as described frequently in decision-making literature.[9]

The multi-armed bandit problem is a question of allocation that deals with the decision between "exploitation" and "exploration."[10] Classically introduced using the example of multiple slot machines, the bandit problem arises when the winning rate for each slot machine is unknown. In this type of problem, an optimal algorithm must strike a balance between playing the machines that have already been successful ("exploitation") and finding new machines that may offer even higher winning rates ("exploration"). Our EPA problem is analogous to the bandit problem because there are so many facilities that have never been inspected before. The EPA knows they must inspect major polluters every few years – what they're interested in is those facilities that have been able to avoid EPA enforcement for years. Thus, while the precise workings of University of Chicago's algorithm is confidential,[11] it is likely that the type of model used for optimizing RCRA inspections will favor those facilities that have had many prior infractions to those that have never been inspected.[12]

To solve this problem, a different approach to the data is necessary. In particular, performing a regression on the entire dataset may not properly account for changes in particular facilities over time. A favorable approach could iterate through the history of enforcement practices at EPA and continually update its parameters based on the changing data. Accounting

---

[9] See Sébastien Bubeck and Nicolò Cesa-Bianchi, "Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems," *Foundations and Trends in Machine Learning, vol. 5, no. 1*, 2012. http://dx.doi.org/10.1561/2200000024.
[10] Ibid., 1-2.
[11] Information about confidentiality comes from email consultation with a Sarah Armstrong, a researcher at University of Chicago's Urban Labs.
[12] Potash,et. al. (6) discuss selection bias and non-random selection of facilities in historical data, and try to re-weight observations to account for the bias. However, this issue is distinct from the Bandit Problem, which arises even when historical data is a uniformly random sample.

for the multi-armed bandit problem might further improve the algorithm's ability to identify non-compliant facilities.

Ultimately, the approach used for this project is a semi-model free, unsupervised machine learning method that tries to tries to implement both potential improvements. For discretized time-steps, the algorithm iterates through the history of available data to build a policy that finds the facilities most likely to be in violation of the Clean Air Act. Then, at each time step, the algorithm can draw facilities that have never been inspected before to balance the exploitation-exploration tradeoff. Hopefully, this model can more aptly describe this particular allocation decision-making problem and, further, inform future work in applied data analytics and forecasting.

The Policy Function Search

One important challenge for a data-driven targeting strategy is that facilities change over time. Therefore, common machine learning strategies like the Random Forest may not adequately model the data when applied to the entire dataset. For this reason, a policy search method was implemented to read through the data chronologically and update parameters to approach an optimal policy.[13]

The policy function search iterates through the data to define a certain *policy*. A policy is a function of the input data that finds an optimal strategy for a given time. The model is derived from the following definitions and formulations:[14]

$X^{(k)}$	State Variable	Matrix of features that describe a facility at time k
${v^{(k)}}_{link}$	Dependent Variable	Boolean – whether an inspection led to enforcement
ϕ	Basis – Decision Var.	A set of functions that determine the strategy (policy)
$\prod_{\phi}(X)$	Policy	The strategy for decision-making
$\zeta(y;\phi)$	Objective Function	Sum of error or reward function: $\sum_{k} f\left({v^{(k)}}_{link} - \prod_{\phi}\left(X^{(k)}\right)\right)$

In this case, the formulation of the policy function is a linear equation of each time-variant predictor. Time was discretized into time-steps *k*, with each Δk equal to 91 days. Therefore, at each time *k*, the following raw data is computed for each facility:

days_since_inspection	Number of days since last air inspection
number_prior_inspections	Number of prior air inspections
never_inspected	Indicator: facility has never been air inspected
days_since_stack_test	Number of days since last stack test
number_prior_stack_tests	Number of prior stack tests
last_stack_test_result	Result of last stack test (pass/fail/other)
never_stack_tested	Indicator: facility has never had a stack test
days_since_rcra_inspection	Number of days since last RCRA (waste) inspection

[13] For more on Policy Search Methods, see Warren B. Powell, "Optimization Under Undertainty: A Unified Framework," *Wiley-Interscience*, September 18, 2017.
[14] This particular formulation of the model was created in consultation with Elahesadat Naghib, Ph.D. candidate in *Operations Research and Financial Engineering* at Princeton University.

number_prior_rcra_inspections	Number of prior RCRA inspections
number_prior_rcra_violations	Number of prior RCRA violations
average_rcra_compliance	$1 - \left(\frac{number\_prior\_rcra\_violations}{number\_prior\_rcra\_inspections}\right)$
never_rcra_inspected	Indicator: facility has never been RCRA inspected
days_since_enforcement	Number of days since last CAA formal enforcement action
number_prior_enforcements	Total number of prior CAA formal enforcement actions
last_penalty_amount	Last non-zero monetary penalty paid for non-compliance
average_air_compliance	$1 - \left(\frac{number\_prior\_enforcements}{number\_prior\_inspections}\right)$
never_enforced	Indicator: facility never had a CAA formal enforcement action
last_reported_emission	Last self-reported Hazardous Air Pollutant (HAP) emission (lbs)
emissions_trend	4-year trend in Hazardous Air Pollutant (HAP) emissions
missing_emissions_observation	Indicator: No self-reporting emissions data for the given year

Ultimately, these variables define our observation matrix $X^{(k)}$. However, in their current form, the features each have probability distributions that are quite variant in mean, range and deviation. Histograms for each non-indicator variable are reported in *figure (2)* to provide insight on the probability distributions of each predictor.

From the raw data defined above, phi-functions were defined to transform the raw-data into more a more balanced form, for analytic and numerical reasons. The phi functions are the basis functions upon which the observation matrix $X^{(k)}$. The ϕ-functions are reported below, and density histograms are reported in *figure (3)*.

ϕ_1	$\log(days\_since\_inspection)$
ϕ_2	$\log((number\_prior\_inspections) + 1)$
ϕ_3	$\log((last\_reported\_emission) + .0005)$
ϕ_4	$(emissions\_trend)$
ϕ_5	$\log(days\_since\_rcra\_inspection)$
ϕ_6	$\log((number\_prior\_rcra\_inspections) + 10)$
ϕ_7	$\log((number\_prior\_rcra\_violations) + 10)$
ϕ_8	$(average\_rcra\_compliance)$
ϕ_9	$\log(days\_since\_enforcement)$
ϕ_{10}	$\log((number\_prior\_enforcements) + 5)$
ϕ_{11}	$\log((last\_penalty\_amount) + 1)$
ϕ_{12}	$(average\_air\_compliance)$

ϕ_{13} $\log(days\_since\_stack\_test)$
ϕ_{14} $\log((number\_prior\_stack\_tests) + 5)$
ϕ_{15} $(days\_since\_inspection * (1 - never\_inspected)$
ϕ_{16} $(never\_rcra\_inspected)$
ϕ_{17} $(days\_since\_rcra\_inspection * (1 - never\_rcra\_inspected)$
ϕ_{18} $(never\_enforced)$
ϕ_{19} $(days\_since\_enforcement * (1 - never\_enforced)$
ϕ_{20} $(missing\_emissions\_observation)$
ϕ_{21} $\log((last\_reported\_emission * (1 - missing\_emissions\_observation) + 1)$

Now that each ϕ-function is set, we can define the model. Using iterative linear regressions, the algorithm finds an optimal set of coefficient parameters $\{\alpha_0, \alpha_1, \alpha_2 \dots\}$. At each time-step *k*, facilities receive a score based on their features that represents the relative likelihood of a CAA violation. Then, we compute which facilities are inspected between times *k* and *(k+1)*, and whether the facilities that are inspected have a subsequent formal enforcement action within 1000 days of inspection.[15] We call this response variable $v^{(k)}{}_{link}$. Comparing the computed scores to the actual response variable, the algorithm updates its policy to minimize error.

$$X^{(k)} = \begin{pmatrix} \phi_1 & \phi_2 & \phi_3 & \phi_4 & \cdots \\ [\vdots] & [\vdots] & [\vdots] & [\vdots] & \end{pmatrix}$$

$$v^{(k)}{}_{link} = \begin{bmatrix} 0 \\ -1 \\ 0 \\ 1 \\ \vdots \end{bmatrix}$$

At each iteration, we compute the changes in scores $\{\Delta\phi_1, \Delta\phi_2, \Delta\phi_3, \dots\}$ and changes in basis functions $\{\Delta\phi_1, \Delta\phi_2, \Delta\phi_3, \dots\}$, and use the stochastic gradient method updating process to scale our updating procedure:

[15] Defining the dependent variable $v^{(k)}{}_{link}$ was an involved process because there wasn't an explicit link in the data between inspections and legal enforcement actions. However, the 1000-day cutoff is not completely random – it is statistically derived. A histogram of the time-delay between inspections and enforcement actions reveals that 90% of delays occur within 1001 days. Additionally, air compliance personnel at EPA confirmed that enforcement actions almost never surpass a 3-year delay after an inspection occurs.

- Regress $v^{(k)}{}_{link}$ on $X^{(k)}$ to attain coefficients $\{\beta_0, \beta_1, \beta_2 \ldots\}$
- Compute discretized differential values $\{\Delta S_1, \Delta S_2, \Delta S_3, \ldots\}$, $\{\Delta\phi_1, \Delta\phi_2, \Delta\phi_3, \ldots\}$
- Use median $\frac{\Delta S_i}{\Delta\phi_i}$ to scale the learning rate θ for each coefficient α

The process to update alpha-values at each time-step is summarized with the following equation:

$$\alpha_i{}^{(k+1)} = \alpha_i{}^{(k)} + \left(\beta_i^{(k+1)} - \alpha_i{}^{(k)}\right) * \theta^{(k)} * \frac{\left(\frac{\Delta S_i}{\Delta\phi_i}\right)}{\left(\frac{\Delta S_{max}}{\Delta\phi_{max}}\right)}$$

In this case, we use learning rate parameter $\theta^{(k)} = \frac{1}{k}$

Iterating through the data starting in the year 2000, we hope to see signs of convergence in our coefficient parameters, which would signify a stable equilibrium relationship between the predictors used and non-compliant polluting.[16]

[16] The iterations begin in 2000 because at this point the quantity and quality of data markedly improves. That said, no data is emitted from the model. The year 2000 is simply when the algorithm begins to look *backwards* at the *history* of enforcement data.

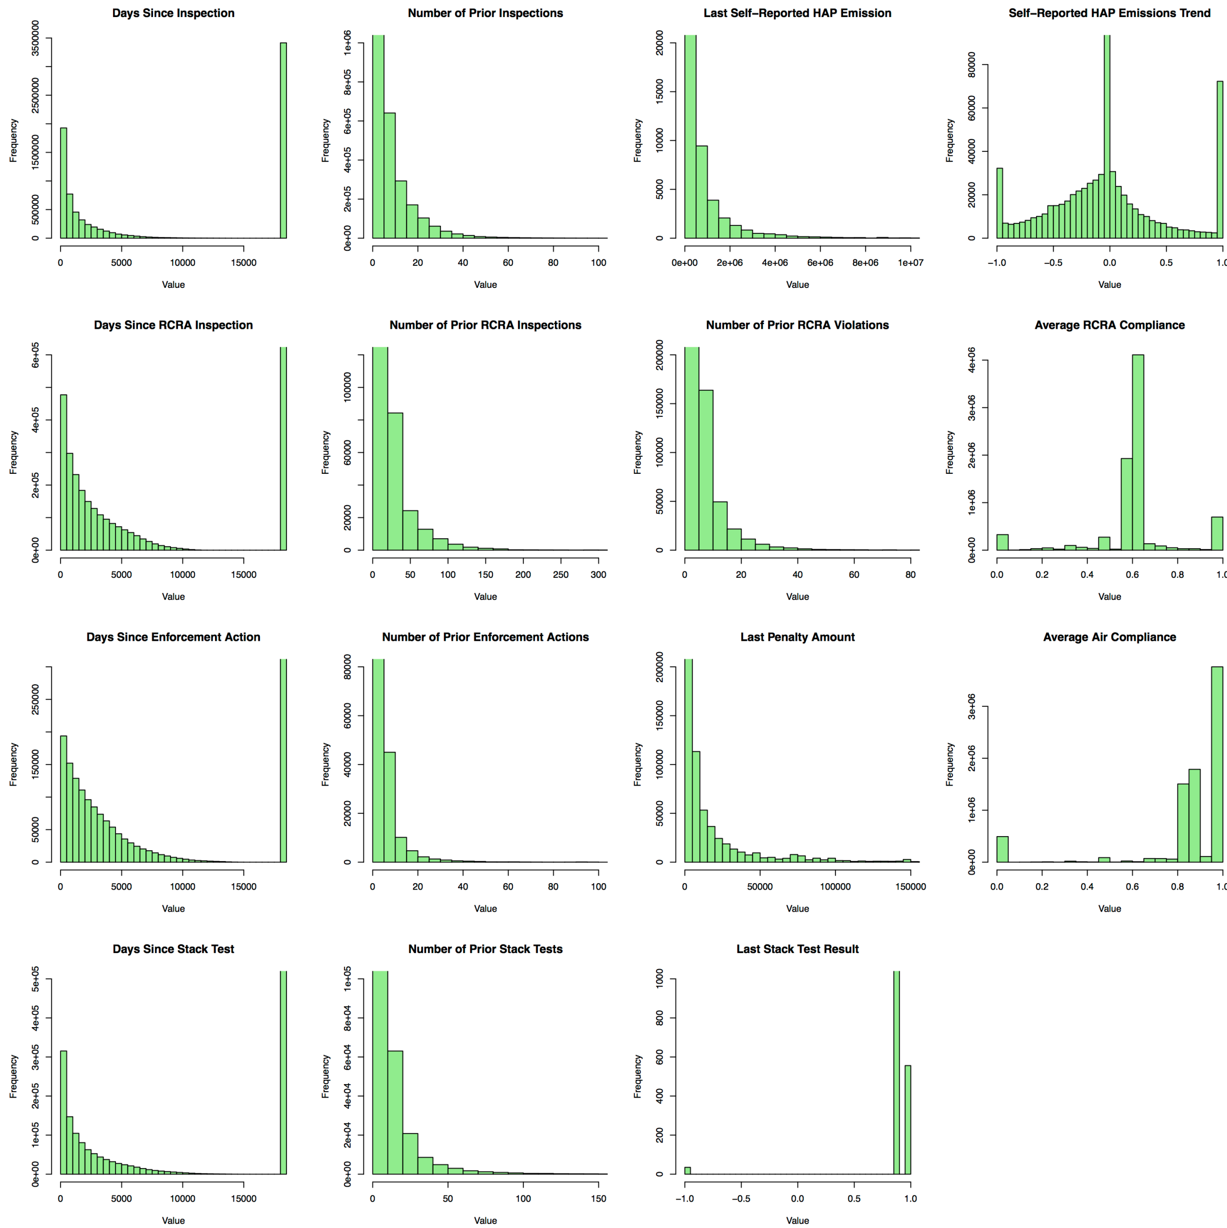

***Figure (2)**: Histograms of time-dependent variables in raw dataset*
Some plots have large outliers because of various imputation methods for missing values.
However, the linear model uses indicator variables to address missing data.

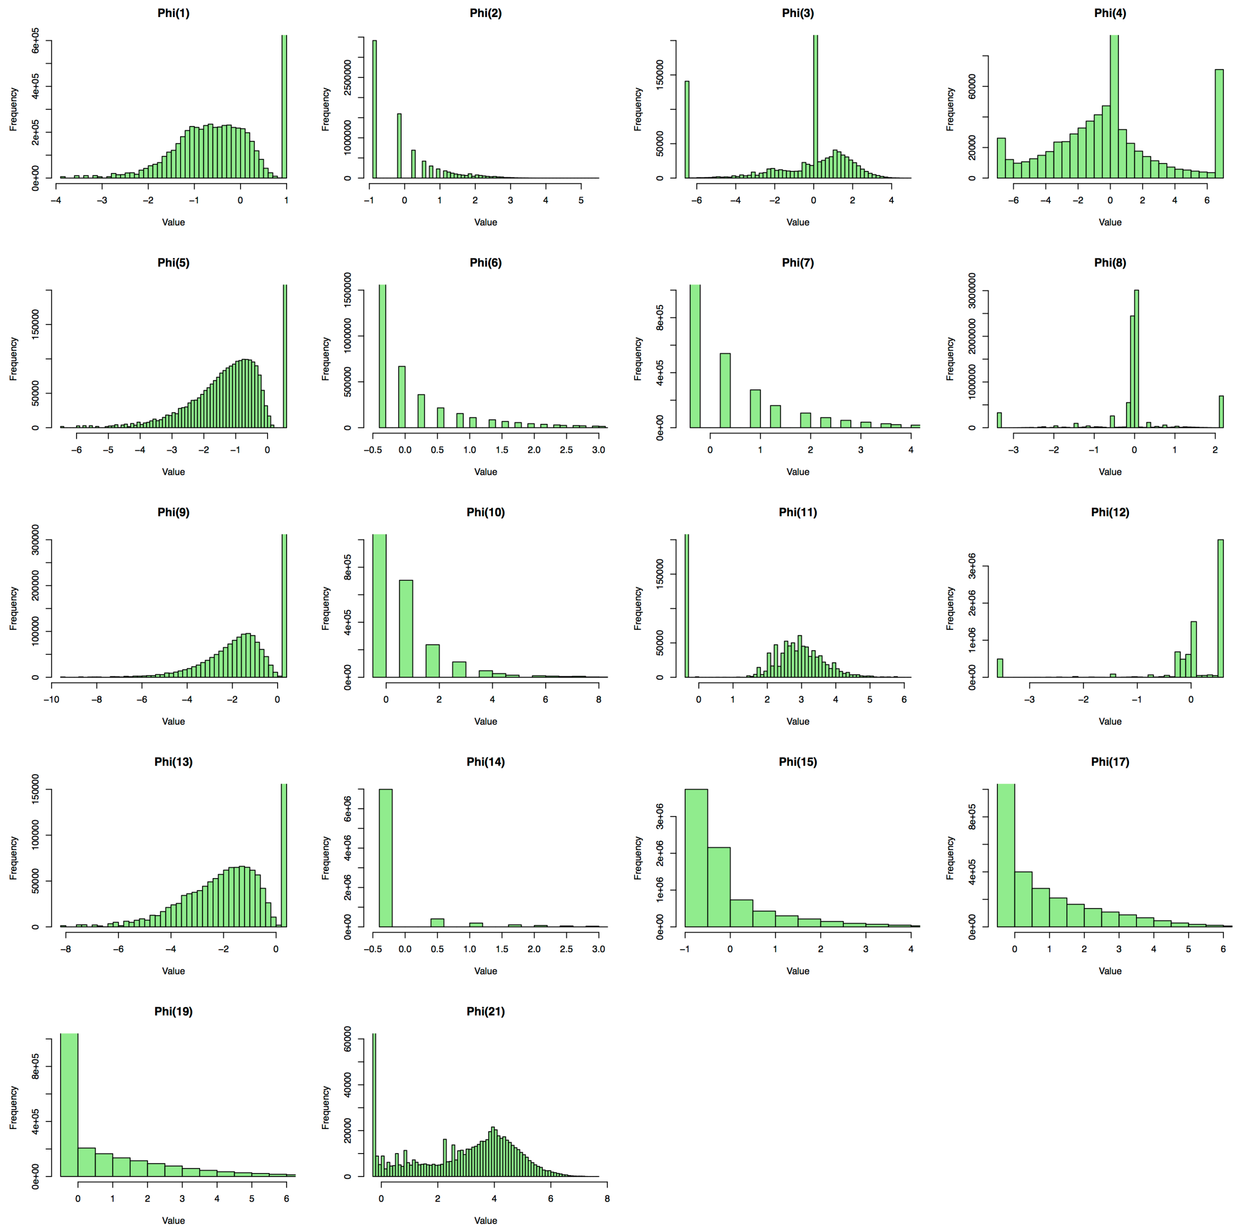

Figure (3)*: Histograms of ϕ-functions for each non-categorical variable in $X^{(\square)}$*
Each distribution is centered at 0 with unit standard deviation. Again, plots with large outliers are the result of missing values that ultimately do not impact the model.

Stationary Facility Characteristics

The stochastic gradient method and policy function described above were used only for the facility features that change over time. As mentioned, a proper algorithm for optimizing inspections must weigh 'exploitation' of findings with 'exploration' of sites that have never been inspected. To include these facilities in our model, we must devote a certain proportion of inspections at each time step to exploring new possible non-compliant facilities.

Since the facilities in question have never been inspected before, little is known about their history of compliance, pollution levels, or stack test results. In other words, the predictors used for the stochastic gradient method will not suffice in deciding which new facilities to inspect. Therefore, we must treat new facilities with a separate statistical procedure. One option is to simply take a random sample of new facilities to target in addition to those output by the policy function. But with tens of thousands of facilities that have never been inspected – many of which are tiny – a purely random approach would not be strategic.

Instead, traditional regression techniques were implemented to help choose which new facilities to inspect at each time step. Since we do have some data on new facilities – just not historical enforcement data – we can use statistical methods to inform target selection of new facilities. More specifically, using the pool of data for facilities that *have* been inspected, we can regress the dependent variable ${v^{(k)}}_{link}$ on facility features that are invariant with time. Then, with coefficients for each variable used, we can pick a weighted random sample from the pool of new facilities to inspect at time k. The predictors available for all facilities – both inspected and non-inspected – are listed below:

state	Which state the facility is located in
EPA_Region	EPA region that the facility is located in
naics_industry_sector	Which industry the facility belongs to. Variable uses first 2 digits of NAICS codes. Manufacturing and energy specified to subsector

facility_type_code	What type of facility it is - a corporation, county government, district, federal facility, tribal government, etc.
air_pollutant_class_code	Whether the air-polluting facility is registered as major, minor, synthetic minor, unknown, not applicable, other.
air_operating_status_code	Whether the facility is operating, temporarily closed, seasonal, under construction, etc.
indian_country	Indicator: Whether the facility is located in Indian Country
non_attainment_area	Indicator: Whether the facility is in a Non-attainment Area
percent_minority	The percentage of the population within a 3-mile radius that is not white (or is of Hispanic origin)
population_density	The number of persons per square mile within a 3-mile radius
environmental_justice	Indicator: Whether the facility is considered by the EPA to be in a community with Environmental Justice concerns
RCRA_permit_types	Set of 6 indicators: If the facility has a RCRA permit, then this field classifies their permit type - TSDF, LQG, SQG, CESQG, etc.
pollutants	Set of 46 indicators: Whether the facility is a registered polluter of the top 46 most common pollutants - benzene, xylene, CO2, etc.
onsite_methods_codes	Set of 7 indicators: What processes the facility uses in emitting air pollutants - flare, condenser, scrubber, absorber, etc.

To create weights for sampling the non-inspected facilities, a random forest machine-learning model was used. The random forest is conveniently equipped to deal with categorical variables, so it was a good option for this particular dataset. Still, it may be possible to use regularized regression techniques at this step of the procedure.

By regressing inspection results on stationary (unchanging) facility characteristics, our goal is to inspect new facilities that are likely candidates for non-compliant pollution practices. The algorithm devotes a certain proportion $\gamma = \frac{1}{4}$ of inspections to exploring new facilities at each time *k*. That said, it is important to note that this procedure would be a departure from the history of EPA decision-making, which has not quantitatively balanced the bandit trade-off issue. Therefore, we simply begin including a proportion of new facilities *now*, and the algorithm begins the ‘exploration’ procedure *from now on*.

Findings

The linear policy function search model was implemented and alpha values were computed for time-dependent features ϕ. The stochastic gradient method yielded alpha-values that show evidence of convergence for $\{\alpha_1 \ldots \alpha_{21}\}$. The computed alpha-values at each time-step k of the process are plotted in *figure (4)*.

The results in *figure (4)* indicate that the policy search process reaches an asymptotic solution for almost every alpha-value. There are some time-dependent changes that are not fully asymptotic – visible in the plots for α_8 and α_{14} in particular. In these plots, there seem to be local minima and maxima where the variable begins to switch direction. Still, these plots show some evidence of leveling off as time passes.

Ultimately, the algorithm produced higher scores for facilities that it deemed most likely to be in violation of CAA standards. The random forest output added a portion of never-before-inspected facilities for EPA targets in 2018. Together, the list of facilities will help inform EPA Region 2's 2018 list of target facilities. While quantitative results, including the list of recommended facilities, are confidential, the methods discussed may be applied to publically available EPA data to attain identical results.

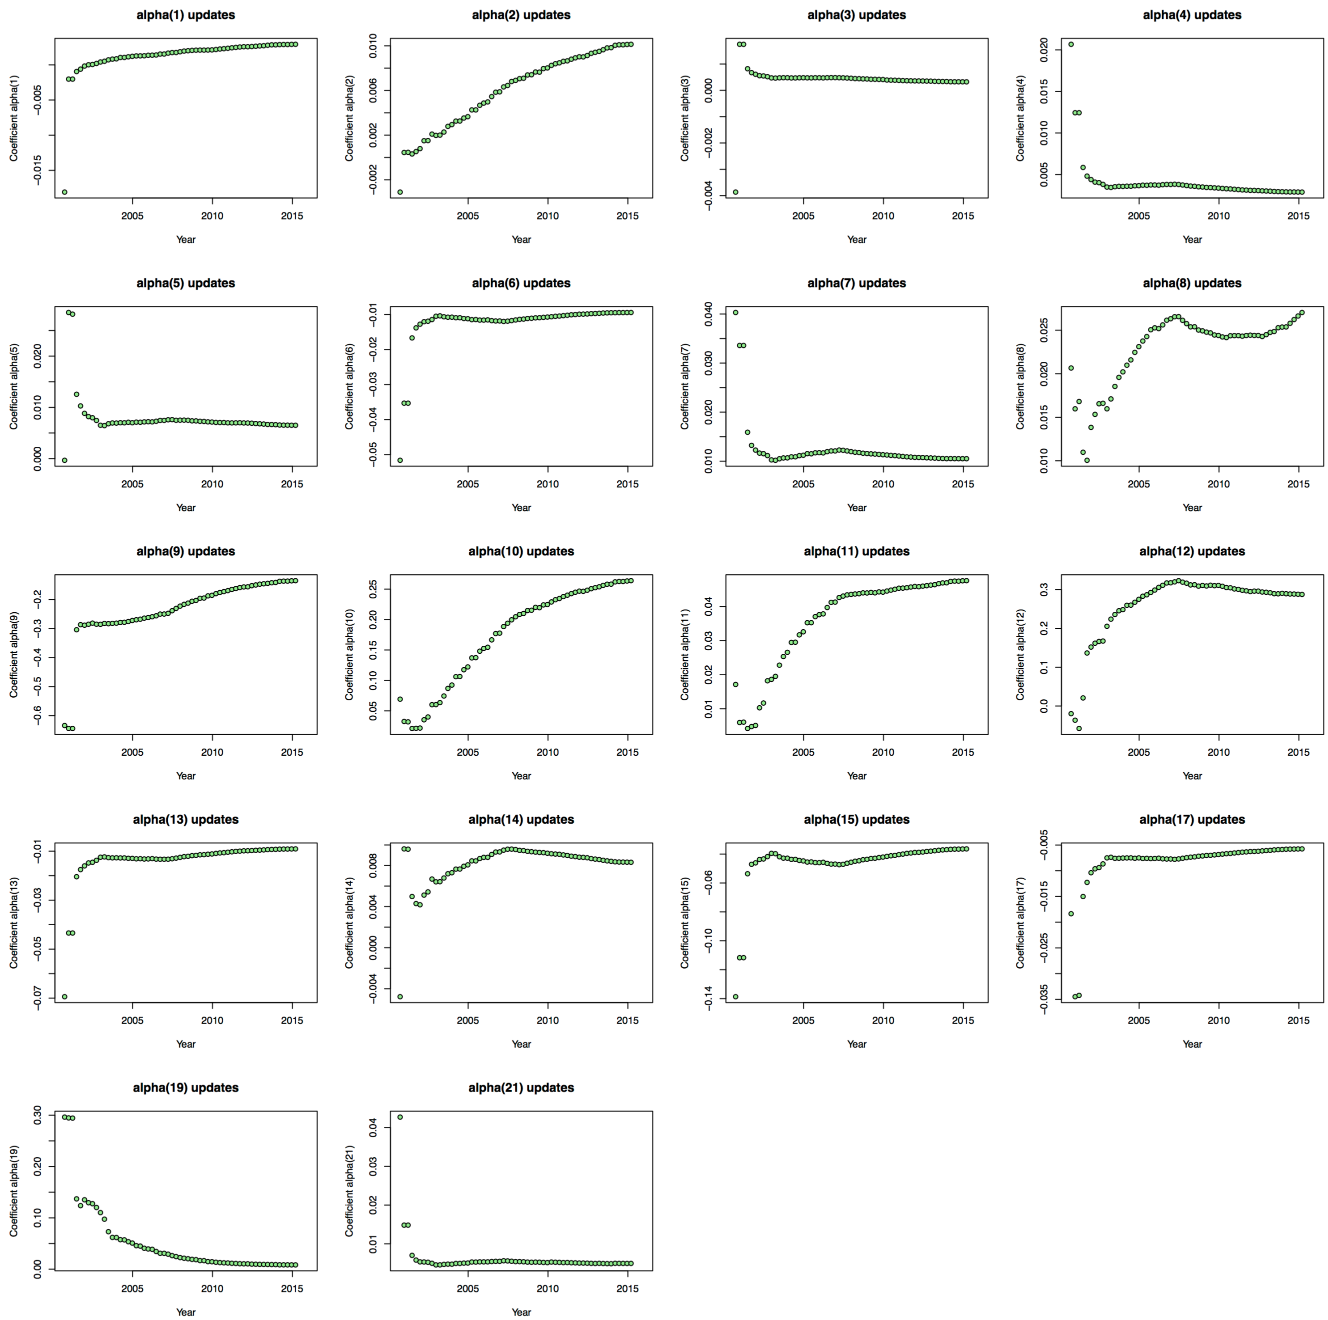

Figure (4)*: Alpha-value definitions as a function of time.*
Convergence implies that the model approaches a stable point for large k.

Further Study & Conclusion

This study suggests that the policy function search is a viable option for targeting and allocation strategy. The linear model and stochastic gradient method were implemented as a simple example of the policy function search in practice. This approach holds certain advantages over existing methods, including its treatment of time-dependent data and never-inspected facilities. That said, newer and more advanced algorithms can further improve the targeting strategy.

In particular, non-linear models may prove advantageous over the model produced in this study. Regularized regression techniques – like a ridge regression – at each time-step could possibly improve the policy. More complex models could further benefit the model – a deep neural net (DNN) policy could be especially powerful at classifying targets. This would involve many more parameters and a much longer computation time, but these drawbacks could pay off in model performance.

Another outlet for further study is applying the machine-learning techniques discussed to other regulatory operations. Other environmental laws – including Clean Water Act and Safe Drinking Water Act – do not currently use machine-learning techniques to inform targeting strategy, to my knowledge. Similar techniques may be applied to health and safety inspections in other areas of government as well.

Despite the complexity of a quantitative approach to inspections strategy, the potential for pay-off is significant. A move towards smarter, data-driven operations would not only save time and resources for EPA inspectors, but would help protect U.S. residents from dangerous environmental pollutants. Discovering violations in environmental law helps minimize human exposure to hazardous air pollutants and chemicals that threaten our environment. It is the goal of this study to inform policy and positively impact U.S. air quality and general well-being.

Acknowledgements

Professor Robert Vanderbei, *Advisor*

Elahesadat Naghib, *Advisor and Mentor*

Daniel Teitelbaum, *EPA Data Team Supervisor*

Stephanie Wilson, *EPA Data Team*

This paper represents my own work in accordance with Princeton University's Rights, Rules and Responsibilities