The spatial evolution of the MA C&I efficiency landscape, 2011 - 2014

Richard Crowley, Greg Gronski, Anthony Davis, DNV GL
Whitney Brougher, National Grid
Riley Hastings, Eversource Energy

ABSTRACT

Understanding customer engagement in efficiency programs—who is and is not engaged, and where they are—is a powerful means for Program Administrators (PAs[1]) to better identify and target commercial and industrial customers. To this end, the Massachusetts PAs are expanding their use of geospatial analysis to explore historical trends and identify new opportunities. The PAs developed and maintain a statewide C&I energy efficiency evaluation database that contains cleaned and standardized program tracking and billing data for analysis and reporting. The database currently contains four years of data aggregated across PAs, including recently added third-party tax assessor data, which increased its geospatial capabilities.

A key element of the database is the geographic data integration and visualization that identifies energy efficiency trends within and across PA service territories. The MA PAs leverage the geographic elements of their data and third party datasets to spatially assess participation and savings, as well as individual building types and energy use intensities. This paper presents examples of the MA PAs applying geospatial analysis and spatial data integration to identify and target specific areas and customers with high savings potential. For example, newly developed EUI maps empowered the PAs to identify areas of higher or lower energy use by industry sector across the state. This presents a more comprehensive picture of the efficiency landscape and illustrates how geospatial analysis expands the PAs' knowledge of individual accounts and ability to more effectively evaluate what is happening, where it is happening, and why it is happening.

Introduction

Quantifying energy efficiency accomplishments is not new. Since programs began offering energy efficiency measures, implementers, strategy and program development teams, and regulatory bodies have all wanted metrics to measure how these programs are performing. This paper focus on the commercial and industrial (C&I) space. Classically, this has been done by looking at metrics such as the number of possible accounts that are participating (e.g. participation rates) and the amount of energy those accounts save relative to their consumption (e.g. savings rates). This data is neatly summarized and presented in tabular or graphic format. Using these statistics, statements are made about PA performance, which sectors offer more savings potential and the general state of program evolution over the past year; this in turn is used to tailor future offerings and approaches.

There is a critical shortcoming to these more classical views of the data—they ignore the underlying spatial data! Customers do not exist as a record in a data warehouse, or a fractional percentage of a bar chart. They exist as a physical entity somewhere in space that consumes a

---

[1] In Massachusetts, the term Program Administrators is used and encompasses both the traditional utilities (National Grid, Eversource, Unitil, Columbia, Berkshire, and Liberty) as well as the program administrator Cape Light Compact.

product—be it energy, efficiency measures, or incentive dollars—supplied by the service providers. The classical presentation lenses effectively assume that the underlying customers represented are spatially homogenous; implicitly assuming that that the geographic space that customer exists in is not an important piece of information for understanding what is occurring. It is unlikely that today's savvy stakeholders would truly believe that there is spatial homogeneity amongst customers. But the reality is that the metrics we as an industry have historically relied on do not fully leverage the geography to help better inform and guide decisions. By explicitly integrating geography into how energy efficiency programs are assessed it is possible to achieve a much deeper understanding of participation and savings potential as well as new views into emerging areas such as demand reduction opportunities in congested transmission areas or areas with higher peak energy loads. This type of tailored geographic insight can aid in more targeted program offerings, and can potentially even facilitate assessing and quantifying the value of offerings beyond the simple kilowatt hours or therms saved – and extending even to service territory wide grid and individual customer non-energy benefits![2]

In this paper we first provide an example of the classical non-geographic way of assessing program participation data and the conclusions that can be drawn from it. Next, we provide the same data while explicitly factoring town level geography to illustrate the variability in participation across service territories and offer observations on how geography can facilitate cross fuel and PA views of otherwise unlinked data records while presenting for the first time the combined gas and electric savings across MA towns in MMBtus. We then further pursue different ways of dissecting town level savings rates using absolute, ratio, and normalized scales to illustrate how geographic visualizations can offer tailored insight into sub-territory trends and changes in savings. Finally we explore how the underlying point level data can be leveraged to present detailed insights into sub-town level trends through spatial interpolation and grid-cell aggregation methods. We conclude with context on the future promises, caveats, and potential concerns that spatial data and analysis methods can offer to help program providers understand how their energy efficiency programs are impacting their service territories.

Discussion

The national leader in energy efficiency for the fifth year in a row, Massachusetts has been using geography for several years as part of their ongoing efforts to quantify, understand and build on historical success and lessons learned (DNV GL 2012, DNV GL 2013, DNV GL 2014, DNV GL 2015, DNV GL and NMR 2015). Most recently, they are making both C&I and Residential savings and utilization data available for public stakeholders though the Mass Saves Data web interface, a new repository for statewide utility data (Mass Save Data). The Massachusetts PAs are not alone in understanding the value of geography; more and more frequently utilities across the nation are utilizing tools like geo-targeting of customers as a way to customize program outreach and strategies to achieve higher participation and savings rates from customers that have traditionally been harder to reach, and, taken in isolation, may have a higher cost to engage (Novie and Teng 2016; Crowley and Brougher 2014; Freeman et al. 2014).

[2] As an example, the value of placing a series of new energy efficient water heaters on a congested line may have the same energy savings as on one that is not congested, but the demand reduction may have a different value to the program administrator or merit a different incentive level. Extending this example further, a reduction in service calls – a non-energy benefit to the facility – may have a greater value in areas where there is not a service facility nearby then in a place where there are many competing providers.

What makes Massachusetts a national leader in how they apply this data is that the Program Administrators have compiled their standardized data into an Evaluation Database. This database is updated annually and includes a geographic element that provides granular insight into the spatial evolution of the underlying confidential customer and program records.[3]

Historical Metrics: Participation Graphs

Consider the historical participation and savings metrics for the Massachusetts electric PAs using the classic view with the data in a bar chart (Figure 1). The left hand chart presents both the participation for each PA within a given year at the account level as well as the unique account participation across four years of time.[4] At the PA level, there has been continual improvement in C&I participation each year since 2011. At the unique account number level nearly all the PAs have engaged around one-in-ten of their C&I accounts in the last four years. The right hand chart shows the consumption-weighted participation which accounts for the relative sizes of the participant accounts rather than treating each participant as equal.[5] Weighting the participants by the consumption they represent those one-in-ten accounts that have participated represent between 40 and 50 percent of the PA's overall consumption. A lower account participation rate, but high consumption-weighted rate is indicative of larger accounts participating. Inferring deeper from the chart, the increase in account-level participation in 2013 and 2014 combined with the relatively stable year-over-year consumption-weighted participation suggests that more small accounts have begun to participate in the programs.

All of this is useful insight and suggests strong (and growing) C&I efficiency programs and offerings that are engaging new customers—but from the implementation standpoint, these views are somewhat limited. Stakeholders receive insight into historical population level accomplishments, but do not know much about where the customers that have been reached are, or more importantly from an implementation and strategy standpoint, where to look next.

[3] Data in the Evaluation Database is confidential and maintained on behalf of the Massachusetts PAs by DNV GL as the third party data manager.

[4] At the year level, accounts can participate each year and count towards that years participation rate (blue bars), while at the population level (gray bar) each account only can be counted once against participation in the time series. If a single different account participated every year from a population of ten accounts, the annual participation for each year would be 10% (1/10), while the four year participation would be 40% (4/10). If the same account participated every year, the annual participation would still be 10% (1/10) but the four year participation would only be 10% (1/10).

[5] Extending the example from the lasts footnote, if the four accounts each represented 20% of the population consumption, then the consumption-weighted participation in each year would be 20%, while the four year consumption weighted participation for the 4 unique accounts would be 80%.

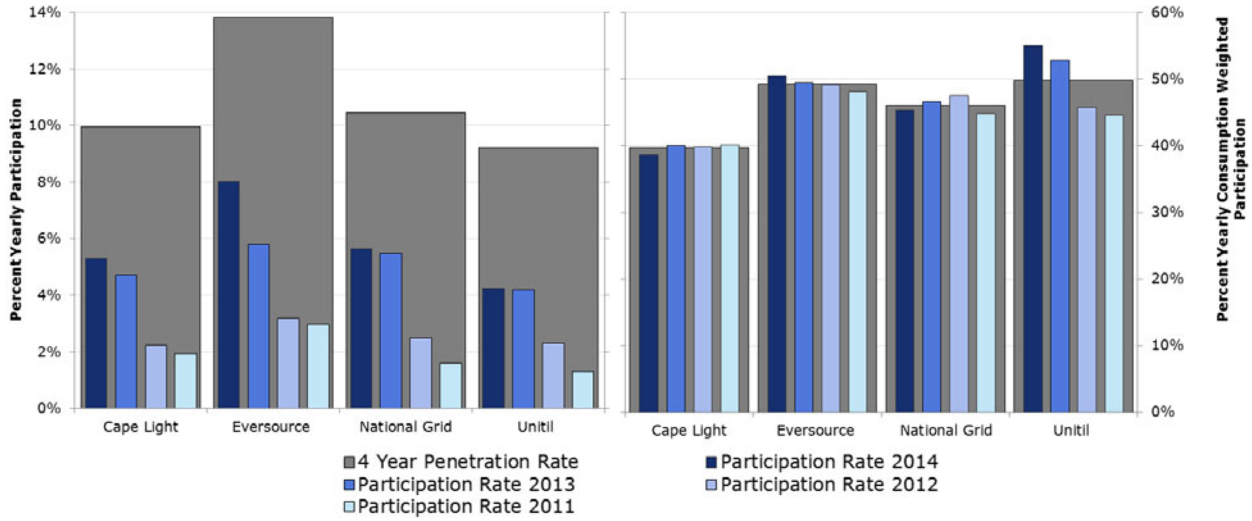

Figure 1. 2011-2014 Participation (left) and Consumption-weighted Participation (right) for the MA Electric PAs - Traditional View

Participation in the Geographic Context

Figure 2 shows the same underlying data as Figure 1, but summarized geographically at the town level.[6] It is immediately apparent that within the towns the PAs serve the assumption of spatial homogeneity is inaccurate! Regardless of how we weight it, participation is highly variable across the state, as well as within the different program administrator territories.[7] As the visual deliverable from geographic analysis maps can communicate large amounts of dense data in an efficient picture that greatly facilitates interpreting the results and leading to actionable insights. For Massachusetts, the same town-level data presented in Figure 2 would require 351 rows per table.

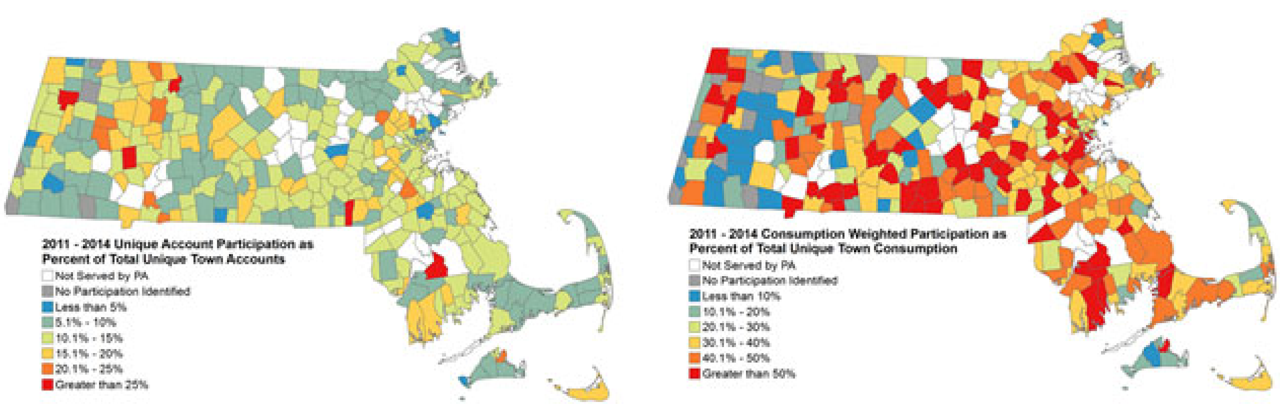

Figure 2. 2011-2014 Account and Consumption-weighted Participation for the MA Electric PAs - Geographic View

[6] Throughout the bulk of this paper, town level is used as it is an intuitively relatable geographic boundary, and with only one exception for electric and relatively few for gas, each town has a PA assigned to it making for clear nesting between town- and PA-level summaries.

[7] A map of the program administrator territories can be found at http://www.mass.gov/anf/research-and-tech/it-serv-and-support/application-serv/office-of-geographic-information-massgis/datalayers/pubutil.html

Depending on the participation metric that is being measured the geographic view offers important the insight that quickly direct the focus to areas of interest.[8] Different strategies are appropriate for different goals, and the maps help to prioritize where those strategies can be applied.

If absolute participation matters—say due to a target goal that each town has at least 10% participation over the four year period—then focusing on towns where there is low participation by using local contractors and community resources in a targeted engagement push for direct install measures at smaller customer sites could be a strategy option. The map quickly shows that that a program and vendor focus on towns in the southeast, southwest, and northern-central parts could quickly improve the town level participation ratios. If the desire is to ensure that the largest accounts in each town have already been engaged, then consumption-weighted participation map illustrates that several western towns have—at least by the town standard—large accounts that have not been engaged.[9] An even more interesting one is identifying towns that have either high participation rates and low consumption-weighted participation as this dynamic suggests there may be larger accounts (relative to the town) that have not been engaged and could represent strategic outreach points for local vendors.

Geography also offers a standardized way to aggregate and report data that may be more effective at communicating a comprehensive picture than current methods. Each Massachusetts PA has its own customer tracking and billing systems, each with its own unique account numbers. These numbers do not link across PAs and also may not link across electric and gas accounts within individual PAs; this means some facilities with a gas account with one PA and an electric account with another cannot be evaluated as one "customer" since there is not a shared key or way to link the two accounts. By aggregating data to the town level it is possible to look at the more holistic picture aggregating electric and gas billing and tracking data and using MMBtu—as the next series of analyses do—rather than having to look at therms or kilowatt hours individually.[10]

Mapping Historical Savings Ratios Across Fuels

Figure 3 shows the savings ratio at the town level for the time series 2011 through 2014. The "savings ratio" is defined as the sum of savings within a town divided by the total sum of consumption within that town. Areas in blue have a lower four year combined savings ratio (<1%), while areas in red have a higher four year savings ratio (>5%) indicating that a larger amount of savings relative to the town's consumption have been undertaken. At this overview level, maps such as Figure 3 can help direct focus on areas where there may be opportunities for towns to generate more savings relative to town consumption (such as north of Boston) and can help to quickly assess if there are towns that at the population level may be lower priorities for outreach. Although it was not assessed at the town level for the report underlying this paper, this type of insight may also offer an opportunity to identify areas where cross fuel or PA

[8] in this town level example they are "how many customers have participated (left map) vs. "how much of the town level consumption do participants represent"

[9] "By the town standard" is a critical distinction – the "large" accounts may actually be small relative to the full state population, and this is often the case in the rural towns. However, while they are small relative to Boston, these accounts represent the "best" accounts from a consumption standpoint for local vendors and implementers.

[10] Extending on this idea, it also is possible to start looking at spatial statistics like regional hot spots and clusters. These are powerful tools for trend identification, but are not discussed in this paper.

coordination has been particularly effective and can be adopted by other program administrators.[11]

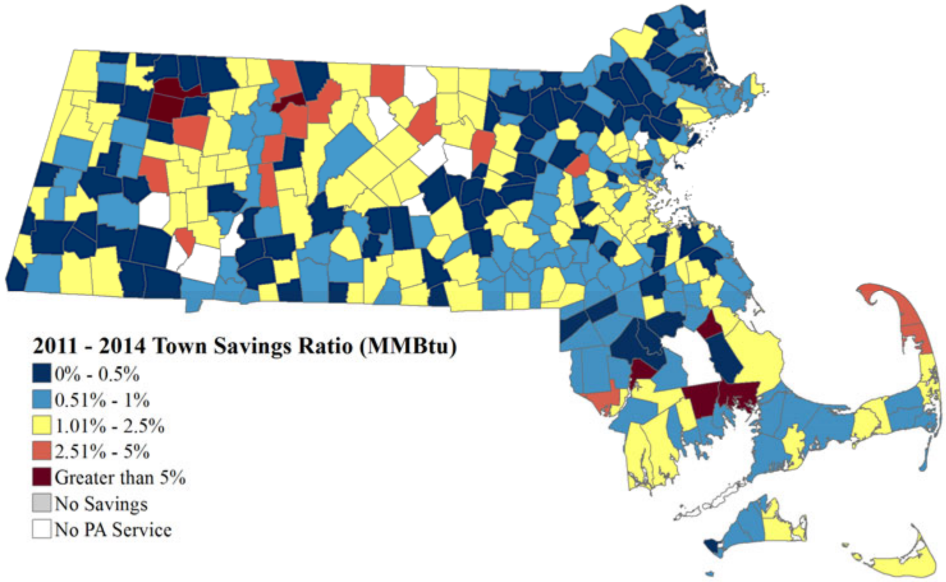

Figure 3 2011-2014 Combined Gas and Electric Town Savings Ratio (MMBtu)

Going Beyond Ratios – Normalized Time Series Trends

A limitation of the town level savings ratio is that it does not capture scale of savings relative to scale of savings in the other towns – this would require a different normalization scheme for a different question.[12] It also does not provide insight into any statistically significant trends that may be occurring. The aggregated time-series data must be broken into discreet annual datasets to accomplish this as town's individual year's savings ratios are highly variable as different customers engage in programs. Understanding where and how these trends are evolving can identify where there have been successful changes in program marketing for towns, shifts in vendor strategies, holes in vendor service coverages, or any number of other non-data drivers. As an example the most recent Massachusetts C&I Customer Profile found town level participation in 2014 is generally higher than historically across the state confirming in part the successful impact of Massachusetts's successful upstream lighting program.

Metrics like savings ratios offer a good success story when the majority of towns have increased their participation since 2012, but taken in isolation the percentages alone do not facilitate assessing how a town's change in the savings ratios fit into the overall picture relative to the other towns. E.g. improving from 1% to 2% is great; a little less interesting though if every other town experienced the same increase though. Visualizing the standard deviations provides a statewide apple-to-apples view across time and facilitates the quick identification of regions

[11] It was however assessed and presented at the PA combination level, and there were some notable differences for gas PAs. The full report can be accessed at

[12] As an example, using the ratio of savings to reach the conclusion that that Mt. Washington (population 167) represents "more" potential than Boston (population 645,966) just because it has a lower ratio would (obviously!) be incorrect. However, the conclusion that Mt. Washington's savings ratio is lower than Boston's over the four year period and therefor those accounts as a town level population have saved less relative energy is accurate. It is very important with geographic analysis to understand the question you are trying to answer when assessing the maps!

within the PA service territory that are contributing higher or lower savings within the year relative to other areas of the state. Figure 4 presents the standard deviations for the town-level total MMBtu savings ratio for the individual years 2012 and 2014.[13]

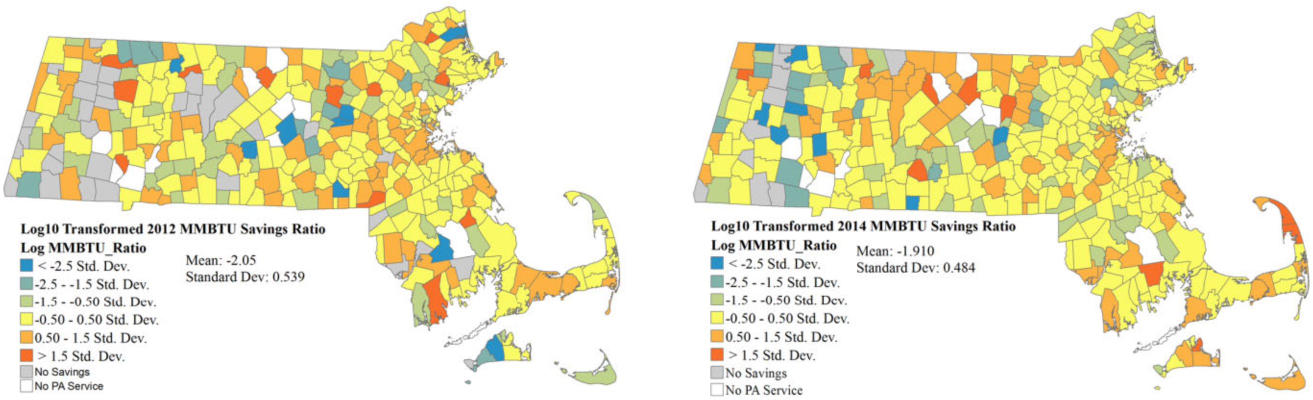

Figure 4. 2012 and 2014 town-level standard deviations for log10 transformed MMBtu savings ratios

Several towns in the central southern portion of the state, including Ludlow and Ware, have savings ratios in 2012 and 2014 that were below their respective year's mean ratios. Findings such as these may help to identify areas where vendors may be overlooking potential savings, such as installing lighting measures but not considering gas HVAC equipment that may be in harder to access areas of buildings. In this manner, maps can also help identify markets where the existing vendor network may not reach.

Alternately, spatial analysis can highlight where programs—particularly new offerings—appear to be reaching previously less-engaged customers. In 2012, New Salem and the surrounding towns in the central northwest exhibited no savings relative to town consumption; however by 2014 they were exceeding the within year mean. A continued scale-up in upstream programs, lighting in particular, may have helped shift the landscape by providing a mechanism to engage these smaller, more geographically remote accounts. In other areas, GIS can help us identify new accounts that have come online that—though smaller relative to the full PA population—represent in aggregate a sizable new amount of consumption for a single town by driving down the town savings ratio.

Time Series Changes in Town Level Share of Savings Geographically

Another question is "where are the savings coming from?" Taken in a single year, the question of savings origin is not in itself a particularly interesting one: chances are the largest share of savings will correlate with the areas that also have the largest amount of consumption.[14] However, the addition of the time element to the maps allows stakeholders to see where changes in the shares of savings occur. This presents a more interesting view of the data and can form the basis for different unique ways of analyzing data (Davis, Crowley and Gautam, 2016). Share of savings is defined as the total of a town's savings divided by the total statewide savings.

Figure 5 shows the absolute change in statewide savings contribution, in MMBtu, in 2014 relative to the statewide savings contribution, in MMBtu, in 2012.While the largest share of the

[13] Savings ratios were log10 transformed to bring to a "normal-like" distribution.

[14] In Massachusetts the gross town level savings generally do correspond with town level consumption.

savings annually may come from the largest areas of consumption, the time series data indicates that Boston and the surrounding towns have actually decreased their share of the statewide savings since 2012; while towns in the central and northern portions of the state have generally increased their share of savings. This suggests that as gross overall savings has increased each year, the PAs are also finding ways to engage customers in more rural areas of the state even though the population centers like Boston, Worchester, and Springfield are major sources of gross savings.

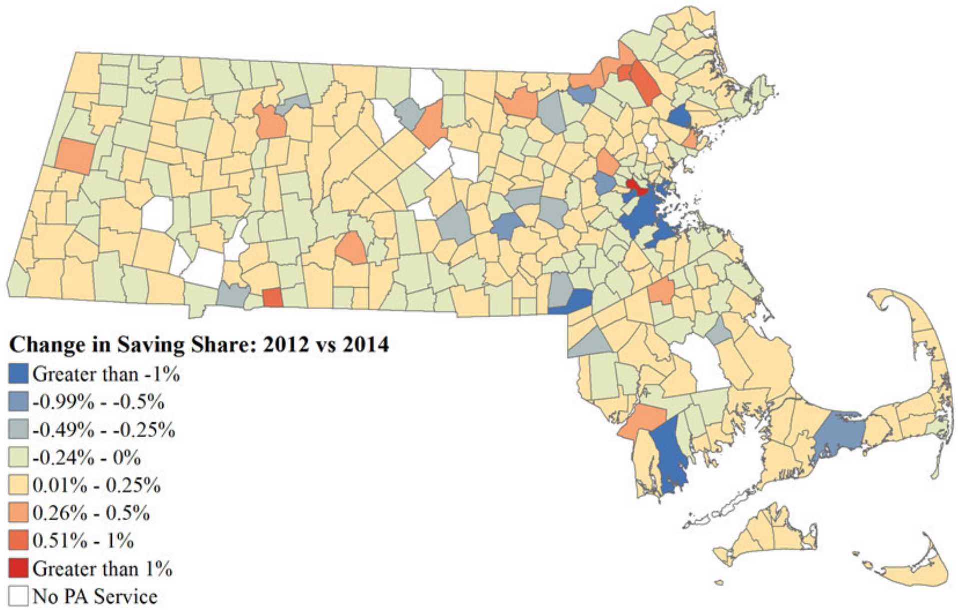

Figure 5. Change in Savings Share from 2012 to 2014, MMBtu

All the looks at data from a purely visual perspective are useful—but what if we wanted to go deeper still? From a strategy perspective, we might be interested in knowing whether there are any clusters of towns where the changes in the metric are statistically significant. If we have knowledge of program or marketing changes in those areas, we can test to see if the changes have had an impact and use the lessons learned to inform how programs are implemented. Even in the absence of specific program knowledge, this type of statistical testing over time can help identify areas where vendors may be finding particularly effective ways to approach customers – this type of insight is often obscured in data.

Cluster analysis and hot spot analysis—particularly when used in tandem as hot and cold spots can contain local variations of the opposite trend—are a useful tool for assessing if and where location variations are statistically significant.[15] These analyses were conducted on the data in Figure 5 and found that the majority of the changes are not statistically significant relative to neighboring towns.[16] However, the changes north of Boston indicate that some of the participation changes from 2012 to 2014 were statistically significant at or above the 95% level, with the cluster analysis identifying Lawrence in particular as an important element of the

[15] E.g. a "hot spot" may actually contain towns that on an individual basis had a decrease in savings. This decrease may or may not be statistically significant, and the cluster analysis can help clarify this.
[16] This is not to say there are not statistically significant through time, this paper does not address that. Merely we are saying that although many towns in Figure 3 had a change in savings, the results do not suggest that the change was statistically different then what was experienced by neighboring areas.

cluster. Conversely, Boston and the cities south of Boston had statistically significant changes in participation and cold spot clusters at or above the 95% level. This is not to say that these cold spots are not critical drivers of savings—Boston, for example, contributes nearly a fifth of all savings and around 15% of all MMBtu consumed. It may be that since share of savings is a zero-sum game[17] the decline is simply that these towns are bearing more of the impact of changes in savings dynamics across the state, but it could also be indicative that savings as a share of statewide total in these towns are declining as consumers have already undertaken low hanging projects. Regardless of the driver, these types of time series analyses of geographic trends of hot spots and clusters can ensure that strategy and implementation teams know which areas in their territory are undergoing statistically significant changes. Armed with this knowledge, stakeholders can then best decide how to react.

The geographic views are not restricted to just the town level. These views help focus on the evolution of the landscape in Massachusetts and offer insight into where large changes have occurred in participation and savings.[18] In this regard, the town-level maps and different metrics serve as a guidepost into the more granular underlying account point-level savings and consumption data, helping strategy and implementation teams quickly assess and discuss areas that may merit a deeper dive into the data. However, the town level maps do not offer insight on the specific accounts or detailed sub-town areas where groups of these accounts exist—this type of insight requires point level data.

Account Level Visualizations: Energy Use Intensity and Within Town Dynamics

Point-level account data can provide guidance for even more targeted actions. It is beneficial to understand where there is opportunity; it is even better to know details about the underlying opportunities prior to heading out into the field. Insights into point-level details like the potential energy use intensities of specific buildings in an area, the types of measure that have been installed, and variables like building vintage as a proxy for building insulation standards at time of construction can help increase the odds of identifying savings. Figure 6 (next page) presents an inverse distance weighted interpolated representation of statewide kWh energy use intensities (EUI) for all accounts that could be matched to tax data. [19],[20]

Metrics like EUIs can offer a normalized view of the energy efficiency potential and identify areas of the state that may have higher or lower locational energy use, while continuing to preserve the underlying customer confidentiality by obscuring the individual data points. By

---

[17] If one town increases its share of savings, others by definition decrease theirs.

[18] By extension, they also confirm that spatial homogeneity does not exist within PA territories.

[19] Per ESRI "*Inverse distance weighted (IDW) interpolation determines cell values using a linearly weighted combination of a set of sample points. The weight is a function of inverse distance. The surface being interpolated should be that of a locationally dependent variable*" For the purposes of this graphic, we assume that that EUI is a spatially correlated variable, this is an oversimplification of the true landscape. However, in the absence of data on behavioral, occupancy, and technology within individual location it provides a transparent way to model out a low granularity visual representation of the landscape as a starting point for discussion. Additional details on the caveats around this map can be found in the *2014 C&I Customer Profile* on the Massachusetts Energy Efficiency Advisory Council website at http://ma-eeac.org/.

[20] Per ESRI "*The best results from IDW are obtained when sampling is sufficiently dense with regard to the local variation you are attempting to simulate. If the sampling of input points is sparse or uneven, the results may not sufficiently represent the desired surface.*" This limitation is particularly pronounced in more rural areas of western Massachusetts where there are far fewer points to use. The result is that with fewer points the influence of any one point carries more of an impact than in densely populated areas like Boston.

understanding where there are higher EUIs stakeholders can focus on groups of customers to better understand drivers behind those energy use intensities and opportunities to capitalize on savings. One such example would be to find a series of office buildings in one town consistently having a higher EUI than similarly sized buildings in a neighboring town, and upon further examination learn that the higher EUI buildings are consistently older, suggesting that building shell measures may be a more attractive offering to these accounts.

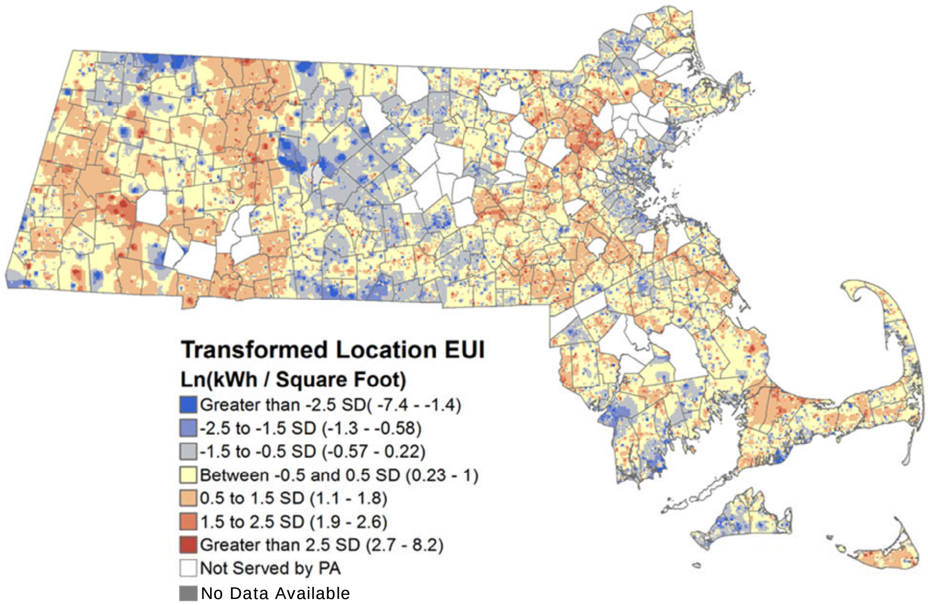

Figure 6. IDW interpolation of statewide energy use intensity for the 2014 C&I landscape.

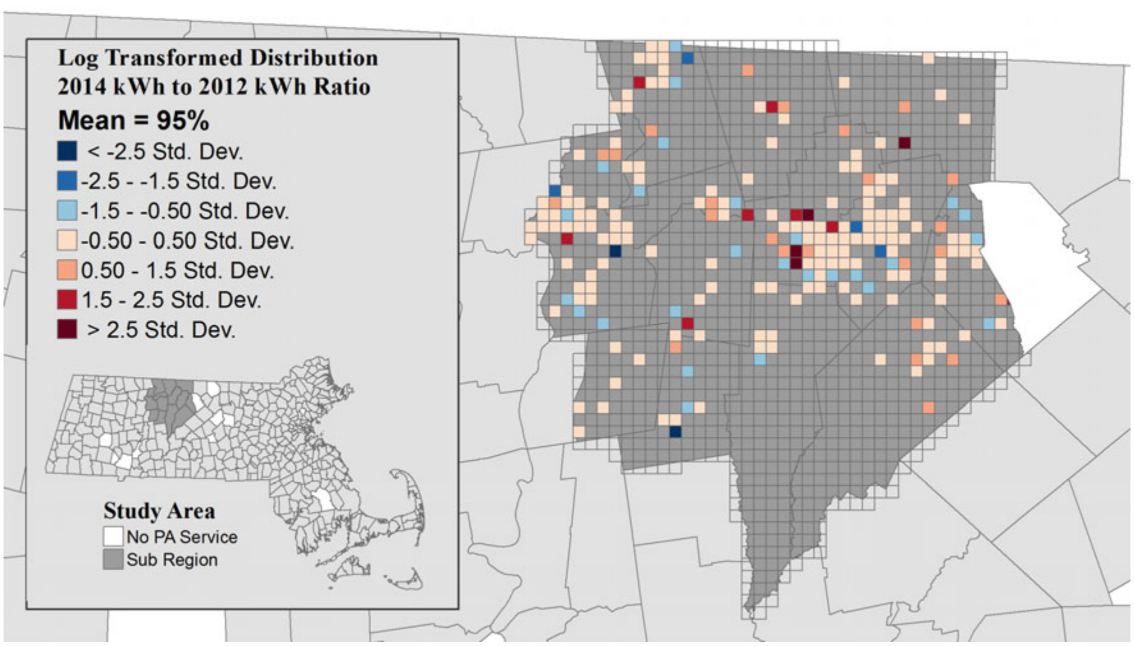

Figure 7. Example 1km cell kWh 2014 consumption differences from 2012

Figure 7 above, provides an exciting example on how time series data, even when not linked at the account level, can be used to help identify target areas at an extremely granular level. A subset of electric accounts in the study area is presented by a one kilometer grid overlay for all cells that had accounts in both 2012 and 2014. The consumption for all accounts within the cell are aggregated and then the change for that cell from 2012 to 2014 is computed. Finally, to bring the data into a normal-ish distribution, the resulting ratio has been log transformed. The result is we can quickly identify areas where the aggregated consumption has changed significantly relative to the rest of the study area since 2012. Diving deeper into the underlying data we can identify if the drives of this change are positives (e.g. implementation of energy efficiency or conservation) or negatives (like the closure of multiple accounts). We can use the same underlying datasets to understand if changes might also be due to new consumption at new accounts – like a new strip mall – in close proximity.[21] This can be particularly important for smaller accounts in close proximity to one another where, as we saw earlier in the paper, the individual accounts may be lower interest, but as a group they represent bigger opportunities for savings and engagement.

The geographic views and metrics presented here are not without their own caveats. Examples include systemically missing or miscoded accounts that create an inaccurate underlying analysis starting point or using inappropriate levels of data for the spatial analysis resolution. A textbook example of the latter of these is using geocoded accounts that are only at the zip code level to generate heat maps – the result can be heat maps showing a high concentration of consumption right at the middle of the zip code! Additionally, it's important to understand the assumptions behind the models used, and to communicate them – in this paper see footnote 19 on spatial interpolation as an example – so that data consumers know the inherent limits of what they are looking at.

Conclusion

As additional data are integrated into spatial databases, data mining and machine learning are likely to open opportunities to build more descriptive and accurate profiles of buildings, accounts, and the customers making the decisions about energy efficiency. This could include beginning to better understand how much energy different types of end uses contribute to a specific building type and vintage, and where, though time, buildings fitting this profile are showing either high or low engagement. This type of insight also facilitates the geographical aggregation of smaller and mid-sized customer accounts that may individually represent lower opportunity for savings compared to the larger accounts, but whose proximity means that as a population they represent stronger group potential. Leveraging this type of insight will become increasingly important as the low hanging efficiency opportunities with the largest customers is capitalized on and the more expensive to engage and harder to reach smaller and mid-size customers become an increasingly important share of savings.

Acknowledgements

DNV GL would like to thank the Massachusetts PA's for providing the raw data used throughout this paper, and the PAs and EEAC Consultants for their feedback throughout the analysis process. We would also like to thank the anonymous reviewer for their feedback in improving earlier drafts of the paper.

[21] At a more insightful level, this material can be further normalized to look at the EUIs, adjusted for degree days, or scaled to reflect scale of consumption rather than a straight ratio.

References

Baab, S. "Program vs Market : Which Approach is Better for Small Business Energy Efficiency?" *Energy in Transition* (blog), *DNV GL*, March 9, 2016, http://blogs.dnvgl.com/energy/program-vs-market-which-approach-is-better-for-small-business-energy-efficiency

Crowley, R., W. Brougher 2014. ""Watts" Where, and Why? Using GIS to Identify Energy Efficiency Opportunities" In *Proceedings of the ACEEE 2014 Summer Study on Energy Efficiency in Buildings*, 11:73–83 . Pacific Grove, CA: ACEEE

Davis, A., R. Crowley, and S. Gautam. 2016. "Beyond participation and savings – leveraging quantitative metrics for insight into C&I project and population trends" In *Proceedings of the ACEEE 2016 Summer Study on Energy Efficiency in Buildings* Pacific Grove, CA: ACEEE.

DNV GL. 2016. "*2014 C&I Customer Profile*" Prepared for: Massachusetts Program Administrators. http://ma-eeac.org/wordpress/wp-content/uploads/FINAL-2014-Customer-Profile-1.pdf

DNV GL. 2015. "*2013 C&I Customer Profile*" Prepared for: Massachusetts Program Administrators. http://ma-eeac.org/wordpress/wp-content/uploads/2013-CI-Customer-Profile-Report.pdf

DNV GL. 2014. "*2012 C&I Customer Profile*" Prepared for: Massachusetts Program Administrators. http://ma-eeac.org/wordpress/wp-content/uploads/CI-Year-2012-Customer-Profile-Final-Report.pdf

DNV GL. 2013. "*2011 C&I Customer Profile*" Prepared for: Massachusetts Program Administrators. http://ma-eeac.org/wordpress/wp-content/uploads/Customer-Profile-Project-Revised-Final-Report-6.20.13.pdf

DNV GL and NMR. 2015. "*Top-down Modeling Methods Study—Final Report*" Prepared for: Massachusetts Program Administrators. http://ma-eeac.org/wordpress/wp-content/uploads/Top-down-Modeling-Methods-Study-Final-Report.pdf

Freeman, L. G. Sadhasivan, R. Crowley, C. Driscoll. 2014. "Searching for Savings: Using Geospatial Analysis to Target Programs" In *Proceedings of the ACEEE 2014 Summer Study on Energy Efficiency in Buildings*, 2:122–133. Pacific Grove, CA: ACEEE

Mass Save Data. *Welcome to Mass Save Data*. Accessed March 2016. http://www.masssavedata.com/public/home

Novie, A., J. Teng. 2016. "Sociodemographics and Small Businesses:
Benchmarking Historical Participation Patterns Across SBDI Programs" In *Proceedings of the ACEEE 2016 Summer Study on Energy Efficiency in Buildings* Pacific Grove, CA: ACEEE.