Supporting Information

Timnit Gebru, Jonathan Krause, Yilun Wang, Duyun Chen, Jia Deng, Erez Lieberman Aiden, Li Fei-Fei

Table of Contents

1. Image Data

In this section, we provide additional detail on the methodology used to acquire annotated image data for our study. This data is required for two steps: to train computer vision models that detect and classify cars, and to apply these models on Street View images of cities of interest. This section proceeds by detailing how we obtained a comprehensive list of car categories, collected a large number of "product shot" images used to train our car classifier, gathered 50 million Street View images used in our analysis, and annotated a subset for training and verifying our model. We conclude with a complete description of the acquired metadata for each car category.

Car Categories. The first step in assembling a dataset of annotated car images is grouping cars into sets of visually indistinguishable classes. For example, while a 2003 Honda Accord coupe ex and a 2005 Honda Accord coupe ls special edition are manufactured in different years and have different trims (ex vs ls special edition), their exteriors look identical. Thus, these two cars should be grouped into the same class. Ideally, the set of classes would contain every type of car in common use. (1) presents a workflow to perform this grouping at minimal cost.

We first retrieved an initial list of 15,213 car types from the car website *Edmunds.com*, collected in August 2012. This forms a generally complete list of all cars commonly used in the United States that were produced from 1990 onward. Throughout this document we use the term "car" to refer to all types of automobiles with four wheels, including sedans, coupes, trucks, vans, SUVs, etc., but not including e.g. semi-trucks or buses.

As a first step toward grouping these categories into a smaller number of visually distinct classes, we used Amazon Mechanical Turk (AMT) to determine whether certain pairs of the 15k car types were distinguishable. The interface is shown in Fig. S10. Within each task we gave six pairs of categories and the user was prompted to determine 1) if the two classes had any visual differences, and 2) if they were different, on which parts they differed. Within each task we had two pairs for which we already knew the correct answer (as determined by hand), and we required that each user on AMT get the answer for those pairs correct in order to count their response. Photos for this task were acquired from the handful of example images that *Edmunds.com* provides. The authors cleaned up the data by hand, resulting in 3,141 categories of cars, with extremely subtle differences between these fine-grained categories. Fig. S11 shows two examples of classes with their constituent groups.

Product Shot Images. After assembling a list of categories consisting of visually indistinguishable sets of cars, we collected training images for each class. These are annotated images containing the car of interest. A commonly used method in the computer vision community is to perform web image searches for each category and cleanup the query images by hand to ensure that they contain the category of interest (2). However, the large number of classes in our dataset makes it infeasible to manually perform this task.

In order to collect training data in a scalable manner, we leveraged e-commerce websites. We crawled images from *cars.com* and *craigslist.org*, two sites where users are heavily incentivized to list the exact type of car they are selling. While these users are not necessarily car experts, they have detailed knowledge about their own car. In the case of *cars.com*, car categories are represented in a very structured format. Thus, after establishing a mapping between our categories and their format, we were able to simply scrape images for each category. For *craigslist.org*, we scraped posts from the "cars+trucks" listings of a variety of U.S. regions, and parsed the post titles to determine which of our categories the posts belonged to. Since these images are from websites with the purpose of selling cars, we call them "product shot" images.

Some product shot images show the car from an extremely close-up angle. Others only depict the interior of the car (Fig. S12). Since our purpose is to recognize cars in Google Street View images, our training set should have cars from view points that can appear in Street View. Thus, we filtered out images which do not contain one central automobile, with its exterior depicted in its entirety. Since this task is relatively simple, we crowdsourced it via AMT, using (3) for quality control. Our interface for this task is shown in Fig. S13.

In the final annotation step, we collected a bounding box (an axis-aligned rectangle tightly enclosing the object of interest) around the car in each image. This ensures that our car classifier is trained using visual information only from the car itself and not extraneous background. Bounding boxes were collected using the labeling methodology and UI of (4), but without the step for determining if there is more than one car in the image. That step is not necessary because the output of the previous AMT task ensures that each image contains exactly one prominent car.

Since some types of cars have many more images than others, we stopped annotating images for each category after collecting 200 labeled photos. Our goal is to build a model that can recognize as many types of cars as possible. Given our limited budget, it is more important to collect annotations for categories with few labeled images than for those with many annotated photos.

In the final step, we removed categories that do not have at least three disparate sources of data per class. We define one source of data as one post on any of the websites we used. This process resulted in our final dataset consisting of 2,657 car categories.

Street View Images. This section outlines our methodology for collecting approximately 50 million Google Street View images and annotating a subset of them to train our car detector and classifier. The process includes selecting GPS (latitude, longitude) points of interest, collecting images for each of these points, enclosing cars in a subset of these images with bounding boxes, and annotating the type of car contained in each box. The final step is performed by car experts.

Selecting GPS Points. Before gathering Google Street View images, we first have to determine which geographical (latitude, longitude) points we want to collect photos for. We call each latitude, longitude pair a GPS point. First, we select 200 cities for our analysis. These are the two largest cities in each state and the next 100 largest cities in the United States as determined by population (see Tab. S1 for a complete list). For each city, we sample potential points of interest within a square grid of length 20km, centered on one point known to lie within the city. There is a 25 meter spacing between points. We reverse geocode each of these points to determine whether they lie within the city of interest and how far away they are to the nearest road. We keep all points within 12.5 meters of the nearest road. This process did not provide full coverage for a handful of cities. Thus, we augmented these points with GPS samples from road data provided by the U.S. Census Bureau (5).

Sampling Images from Street View. For each GPS point, we attempt to sample 6 images from Google Street View, one for each of 6 different camera rotations. This was done via browser emulation and requires only the latitude and longitude of each point. However, we cannot immediately use photos retrieved with this process as they appear warped: an equirectangular projection is applied to images in a spherical panorama. We apply the reverse transformation before all subsequent tasks using the images (see Fig. S14 for an example).

Annotations on Amazon Mechanical Turk. While our product shot images can be used to train a car classifier, we cannot utilize them to train a car detector: a model that learns to localize all the cars in an image. This is because all of our product shot images include only one prominently featured car in each image.

Using the system of (4), we collected bounding box annotations in a subset of our Street View images. To increase the efficiency of this process, we first filtered out all images containing either zero or more than 10 cars via AMT, using the same interface and pipeline described in the section pertaining to product shot images. A randomly selected subset of 399,331 Street View images were annotated in this manner. We found that 26.6% of images were annotated as having no visible cars and 12.4% had more than 10 cars. The distribution of the number of cars in the remaining images is shown in Fig. S6A.

Fig. S6B plots bounding box size versus location. Cars located closer to the bottom of the image tend to occupy more space than those near the top. This agrees with the intuition that cars lower in the image are closer to the camera and therefore appear larger. Similarly, Fig. S6C shows a heatmap of bounding box location for cars in Street View. Most automobiles are located near the horizon line because that part of the image occupies more 3D space, i.e., more space in the real world. There is a sharp dropoff in the distribution of cars above the horizon line.

Expert Class Annotations. To learn to recognize automobiles in Street View images, a classifier needs to be trained with cars from these images. To this end, we labeled a subset of the bounding boxes from Street View images with the types of cars contained in them. This annotated data also enables us to quantitatively evaluate how well our classifier works. In contrast to product shot images, we do not know the types of cars contained in Street View photos. Therefore, we hired expert car annotators to label these images. Experts were primarily solicited via Craigslist ads. Those who were interested in performing our task were first asked to annotate cars in Street View images for one hour, and only those who could annotate at a speed of 1 car per minute and a precision of at least 80% were allowed to annotate further. 110 expert human annotators worked for a total of approximately two thousand hours to label our images.

Very small images typically do not contain enough visual information to discriminate fine levels of detail. Thus, annotators were only shown cars in bounding boxes whose height exceeded 50 pixels. 32.89% of bounding boxes in our dataset fulfill this criteria. The annotation task itself proceeded hierarchically: Fig. S15 shows the user interface for the task. Given a Street View bounding box, annotators were first asked to select the make of the car (Fig. S15(A)). They were then presented with a list of body types for the chosen make (Fig. S15(B)). After selecting the right body type, experts were shown a list of options for the car model, and finally, the trims and years associated with each model.

Since differences between categories can be extremely subtle at that final level, we also provided example images from each trim and year grouping for the annotator's benefit (Fig. S15(C)). At any point in the process, the annotator could declare that he or she did not have enough information to make a selection. Thus, each label at this finest level of detail represents a confident selection by a car expert. We collected a total of 69,562 car category annotations in this manner.

Car Metadata. In addition to the images, category labels, and bounding boxes, we also have metadata pertaining to each class, listed below.

- Make: The make of the car, of 58 possible makes. The makes we consider are: Acura, AM General, Aston Martin, Audi, Bentley, BMW, Buick, Cadillac, Chevrolet, Chrysler, Daewoo, Dodge, Eagle, Ferrari, Fiat, Fisker, Ford, Geo, GMC, Honda, Hummer, Hyundai, Infiniti, Isuzu, Jaguar, Jeep, Kia, Lamborghini, Land Rover, Lexus, Lincoln, Lotus, Maserati, Maybach, Mazda, McLaren, Mercedes-Benz, Mercury, Mini, Mitsubishi, Nissan, Oldsmobile, Panoz, Plymouth, Pontiac, Porsche, Ram, Rolls-Royce, Saab, Saturn, Scion, Smart, Subaru, Suzuki, Tesla, Toyota, Volkswagen, and Volvo.

- Model: The model of the car, of 777 possible models.

- Year: The manufacturing year of the automobile. Since cars might not change appearance over a small number of years, this is typically listed as a range of years. The minimum year in our dataset is 1990, and the maximum year is 2014.

- Body Type: The body type of the car. The 11 possible values are: convertible, coupe, hatchback, minivan, sedan, SUV, truck (regular-sized cab), truck (extended cab), truck (crew cab), wagon, and van.

- Country: The manufacturing country of the automobile. The 7 possible countries are: England, Germany, Italy, Japan, South Korea, Sweden, and USA.

- Highway MPG: The typical miles per gallon of the car when driven on highways. If a class contains cars with multiple years, it is annotated with the highway MPG of the oldest car in the group.

- City MPG: The typical miles per gallon of the car when driven on non-highway streets.

- Price: the price of the car in 2012.

This metadata was acquired via *Edmunds.com* in August 2012, with some missing data (a handful of car prices) filled in by car experts afterward. In cases where a class consists of multiple visually indistinguishable types of cars, it is annotated with the metadata of the oldest car in the set.

Dataset Summary. Tab. S2 provides a summary of the annotations collected for both product shot and Street View images, which we split into training (50%), validation (10%), and test (40%) sets for use in training our car detector and classifier.

2. Demographic Data

Income. Data for median household income was obtained from the American Community Survey (ACS) (6), and was collected between 2008-2012. We used census variable `B19013_001E`, "Median household income in the past 12 months (in 2013 inflation-adjusted dollars)".

Education. Education data was also obtained from the ACS (6). Education levels are split into the following mutually exclusive categories (census codes in parentheses):

- Less than high school graduate (`B06009_002E`)
- High school graduate (includes equivalency) (`B06009_003E`)
- Some college or associate's degree (`B06009_004E`)
- Bachelor's degree (`B06009_005E`)
- Graduate or professional degree (`B06009_006E`)

Race. Racial demographic data was also obtained from the ACS (6), and corresponds to census codes `B02001_002E` ("White alone"), `B02001_003E` ("Black or African American alone"), and `B02001_005E` ("Asian alone").

Voting. Data for the 2008 U.S. presidential election was provided to us by the authors of (7) and consists of precinct-level vote counts for Barack Obama and John McCain. For all of our analyses, we ignore votes cast for any other person, i.e. the count of total votes is determined solely by votes for Obama and McCain. We visualize this raw data in Fig. S16.

Obama received greater than 50% of the votes in most of the precincts in our dataset. This can partially be attributed to the fact that he won the popular vote in the 2008 election. Precincts in our dataset are also located in major cities which favor candidates from the Democratic party. Interestingly, Obama received an extremely high percentage ($\geq$ 95%) of the votes in many precincts in our dataset. A large portion of these precincts have high concentrations of African Americans, who overwhelmingly voted for him during the 2008 election.

3. Additional Details for Car Detection and Classification

Isotonic Regression. Our car detection model outputs bounding boxes and scores associated with each box. We use isotonic regression to convert these scores to probabilities depicting the likelihood of containing a car. Isotonic regression learns a probability for each detection score subject to a monotonicity constraint. Concretely, after sorting n validation detection scores $s_1, \ldots, s_n$ such that $s_i \leq s_{i+1}$, and with y_i a binary variable denoting whether detection i is correct (has Jaccard similarity of at least 0.5 with a ground truth car bounding box), isotonic regression solves the following optimization problem:

$$\begin{array}{ll} \underset{p_1,\ldots,p_n}{\text{minimize}} & \sum_{i=1}^{n} \|y_i - p_i\|_2^2 \\ \text{subject to} & p_i \leq p_{i+1}, \qquad 1 \leq i \leq n-1 \end{array} \qquad [1]$$

Given a new detection score, a probability is estimated by linear interpolation of the p_i. We plot the learned mapping from detection scores to probabilities in Fig. S7A.

Additional Design Considerations.

Car Detection. We made a number of additional design choices while training and running this car detector in practice. First, we only detected cars that are 50 pixels or greater in width and height. The output of our detector is fed into the input of our car classifier. Thus, detected cars need to have sufficient resolution and detail to enable the classifier to differentiate between 2,657 categories of automobiles. Similarly, we trained our detector using cars with greater than 50 pixels width and height. Our DPM is trained on a subset of 13,105 bounding boxes, reducing training time from a week (projected) to 15 hours. Using this subset instead of all ground truth bounding boxes results in negligible changes in accuracy.

Car Classification. One further challenge while classifying Street View images is that our input consists of noisy detection bounding boxes. This stands in contrast to what would otherwise be the default for training a classifier – ground truth bounding boxes that are tight around each car. To tackle this challenge, we first measured the distribution of the intersection over union (IOU) overlap between bounding boxes produced by our car detector and ground truth boxes in the validation data. Then, we randomly sampled the Street View image region input into the CNN according to this IOU distribution. This simulates detections as inputs to the CNN and ensures that the classifier is trained with similar images to those we encounter during testing.

4. Raw Correlations Between Car Attributes and Demographics

Fig. S2 shows the magnitude of weights learned by our model for inferring various demographic attributes. That is, we sort the coefficients of the regression model in descending order. Each coefficient is uniquely associated with one of the 88 car features used in our model. We then plot the top 5 and bottom 5 values. However, looking at the model weights may not always be informative since some car features are highly correlated (e.g. Lamborghini and car price). A linear model distributes the magnitude of its coefficients among highly correlated features. Thus, the weights of highly predictive car features might still be small. Thus, to investigate the relationship between various demographic variables and cars, we list the raw correlations between all of our car features and ground truth demographic data.

Income. Correlations between median household income and each of our car attributes are given in Tab. S4. The five car attributes that correlate most positively with median household income are %Foreign (r=0.47), %Country: Japan (r=0.45), Price (r=0.44), %Make: Lexus (0.44), and %Country:

Germany (r=0.43). The five car attributes that correlate most negatively with median household income are %Country: USA (r=-0.47), %Year: 1995-1999 (r=-0.42), %Make: Buick (r=-0.40), %Make: Oldsmobile (r=-0.40), and %Make: Dodge (r=-0.38).

Education. We show correlations between each of our car attributes and education levels in Tab. S5, Tab. S6, Tab. S7, Tab. S8, and Tab. S9, and the five car attributes that correlate most positively and most negatively with each race are given in Tab. S10.

Race. Correlations between our car attributes and the percentage of each race considered (White, Black, and Asian) are given in Tab. S11, Tab. S12, and Tab. S13, respectively. The five car attributes that correlate most positively and most negatively with each race are given in Tab. S14.

Voting. We show correlations between %Obama and all of our car-centric variables in Tab. S15, and plot our predictions versus actual voting percentages in Fig. S18. The five car attributes that correlate most positively with Obama's percent of votes are Body Type: Sedan (r=0.48), #Cars/Image (r=0.37), MPG Highway (r=0.33), and MPG City (r=0.26). The five car attributes that correlate most negatively are Body Type: Crew Cab (r=-0.48), Body Type: Extended Cab (r=-0.43), Body Type: Regular Cab (r=-0.30), Price (r=-0.28), and Body Type: SUV (r=-0.22).

5. Cross Validated Performance with Randomly Split Training Data

In the main text, we chose zip codes and precincts in counties starting with "A", "B" or "C" to train our model and evaluated the model on the rest of our data. This was done to show that we could train a model that can infer demographics with reasonable accuracy, using very little data (approximately 10% of our data). We also wanted to ensure that zip codes in the same city were not used in training and testing.

Below, we present demographic inference results using a different training methodology. We randomly partitioned our zip codes and precincts into five sets, iteratively training a model on four of the parts and predicting on the held out set. As before, we normalize the car features to have zero mean and unit standard deviation (parameters determined on the training set of four parts). We furthermore clip predictions to be within the range of the current training data, preventing predictions from becoming too extreme. In all experiments at the zip code level we restricted the zip codes used to be ones with a population of at least 5,000 and at least 500 detected cars, which reduces the number of zip codes under consideration from 3,068 to 2,430, mostly as a result of the restriction on the number of detected cars.

Fig. S17 and Fig. S18 show scatter plots and Pearson correlation coefficients for actual vs. predicted values of income, education, race and voting patterns respectively. The scatter plots show results at the highest level of spatial granularity our analysis is performed in (precinct level for voting patterns and zip code level for everything else). (The r-values for the correlations were: median household income, $r = 0.79$; percentage of Asians, $r = 0.79$; percentage of Blacks, $r = 0.81$; percentage of Whites, $r = 0.76$; percentage of people with a graduate degree, $r = 0.78$; percentage of people with a bachelor's degree, $r = 0.76$, percentage of people with some college degree, $r = 0.67$, percentage of people with a high school degree, $r = 0.76$; percentage of people with less than a high school degree, $r = 0.73$; percentage of people who voted for Barack Obama during the 2008 presidential election, $r = 0.67$.

6. City Car Attributes

Using our car detections, we can answer specific questions about cars and cities. For example, we can ask what the average age of a car on the road is, what the average car price is (in the US as a whole and in each city), which city has the most expensive cars on average (New York, NY), or the highest percentage of foreign cars (San Francisco, CA - 60.02%), etc...We show maps comparing a subset of these attributes across our 200 cities (average car price, the percentage of foreign cars, BMWs, Chevrolets, Toyota Prius, and Ford F-150s) in Fig. S19.

7. Additional ACS Variables

In this section, we report results for an additional 28 ACS variables that were inferred using our Google Street View based methodology. While the ACS has many variables, we obtained a subset of 28 attributes that are indicators of income, race, education and occupation levels as well as other characteristics of neighborhoods such as owner occupancy of housing units. Fig. S20 shows scatter plots of actual vs. predicted values (cross validated performance with randomly split training data).

Some variables can be inferred with high accuracy using our methodology (e.g the Pearson correlation coefficient between actual vs. predicted values for median household income for units with a mortgage is $r = 0.80$). Variables such as the age of one's children can be inferred with moderate accuracy ($r = 0.54$). On the other hand the *Farming* variable cannot be inferred from cars at all $r = 0.0$ ($p \ll 1e-7$ for all variables). This is in part due to the fact that our car features reflect percentages and therefore are most suited to infer the percentage of something as opposed to the actual value. E.g., we have setup our methodology to infer the percentage of inhabitants in a neighborhood with a bachelor's degree as opposed to the total number of citizens who have obtained a bachelor's degree. Note that we also do not include all vehicles in our dataset (e.g. we omit tracktors and large trucks). Including these might improve our accuracy in estimating farming related ACS variables. It is also important to note that we have not refined our methodology to be able to infer these additional variables and have simply applied our current methodology to them and reported our results. However, some of these results indicate that not all ACS variables can be inferred from analyzing cars in Google Street View images. Thus, our methodology is only applicable to those variables that are most strongly correlated with car preferences.

This additional data was also obtained using the ACS API (6). We list the variables of interest below (census codes in parentheses):

- Median Age by Sex-Total (`B01002_001E`)
- Median Age by Sex-Male (`B01002_002E`)

- Median Age by Sex-Female (B01002_003E)
- Median Household Income by Age of Householder-Total (B19049_001E)
- Housing Units-Total (B25001_001E)
- Occupancy Status-Total (B25002_001E)
- Occupancy Status-Occupied (B25002_002E)
- Occupancy Status-Vacant (B25002_003E)
- Median Number of Rooms-Median number of rooms (B25018_001E)
- TENURE BY UNITS IN STRUCTURE-Owner-occupied housing units (B25032_002E)
- TENURE BY UNITS IN STRUCTURE-Renter-occupied housing units (B25032_013E)
- Median household income for units with a mortgage (B25099_002E)
- Median household income for units without a mortgage (B25099_003E)
- Bedrooms-Total (B25041_001E)
- Total Population (B01003_001E)
- Total Race (B02001_001E)
- Total Education (B06009_001E)
- American Indian and Alaska Native alone (B02001_004E)
- aggregate number of vehicles for travel (B08015_001E)
- age of own children (B05009_001E)
- own children under 6 years (B05009_002E)
- 6-17 years (B05009_020E)
- management (B24021_002E)
- service (B24021_018E)
- farming (B24021_030E)

8. Alternate Sources of Data

Department of Motor Vehicles Registration Data. Cars in Google Street View images capture the types of automobiles that are parked, or pass through a neighborhood in a given snapshot of time. If it is near a freeway or a parking lot, a high density of cars will be detected. In our work, we use this information to infer the demographic characteristics of neighborhoods. How do our detected automobiles compare to those from DMV data? We do not expect them to be exactly the same because cars in Google Street View do not always belong to inhabitants of the neighborhood they are captured in.

We compared the distribution of cars we detect in Street View images with the distribution in Boston, Worcester and Springfield, MA (the three Massachusetts cities in our dataset), released by the Massachusetts DMV, the only state to release extensive vehicle registration data (8). We measured the Pearson correlation coefficient between each detected and registered make's distribution across zip codes (Fig. S21 (A)). Twenty five of the top thirty makes have a Pearson's r correlation of r>0.5. Conversely, classifying according to the 2011 national auto sales distribution (9) results in correlation r=0 with DMV data. Beyond Massachusetts, we measure the correlation between our detected car make distribution and the 2011 national distribution of car makes as r=0.97. Fig. S21 (B) plots the DMV values vs. our predicted percentages for the distribution of Hondas in each zip code. Fig. S22 —Fig. S23 show the latter plot for all makes instead of only Hondas.

Vehicle data for Massachusetts is available from the Massachusetts Vehicle Census (8) and contains anonymized zip code and model information for all vehicles registered in Massachusetts between 2008 and 2011. Since our comparison with this data was done at the make level, aligning their list of car classes and ours entailed only aligning the list of makes, which was done by hand. We performed our experiments on the intersection of detected and registered makes resulting in a total of 45 makes.

To calculate the distribution of car makes in each zip code, we compute the expected number of each of the 2,657 car classes across the zip code, then simply use the make metadata associated with each car class to calculate the expected number of cars for each make within the zip code. Since the expected number of instances of a particular car across a zip code is the sum of the expected number of instances of the car across all images within that zip code, the problem reduces to calculating this expectation for a single image.

With I an image and c one of the 2,657 classes, we decompose the expectation for a single image as

$$\mathbb{E}[\#\text{class}\, c|I] = \sum_{\text{bbox}\, b} P(\text{car}|b, I)P(\text{class}\, c|\text{car}, b, I) \qquad [2]$$

where we are summing over all bounding boxes b for generic cars detected by our model. We model $P(\text{car}|b, I)$ using isotonic regression (described in Methods), and $P(\text{class}\, c|\text{car}, b, I)$ corresponds to the conditional probabilities output by the softmax layer of our CNN classifier.

To obtain the percentage of each make, we aggregate these category-level expectations by car make and compute percentages using the make-level expectations.

The Pearson coefficient for each make is calculated by taking the percentage of that make in one zip code as a single data point. We chose zip codes with greater than 5,000 inhabitants and 500 detected cars in the three Massachusetts cities in our dataset (Boston, Springfield, Worcester), resulting in a total of 37 zip codes. For all experiments, we used registration records that were valid during the second quarter of 2010. As outlined by (8), registration datasets are most complete at that point.

Fig. S24A shows the Pearson correlation coefficient between the distribution of registered and detected cars across zip codes for all makes that are in the intersection of our dataset and Massachusetts registration data: in contrast to Fig. S21 (A), this shows all makes instead of the top 30. In addition to comparing the distribution of each registered and detected make across zip codes, we performed two additional experiments. First, we compared the distribution of all detected and registered makes per zip code, computing the Pearson correlation coefficient for each zip code (Fig. S24B). All zip codes have correlation greater than 0.8. In contrast, classifying according to the national distribution only results in correlations greater than 0.45 for all zip codes. Next, we compared the total distribution of registered and detected cars in all 37 zip codes (Fig. S25) and measured a correlation coefficient of 0.94. Prediction using the national sales distribution instead of our approach only has correlation 0.82.

Since we have very few zip codes with DMV data, we did not train a model using DMV data to infer demographics. However, our experiments show that there is significant overlap between the information we collect and DMV data. But we do not capture exactly the same information. The question we ask is how the look of a neighborhood, as captured by the cars that one sees in it, is related to demographics. If the neighborhood is near a freeway, we would detect a lot of cars. If the neighborhood is one with many cars parked on the street, our method would take this into account. On the other hand, performing this analysis with DMV data associates the cars owned by residents of a particular neighborhood with the demographic makeup of that neighborhood. Our work also presents a pipeline to measure various attributes using publicly available visual data. If one were to, for example, study the relationship between tree species in a neighborhood and the health of its inhabitants, they can use our methodology (data collection, detection, classification etc...) to perform their study.

Satellite Night Lights. Many works, e.g. (10–13), have studied the relationship between nighttime lights observed through satellite imagery and total population, population density, GDP and a few other variables such as the number and density of establishments. While most of these works have focused on using night lights to predict population density and income levels in developing countries with very little census data, (11) investigates the use of this technique to estimate income in nations such as Sweden with near uniform distribution of electricity among its population.

What (10) found was that in countries such as Sweden where living standards are much more uniform than the developing world, the correlation between night time luminosity and total wage values was not as high. Using a sophisticated statistical model (rather than simple linear regression used in our paper), the correlation reported by (10) between actual and predicted total wages was 0.52.

While, in the developing world, the presence/absence of lights in one's household is a strong indicator of income levels, this is not necessarily the case in countries like the United States where almost all citizens have access to electricity. As (13) notes, in the United States where living standards are much more uniform than the developing world, the higher concentration of lights in coastal areas near the oceans and the Great Lakes reflects the higher population densities there as opposed to higher income.

To our knowledge, there is no prior work attempting to infer race, voting affiliations of education levels purely using night time data. This is because even the correlation between economic output and night time data has been found to be relatively weak in the developed world (10, 11). While certain races and political affiliations have shown to have historical preferences for different car brands (14–17), the relationship between night lights and these variables in the United States is unclear. However, as discussed in the main text, our paper presents a proof of concept that can be expanded upon. Our future work plans to incorporate all other publicly available data (including night time lights, CNN features, other detected objects such as trees and pedestrians) into our model to improve its accuracy.

9. Related Work Using Google Street View

A number of prior works have asked similar questions to our work about cities. Salesses et al. (18) collected a dataset of approximately 1M Google Street View images labeled with annotators' perception of the safety, uniqueness and wealth of the locations portrayed by each image. Subsequent works (19–21) infer these labels using global image features (20) and CNN features (19, 21). Dubey et al. (21) performs this analysis at a large scale. In these works, a location's perceived wealth/safety/location is given a number between 0 and 10 where 10 corresponds to the wealthiest/safest/most unique location.

While the goal of all of these works is to predict people's perception of a location's wealth, uniqueness and safety using images, our goal is different. Our goal is to predict the actual median household income, racial makeup, education level, voting patterns of a certain location given its street view images. Thus, instead of inferring the perception of a particular neighborhood's safety, uniqueness and wealth given its photos, we are interested in predicting its true characteristics as recorded by government agencies such as the ACS and voter polling stations.

Arietta et al (22) uses features from convolutional neural networks and other global image features (HOG and GIST) to predict housing prices and violent crime rates in San Francisco, Chicago, Boston, Oakland, Seattle and Philadelphia. While they obtain high correlations between predicted and ground truth values when locations in the same city are used for training and testing purposes, the accuracy is low while using different cities for training and testing. For example, while high Pearson r values are achieved between actual and predicted housing prices when the same city is used for training and testing (e.g. 0.815 for Boston), the Pearson r for training and testing across different cities is much lower (e.g. 0.444 while training on Boston and testing on Seattle). Thus, these visual attributes do not necessarily generalize across cities.

10. Baselines

In this section, we compare our approach to a number of baselines that infer demographics from various global image features and course grained census data.

Projecting Course Census Data to Fine Geographic Locations. Here, we investigate the predictive power of features derived from ACS data at course spatial granularity, to infer demographics at finer spatial granularity. Specifically, we train a regression model to infer demographics at the zip code level, using ACS data at the city level. In order to have a baseline that performs better than simply assigning all zip codes in a city to the demographic data of the city, one must assume that some demographic data is available at the zip code level. In our case, we assume access to the total population of each zip code. Thus, our model is trained using 9 city level and 1 zip code level ACS data. That is, each zip code is now represented by a 10 dimensional vector consisting of 9 demographic variables at the city level and the number of inhabitants in the zip code. We then use the exact same procedure as our methodology described in the main text to train a model inferring zip code level demographic variables. The only difference is that we use these 10 ACS variables as our features instead of an 88 dimensional vector representing each zip code's car related attributes. This baseline was performed using all of the zip codes in our data.

Figs. S28— S29 plots our results (actual vs inferred values). The Pearson r values and p values for all variables are listed on the plots. They are $r = 0.05, p = 0.007$ for median household income; $r = 0.174, p \ll 1e-7$ for the percentage of people with less than a high school education; $r = 0.1p \ll 1e-7$ for the percentage of people with a highschool education; $r = 0.08, p = 7e-5$ for the percentage of people with some college education; $r = 0.12, p \ll 1e-7$ for the percentage of people with a bachelor's degree; $r = 0.19, p \ll 1e-7$ for the percentage of people with a graduate degree; $r = 0.09, p = 0.1$ for the percentage of Whites; $r = 0.04, p = 0.04$ for the percentage of Blacks; $r = 0.08, p = 9.8e-5$ for the percentage of Asians. Given these results, coarse level census values seem to have very little predictive power for inferring education levels, and no ability to infer income or race at a more granular level.

Pretrained Convolutional Neural Network Features. Here, we infer demographics with the same methodology as before but construct our model using pretrained CNN features instead of car attributes detected in Google Street View images. That is, our zip code level features consist of $fc6$ activations from an AlexNet (23) CNN instead of the 88 car attributes we used before. As first shown by (24) in 2014, features from a CNN pretrained on ImageNet have been proven to be much more discriminative than handcrafted ones such as SIFT, GIST or HOG (24).

We represent each zip code as a 4096 dimensional vector which is the average of $fc6$ activations obtained from all images in a particular zip code. That is, we input each Google Street View image in a particular zip code to a CNN pretrained on ImageNet. We then take the 4096 $fc6$ activations from each image and average those in the same zip code to obtain a single feature representation for each geographic location of interest. We subsequently use the same methodology (ridge regression) to train a model inferring race, education and income levels obtained from the ACS. We do this analysis on all images from states that start with "A" in our data (i.e., cities in Alabama, Alaska, Arizona, Arkansas). This consists of the cities: Birmingham (Alabama), Montgomery (Alabama), Anchorage (Alaska), Fairbankds (Alaska), Phoenix (Arizona), Tucson (Arizona), Little Rock (Arkansas) and Fort Smith (Arkansas). This comprises of 5,144,334 images which is slightly higher than 10% of the 50 million images in our data. We performed this analysis on a subset of our dataset due to the large amount of time it takes to obtain CNN features using our GPU cluster.

We compare our methodology to this baseline while using a randomly selected subset of 10, 20, 30, 50 and 80 percent of training data, and testing on the rest of the data. This allows us to investigate how the two approaches compare as we have access to little or plentiful training data. Figs. S26— S27 show how the two methods compare with increasing amounts of training data. We plot the Pearson r between actual and inferred values using car based and pretrained CNN based features. The CNN based features have no predictive power with little training data while the car based features have some (albeit little) predictive power even while using 10% of the data for training amounting to just 12 zip codes. However, CNN based features approach the performance of our car based method when large amounts of training data (i.e. over 50%) are available.

We hypothesize that this is due to the low generalizability of Google Street View CNN features to new locations (as mentioned in the related works section above). While using small amounts of training data, no CNN features from locations in visually similar zip codes or cities are available. However, by the time we use more than 50% of our data for training, there are images from nearby zip codes in the training set, which ensures high visual similarity between images in the train and test set. These experiments show that car based features are more predictive and generalizable when little training data is available (such as in our case when we would like to minimize the amount of ACS ground truth data required to train a good model). However, it is clear that CNN features have predictive power and would probably enhance our model's performance if our method is augmented to include them. Nevertheless, our simple model trained using car features is much more interpretable than that trained using 4096 dimensional $fc6$ activations. One can simply look at car attributes where our model places high weights, learn the relationship between various demographic variables and car attributes and gain further insight into American culture.

11. Timelapse Experiments–Inferring Demographics Across Time

Our work so far has used ground truth data from one geographic location to infer demographics in another location at the same time point. In this section, we perform preliminary experiments to test the feasibility of inferring future demographic trends in a particular neighborhood given its current and past ACS data. To perform our experiments, we use the recently introduced Google Street View timelapse tool which shows time-lapse images of the same location over time. Fig. S30 shows the dramatic economic development of a particular address in Brooklyn, New York over time.

We retrieved Google Street View timelapse images for New York city using the same methodology to gather Street View images discussed before. Figs. S31— S32 show maps of the coverage at the district level for each year. Fig. S33 plots the number of images per year for each district. As these maps and plots show, the number of images is nonuniform across years and districts. For example, there is no Statin Island coverage for 2014. This nonuniformity in sampling can cause errors in our estimations. Nevertheless, we conducted preliminary experiments to asses the feasibility of performing future research in this area.

Once we gathered the images across time, we detected and classified all the cars in these images following the same procedure as our retrieved Street View images. For each year, we represented each district as a collection of the car features in that district using the methodology described in the main text.

We retrieved 2011, 2012, 2013 and 2014 yearly ACS data for each of the 13 congressional districts in New York city. Yearly ACS data is only available at the district and county levels and for years after 2011. Figs. S34— S35 plots the data of interest (race, education, income) across time for each district. There is very little change from year to year.

We used past ACS data and images to train a model predicting future ACS data. Specifically, we trained a ridge regression model using 2013 ACS data as ground truth and tested the model on timelapse images from 2014. Figs. S36— S38 plots our results. We achieve a very high correlation between ground truth results and our predictions (e.g. Pearson r =0.93 for the percentage of Asians. All the Pearson coefficients and p values are listed on the plots). However, since we only have data at the district level and the change from year to year is very small, a baseline assuming constant ACS data (i.e. no change from the prior year) achieves even higher Pearson r (r=0.99 for the percentage of Asians). We could not perform this experiment at the zip code or precinct level for lack of data. The ACS does not have yearly data at the zip code level. We hypothesize that more demographic changes occur at the zip code level from year to year. Then, district level ACS data could be used to calibrate car preferences over time and infer ACS data at the zip code level.

In our second experiment, we wanted to specifically predict the change in ACS data. To do this, we trained a ridge regression model using the difference in ACS data between consecutive years as ground truth. To encode the change in detected cars, we subtracted the car features for consecutive years in each district. These features were used to train the regression model. We show some preliminary results in Fig. S39 showing the change in the number of high school educated people in NYC (one of the variables that showed some change over the years). While the potential to detect demographic trends is present, our ability to detect small changes at the district level is not currently strong. And more research and experimentation is necessary to have conclusive results.

To examine the stability of correlations between car types and neighborhood inhabitants, we examine the correlation between various car attributes and ACS data over time. In Fig. S40, we plot the correlation between median household income and various car attributes over time, as well as the associated p values. We see that higher level attributes (such as adjusted car price and age) consistently have a statistically significant positive/negative correlations across time. Some makes (e.g. Lamborghinis) also have a consistently statistically significant high correlation with income. E.g. $r > 0.65$ for 2012, 2013, 2014). On the other hand, Porsches have a statistically significantly high positive correlation with income in 2013 and 2014 ($r > 0.65$) but that correlation is 0.2 for 2012 with a p value of 0.5. The latter is not a statistically significant correlation. This discrepancy is most probably due to errors introduced by the non uniform sampling of images across years and neighborhoods. However, it is still important to note that car preferences in a particular location could change over time. To study this effect, these analyses should be repeated using images from GPS points that are evenly sampled across New York to minimize errors. Nevertheless, these preliminary results show that socioeconomic trends could potentially be inferred using images to capture the change in car types across neighborhoods.

12. Sources of Error

One source of error lies in the quality of the Street View images. Images from Street View may have image stitching artifacts, have street names in the images, and cars in Street View images might not be entirely visible, either due to occlusions with other objects or simply being cut off by the edge of the image. These factors make car detection and classification in Street View images more difficult. However, images used in our performance analysis are also subject to these issues, and thus our demonstrated performance holds despite these challenges.

Another source of error consists of biases in sampling, consisting of either sampling at certain roads or regions preferentially (due to having incomplete or out of date information about roads in cities) or sampling images taken at different dates. Although being able to properly account for these types of errors would undoubtedly improve and strengthen the analysis, these factors do not diminish or weaken the results we already have, and can be considered a source of noise in the data. For example, addressing these limitations might result in stronger correlations and lower p-values. But unless these are systematic errors across the entire United States, they do not affect the validity of the results presented in this work.

Errors in detection and classification of cars can also contribute to inacurraccies in the final results. Our car detection system does not have perfect precision and recall. Similarly, the accuracy with which our car classifies all 2,657 classes is not the same. However, we aggregate these classes by make, body type and other metadata before performing demographic inference. As shown in Fig. S9, the errors we make at levels such as the price and body type are reasonable. I.e., we rarely mistake a very expensive car for a very cheap one, or vice versa. This alleviates the level of systematic error in our demographic predictions.

The data used for demographic, crime, and voting analysis were not collected at the same point in time as images taken in Street View, and thus any drift in those sources over time is also a potential source of error. For example, 2008 presidential election data was used in our analyses, but the majority of Street View images were taken after 2008. This is due to the fact that precinct level election data for 2012 was not available for all of our 200 cities. While this is an unrecoverable source of error, it is primarily a problem when such statistics change rapidly over time, and when demographic projections are performed across time. I.e., when inferences for future years are performed based on ground truth data for current and prior years. Thus, our preliminary timelapse experiments do not use voting data and make inferences based on images and ACS data from the appropriate years.

1. Gebru T, Krause J, Deng J, Fei-Fei L (in press 2017) Scalable annotation of fine-grained categories without experts in *CHI*.
2. Deng J, et al. (2009) Imagenet: A large-scale hierarchical image database in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on.* (IEEE), pp. 248–255.
3. Sheng VS, Provost F, Ipeirotis PG (2008) Get another label? improving data quality and data mining using multiple, noisy labelers in *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining.* (ACM), pp. 614–622.
4. Su H, Deng J, Fei-Fei L (2012) Crowdsourcing annotations for visual object detection in *Workshops at the Twenty-Sixth AAAI Conference on Artificial Intelligence.*
5. (2012) TIGER/Line - Geography - U.S. Census Bureau (https://www.census.gov/geo/maps-data/data/tiger-line.html). Accessed: 2014-11.
6. (2012) American Community Survey 5 Year Data (2008-2012) (http://www.census.gov/data/developers/data-sets/acs-survey-5-year-data.html). Accessed: 2014-9.
7. Ansolabehere S, Palmer M, Lee A (2014) Precinct-Level Election Data (http://hdl.handle.net/1903.1/21919). Accessed: 2015-1.
8. (year?) (http://www.37billionmilechallenge.org/). Accessed: 2014-10.

9. Cain T (year?) U.S. Auto Sales By Brand - 2011 Year End (http://www.goodcarbadcar.net/2012/01/us-auto-sales-by-brand-2011-year-end.html/).
10. Mellander C, Lobo J, Stolarick K, Matheson Z (2015) Night-time light data: A good proxy measure for economic activity? *PloS one* 10(10):e0139779.
11. Chen X, Nordhaus WD (2011) Using luminosity data as a proxy for economic statistics. *Proceedings of the National Academy of Sciences* 108(21):8589–8594.
12. Elvidge CD, et al. (1997) Relation between satellite observed visible-near infrared emissions, population, economic activity and electric power consumption. *International Journal of Remote Sensing* 18(6):1373–1379.
13. Henderson JV, Storeygard A, Weil DN (2012) Measuring economic growth from outer space. *The American Economic Review* 102(2):994–1028.
14. (year?) Differences in new vehicle buyers' ethnicities predict where future sales will grow (http://www.strategicvision.com/auto_2014_ethnic_release.php).
15. Myers M, Dean SG (2007) Cadillac flambé": Race and brand identity in *Proceedings of the 23rd conference on historical analysis and research in marketing.* Vol. 13, pp. 158–161.
16. Tierney J (2005) Your car: politics on wheels. *New York Times* 1.
17. Florida R, Mellander C (2012) Democrat vs. Republican: Who's Buying What Car?, (Strategic Vision), Technical report.
18. Salesses P, Schechtner K, Hidalgo CA (2013) The collaborative image of the city: mapping the inequality of urban perception. *PloS one* 8(7):e68400.
19. Ordonez V, Berg TL (2014) Learning high-level judgments of urban perception in *European Conference on Computer Vision.* (Springer), pp. 494–510.
20. Naik N, Philipoom J, Raskar R, Hidalgo C (2014) Streetscore–predicting the perceived safety of one million streetscapes in *2014 IEEE Conference on Computer Vision and Pattern Recognition Workshops.* (IEEE), pp. 793–799.
21. Dubey A, Naik N, Parikh D, Raskar R, Hidalgo CA (2016) Deep learning the city: Quantifying urban perception at a global scale in *European Conference on Computer Vision.* (Springer), pp. 196–212.
22. Arietta SM, Efros AA, Ramamoorthi R, Agrawala M (2014) City forensics: Using visual elements to predict non-visual city attributes. *IEEE transactions on visualization and computer graphics* 20(12):2624–2633.
23. Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks in *Advances in neural information processing systems.* pp. 1097–1105.
24. Donahue J, et al. (2014) Decaf: A deep convolutional activation feature for generic visual recognition. in *Icml.* Vol. 32, pp. 647–655.

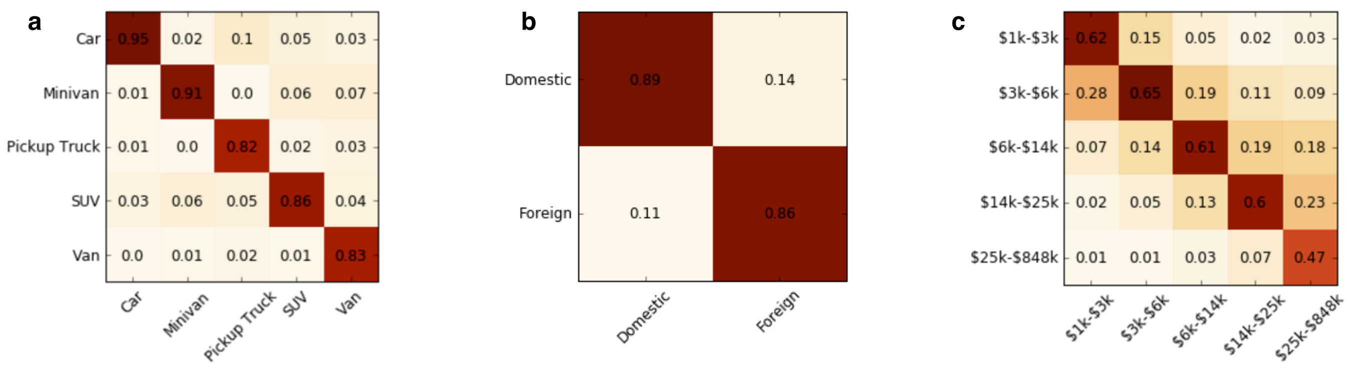

Fig. S1. Confusion matrices show the accuracy with which we classify various car attributes such as type of vehicle in a, whether or not it is domestic in b, and its price in c. The entry in row i and column j indicates the percentage of times ground truth attribute j was classified as attribute i. Thus, the values for all rows in a single column should add up to 1.

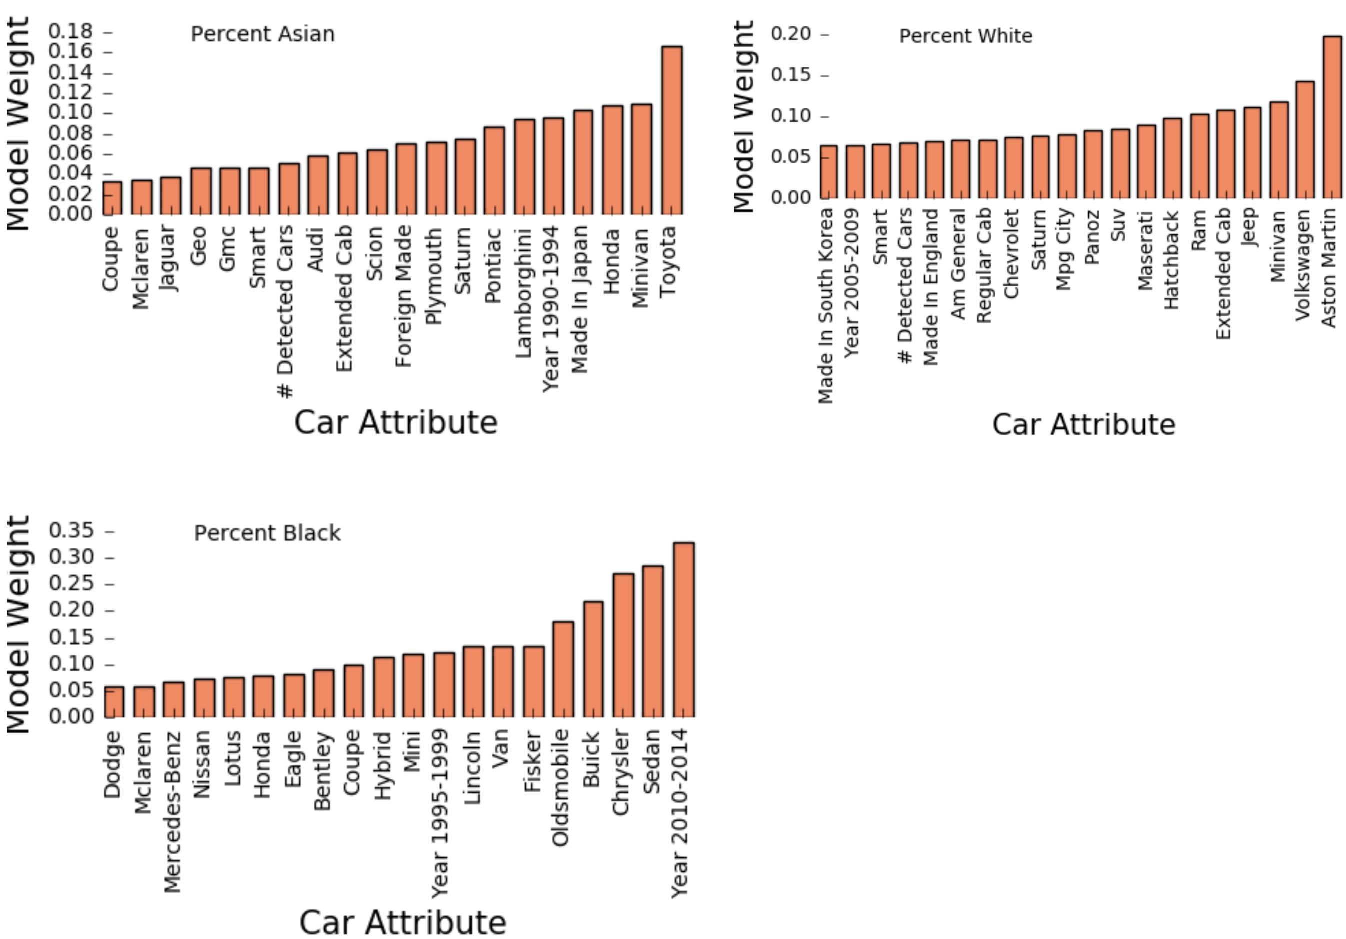

Fig. S2. Bar plots showing the top 10 car features with high positive weight in our race estimation model.

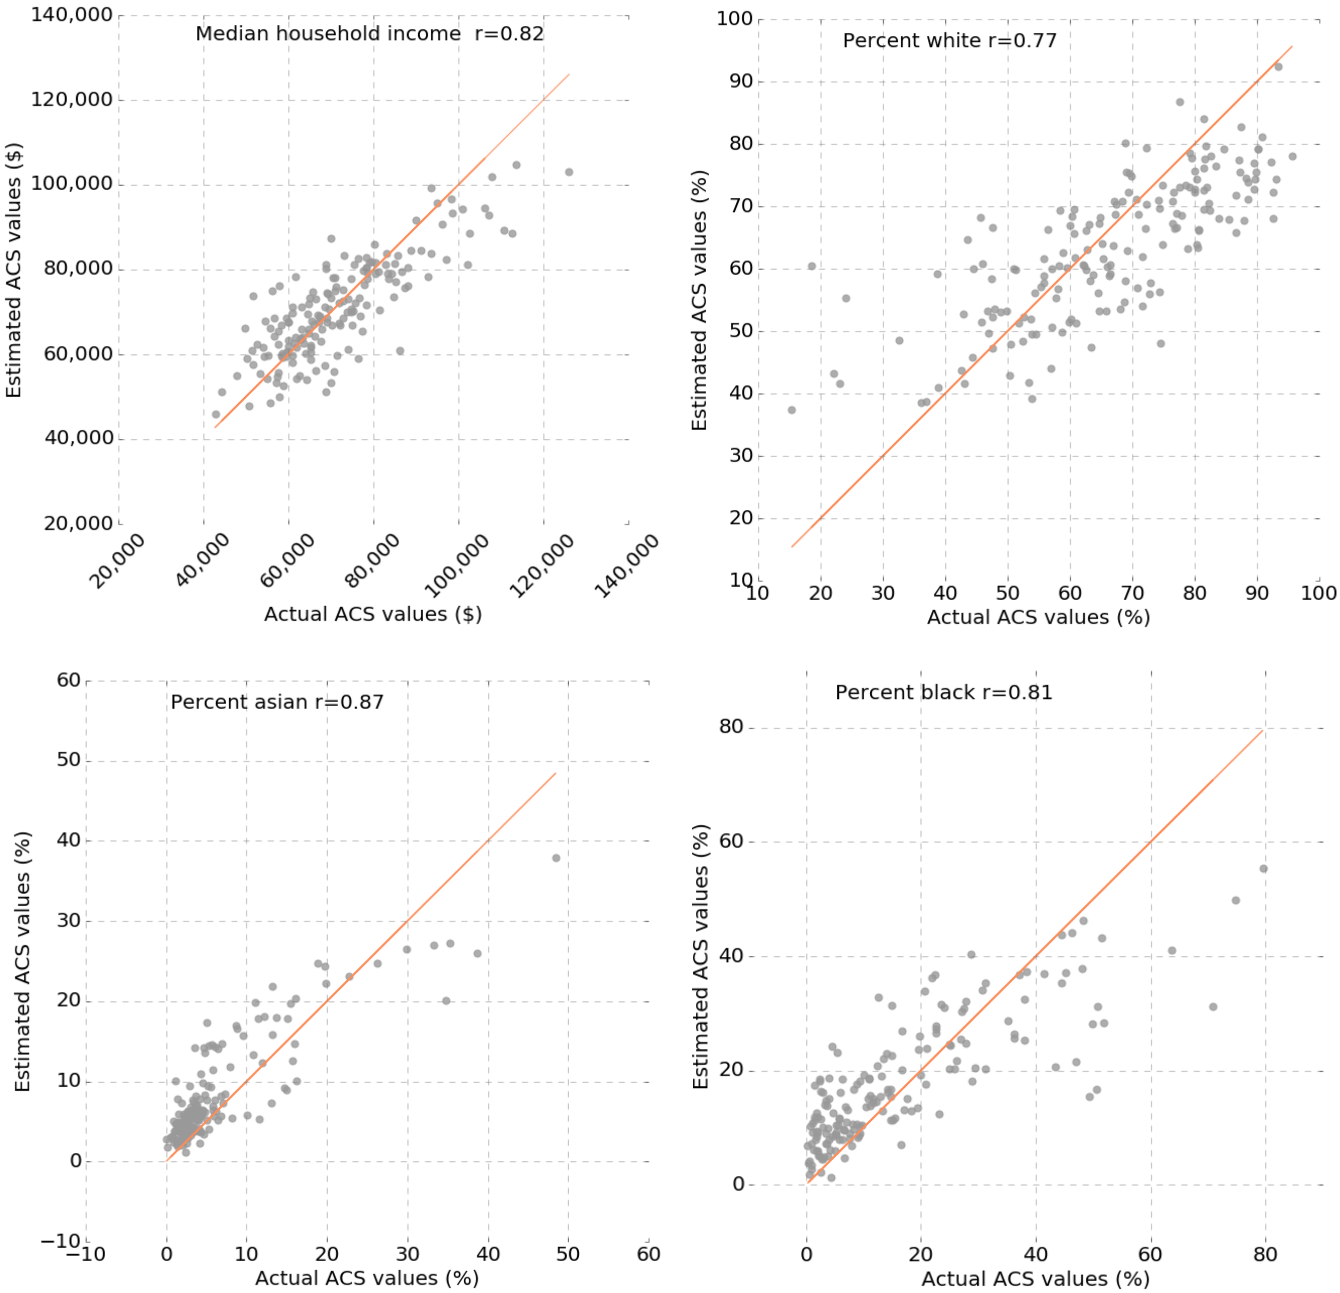

Fig. S3. Scatter plots of ground truth income and race values vs our estimations. Also shown on each plot is the line $y=x$ which corresponds to a perfect predictor.

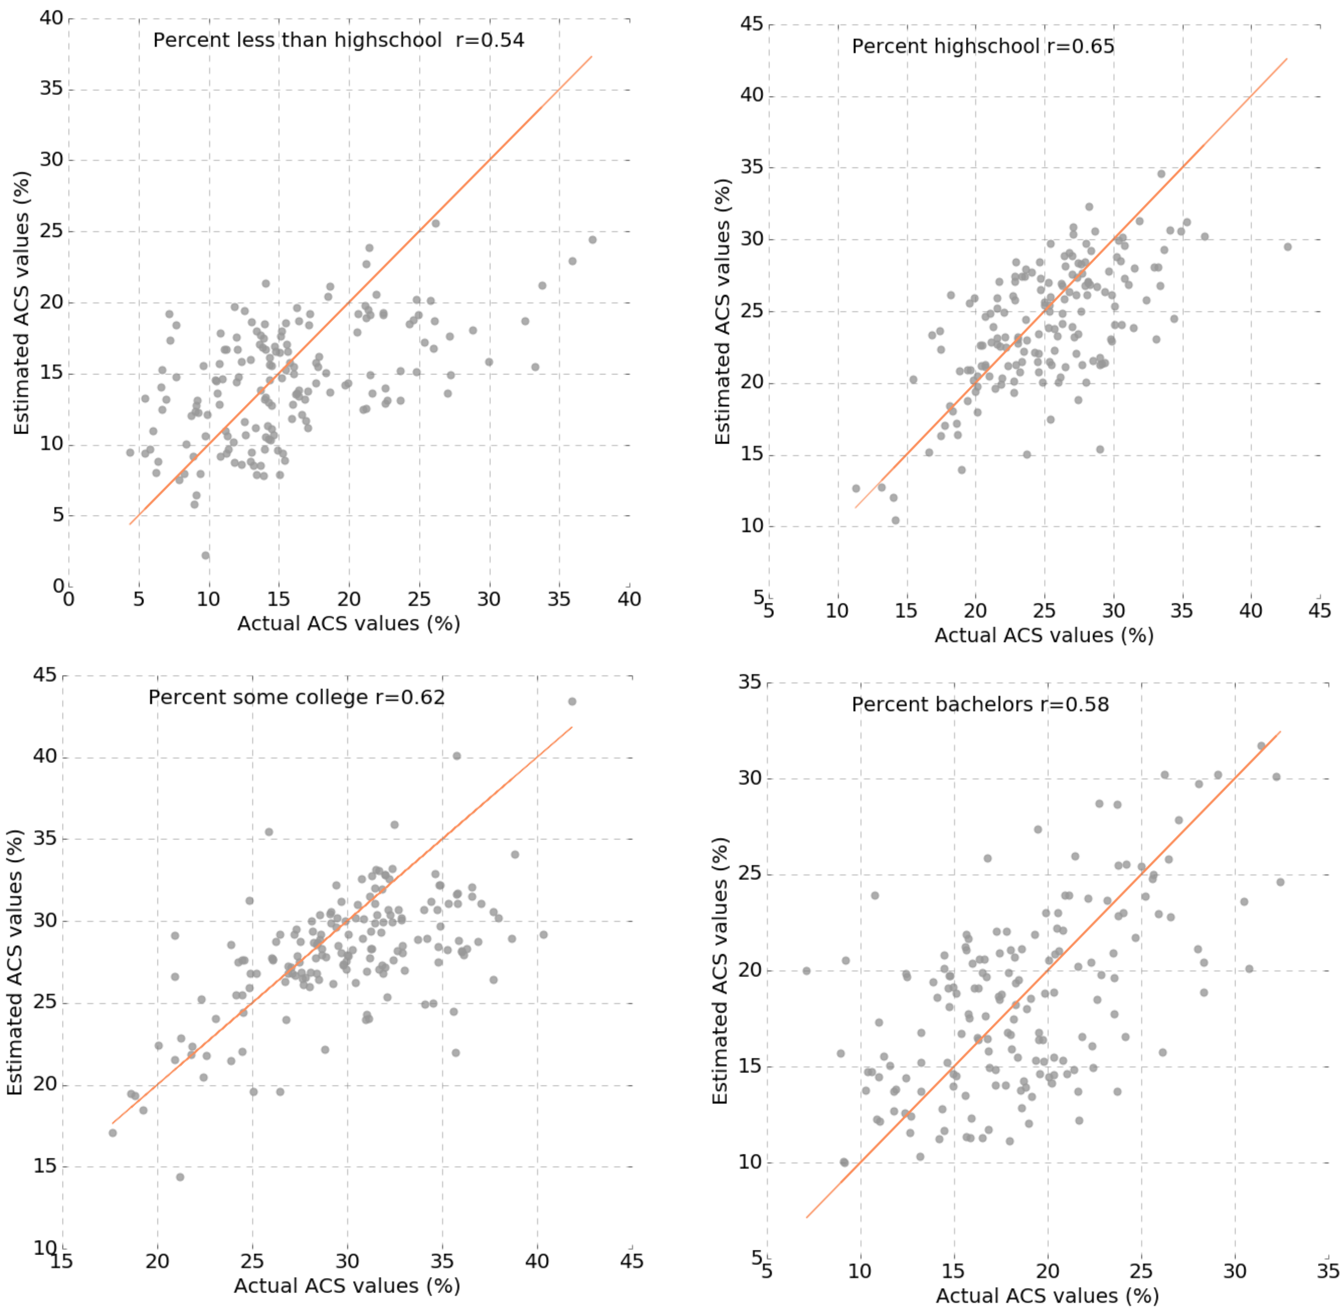

Fig. S4. Scatter plots of ground truth data vs our estimations of educational attainment. Also shown on each plot is the line $y=x$ which corresponds to a perfect predictor.

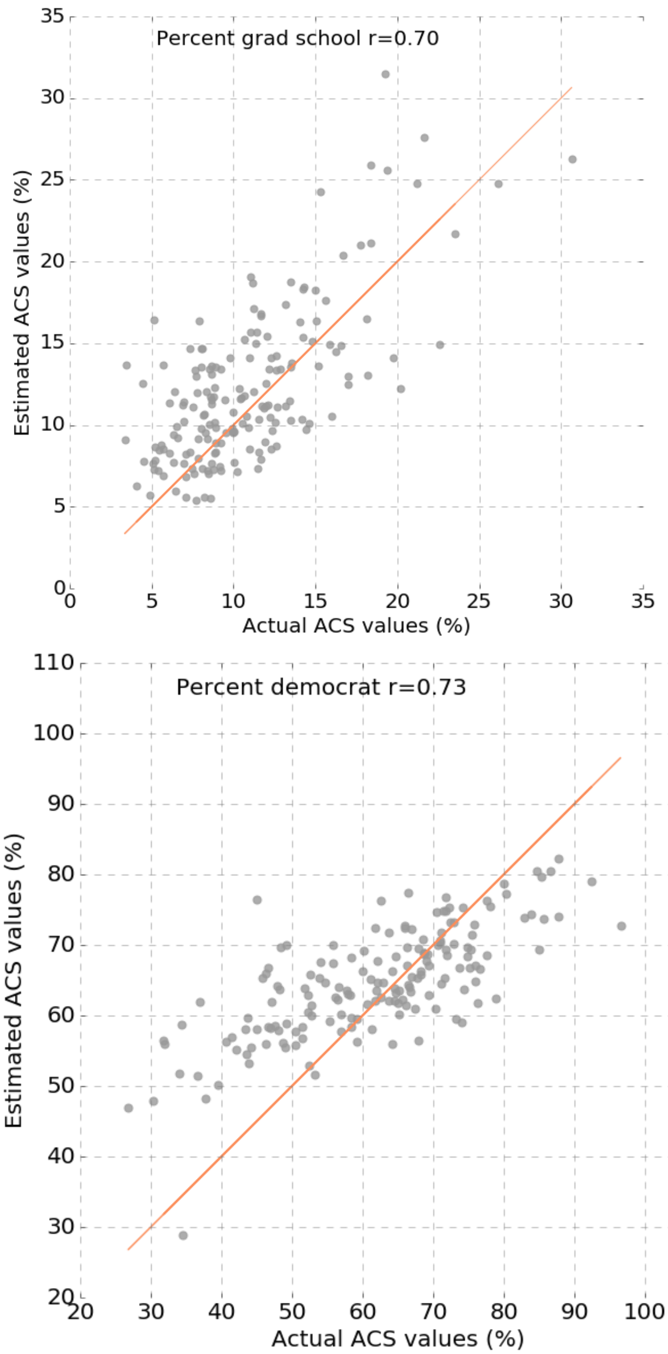

Fig. S5. Scatter plots of ground truth data showing the percentage of people with a graduate school degree vs our estimations, and the percentage of people who voted for Barack Obama in the 2008 presidential election vs our estimations. Also shown on each plot is the line y=x which corresponds to a perfect predictor.

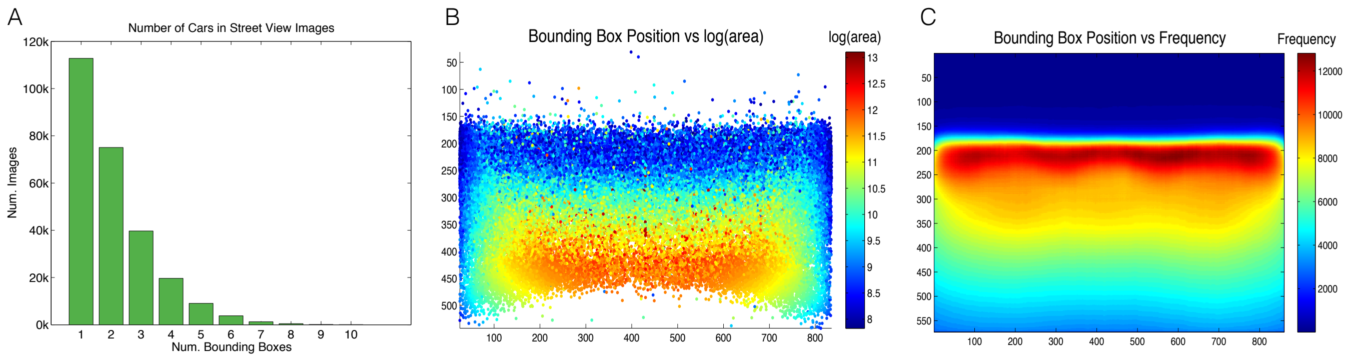

Fig. S6. (A) Histogram of the number of cars annotated in the Street View images, represented by the number of annotated bounding boxes in each image. Images included in these numbers are those images annotated as containing more than zero and less than 11 cars. (B) Bounding box position vs log(area). Each point corresponds to a single bounding box in our training set of Street View images, and the color corresponds to the log of the number of pixels in the bounding box. (C) Bounding box position vs frequency. The color of each pixel indicates the number of bounding boxes in the training set which overlap with that pixel.

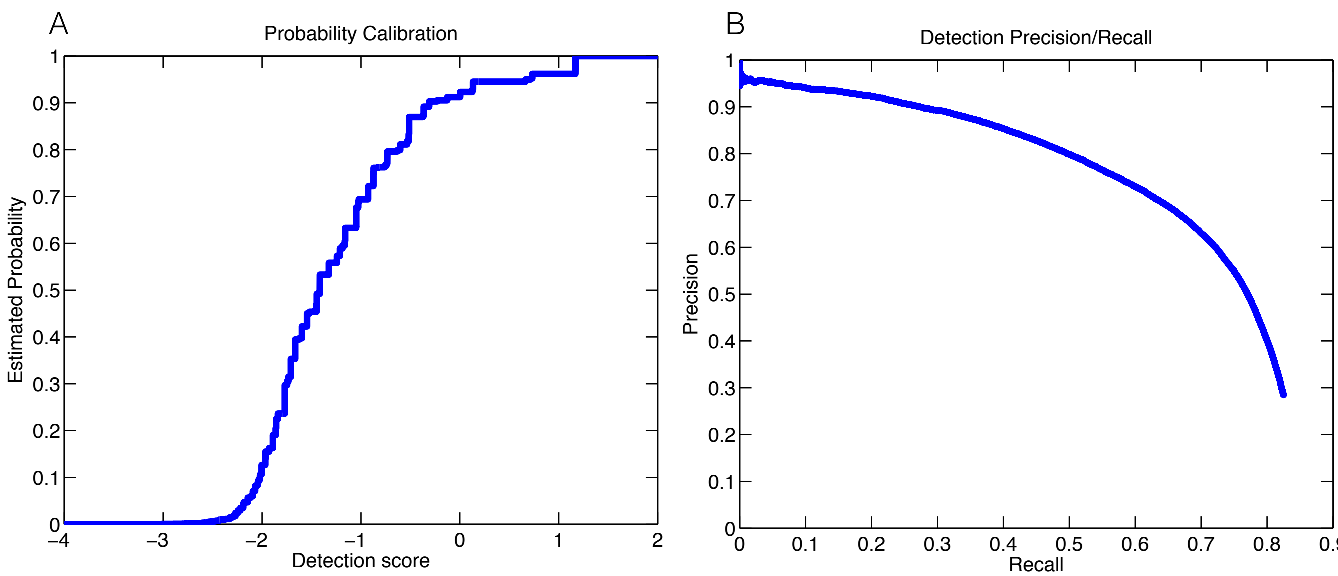

Fig. S7. A. The transformation from detection scores to the probability of the detection being correct (i.e. probability of correctly detecting a car), learned with isotonic regression on the validation set. B. Precision/recall curve for our final detection model on the test set.

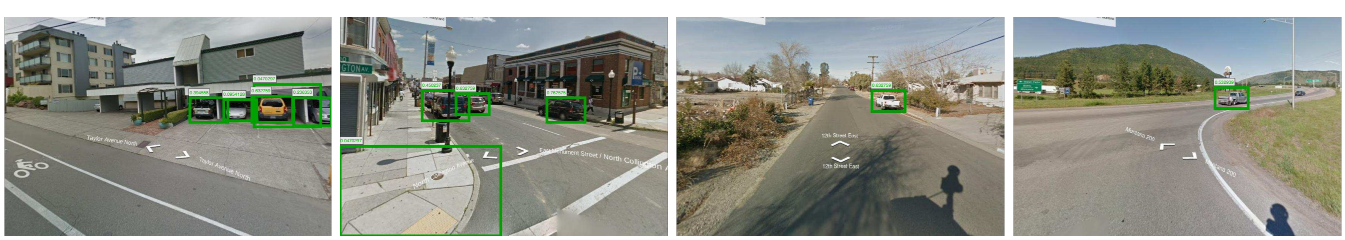

Fig. S8. Example detections with our model on our testing set. Shown in the box around each detection is our estimated probability of the detection having intersection over union greater than 0.5, i.e. counted as correct during detection evaluation.

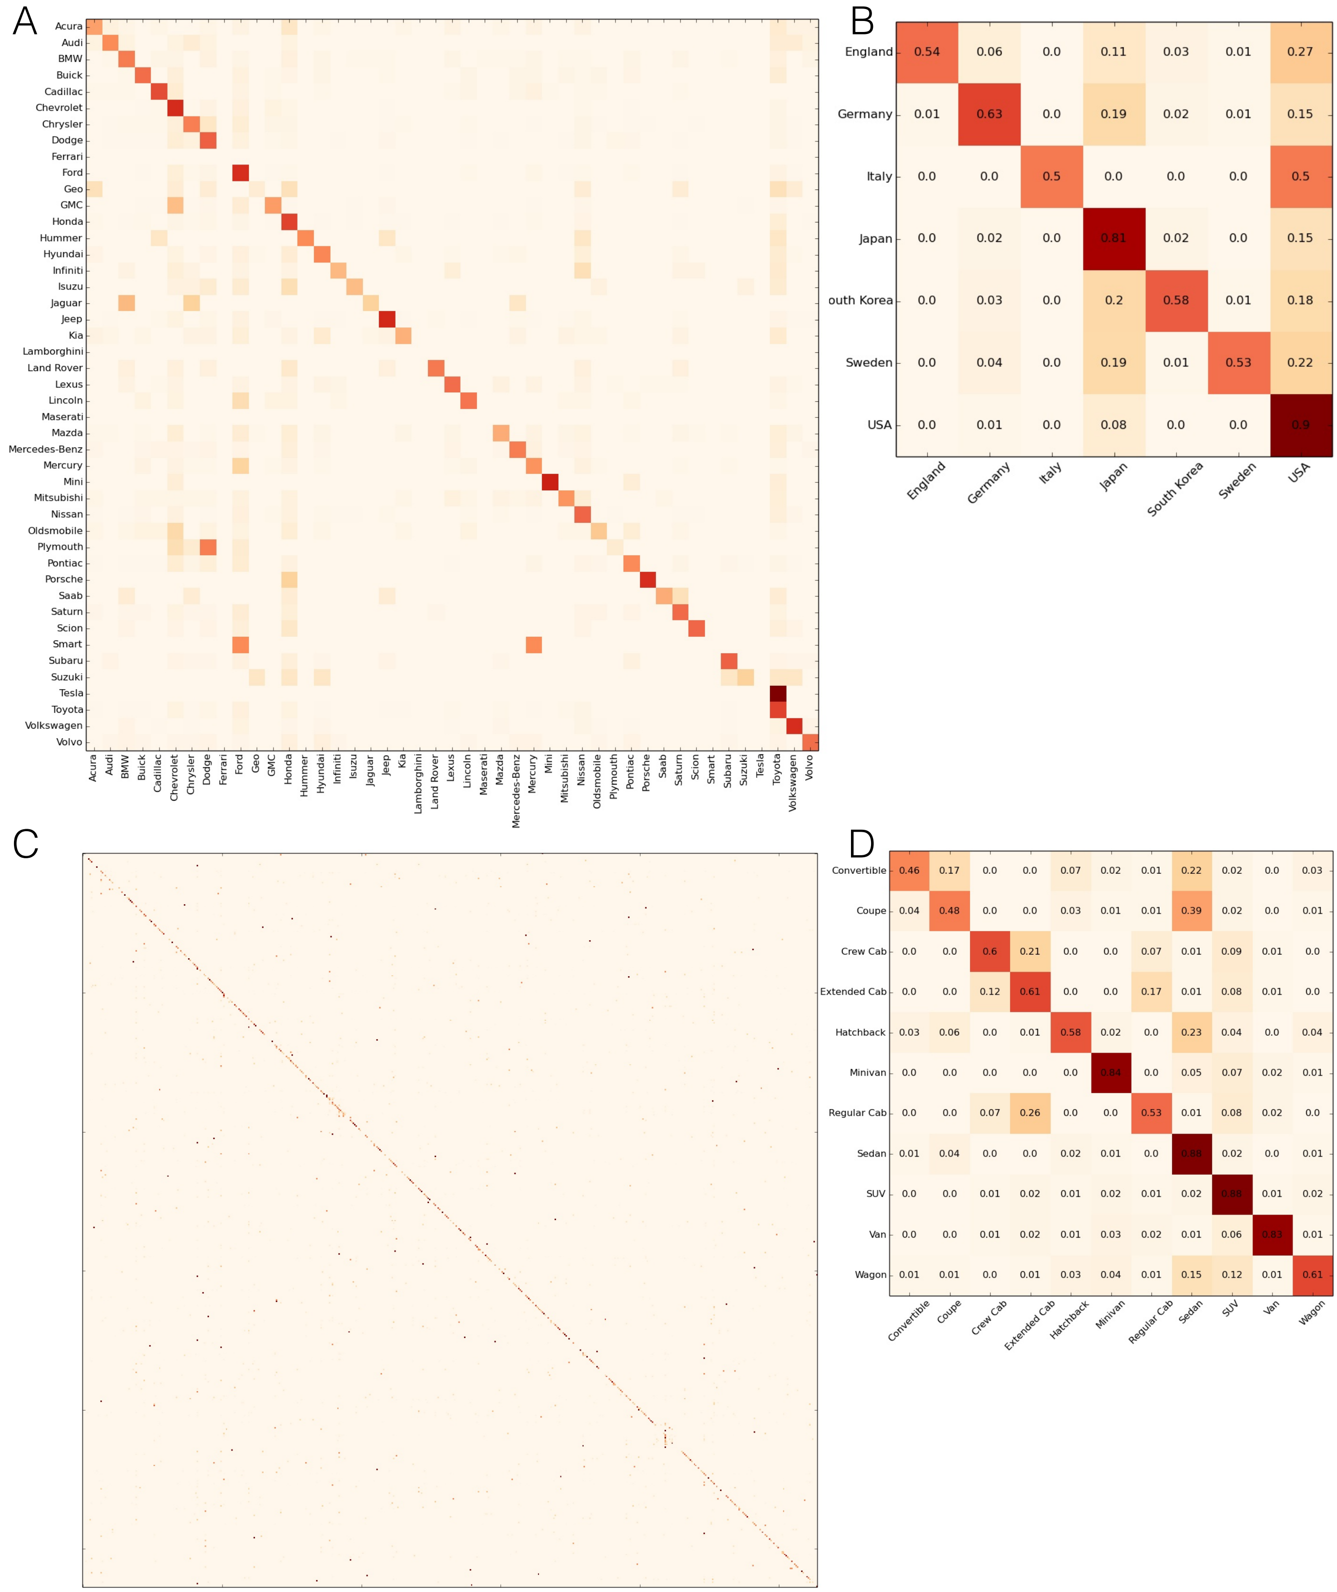

Fig. S9. Confusion matricies of predictions. The entry in row i and column j indicates how many times ground truth attribute i was classified as attribute j. The attributes are A. the make of the car, B. the manufacturing country of the car, C. the model of the car, and D. the body type of the car.

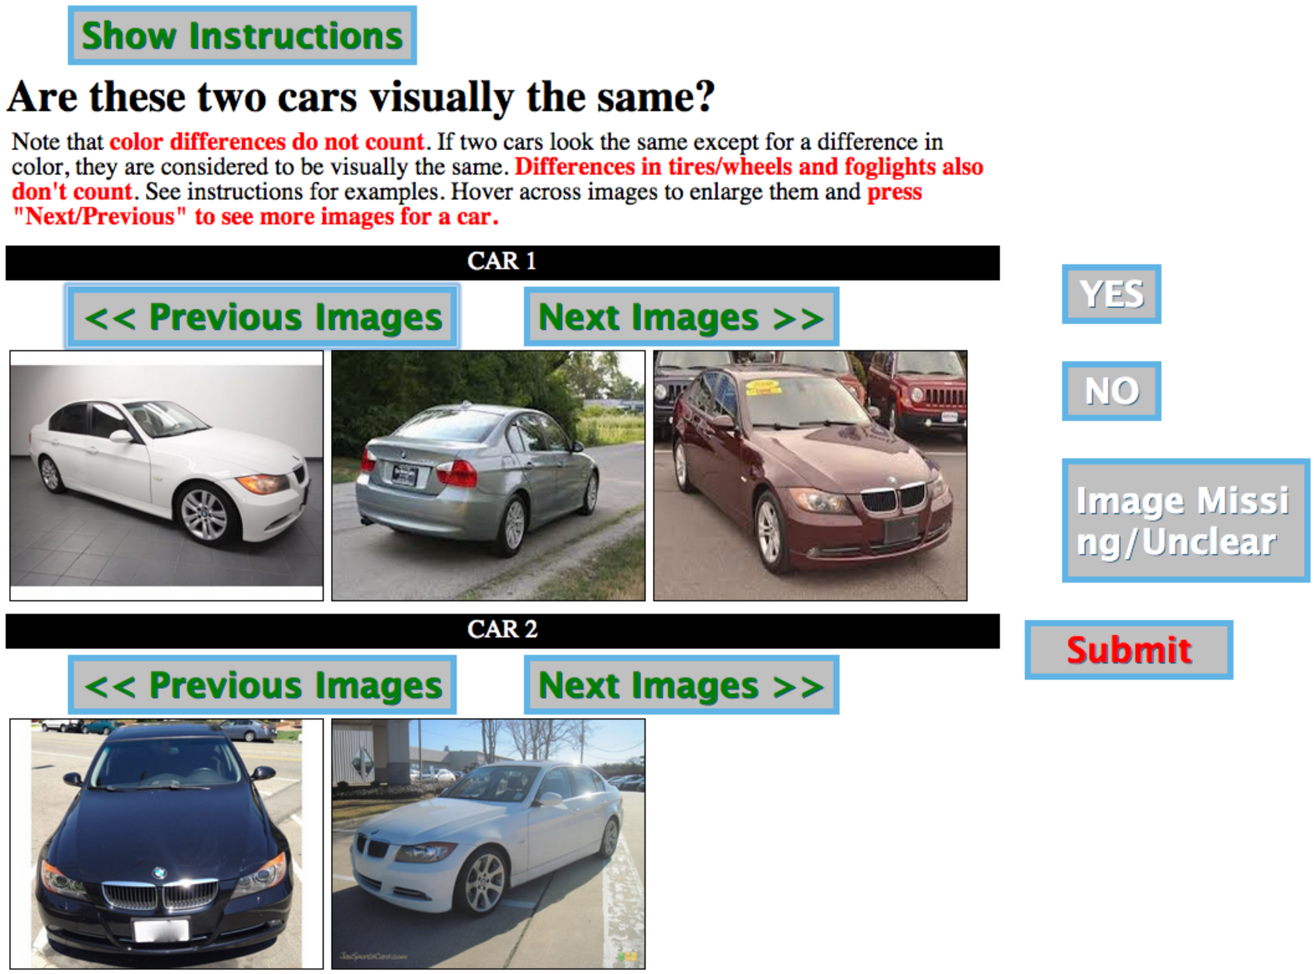

Fig. S10. The Amazon Mechanical Turk (AMT) user interface for grouping visually indistinguishable pairs of classes. The user is asked whether or not the two cars are visually distinct with an option to view more detailed instructions.

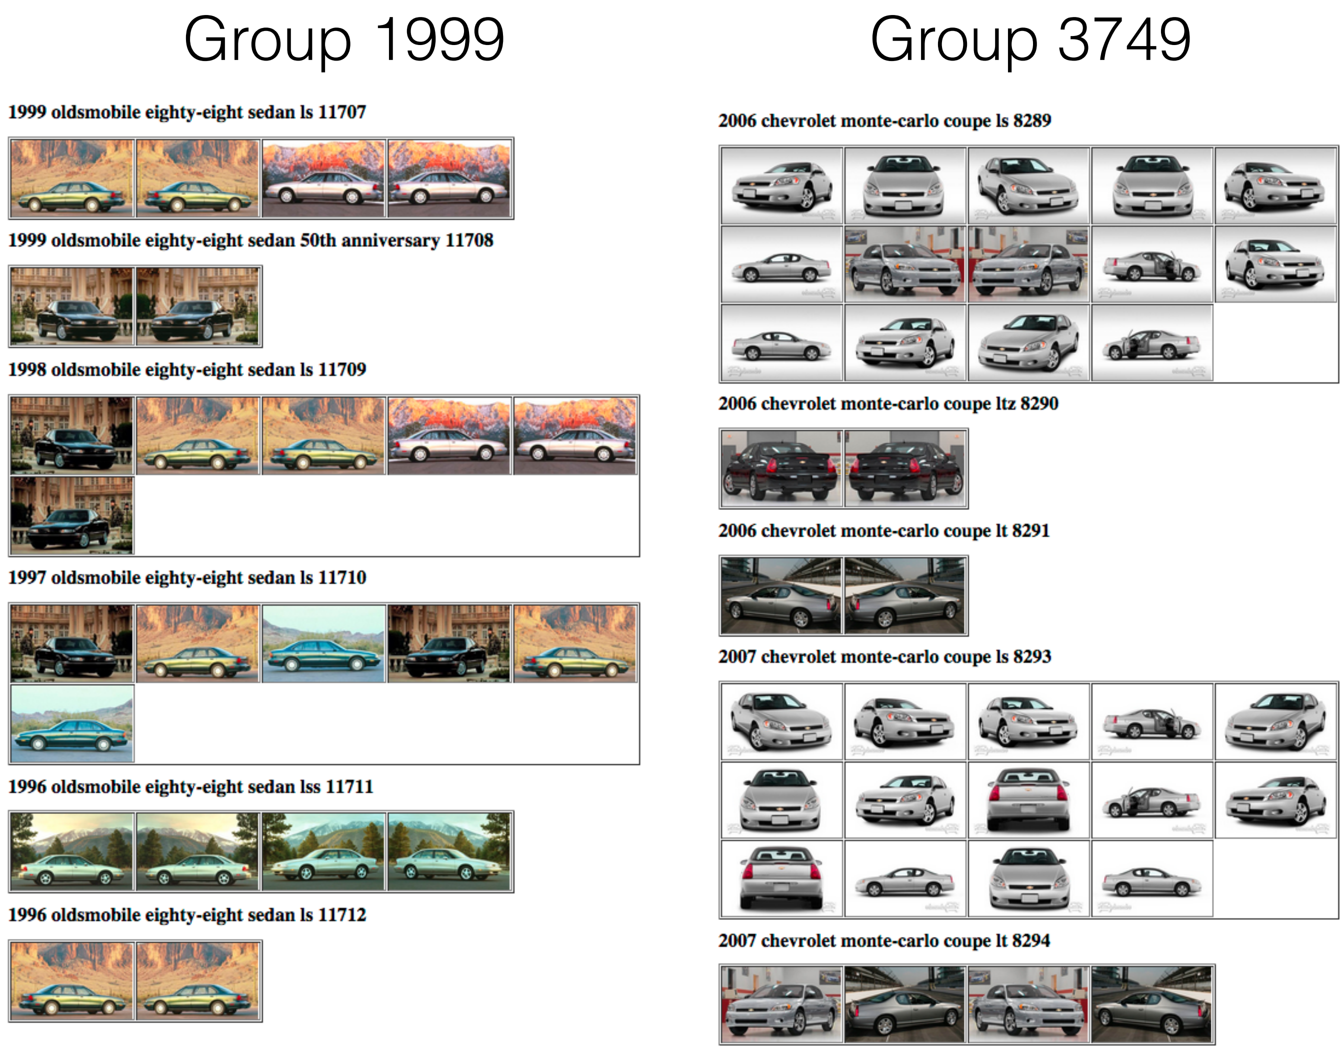

Fig. S11. Two examples of classes and the different types of visually indistinguishable cars in each class. Each column is a unique class. The first column shows cars assembled into group 1999 whereas the second column shows those in group 3749.

Fig. S12. Left, Middle: Examples of product shot images unsuitable for our dataset, as they are either extremely close up (left) or are of the interior of the car (middle). In order to be suitable for recognition, an image must be of the exterior of the car and the car must be entirely visible (right).

Fig. S13. Screenshot of the user interface for labeling images containing a car viewed from the exterior, deployed on Amazon Mechanical Turk. Below the instructions are a set of images, and the user is tasked with clicking on the images containing a single prominent vehicle, viewed from the outside. Images the user clicks are moved to the panel on the right side of the screen, and clicks can be undone by clicking on the image in the right panel.

Fig. S14. An example of the unwarping that needs to be done on images retrieved from Street View. Left: an image from Street View as initially scraped. The image appears warped (e.g. straight lines in the real world are not straight in the image) due to the equirectangular projection used to store spherical panoramas. Right: the result of undoing this projection, which we do before using the images any further.

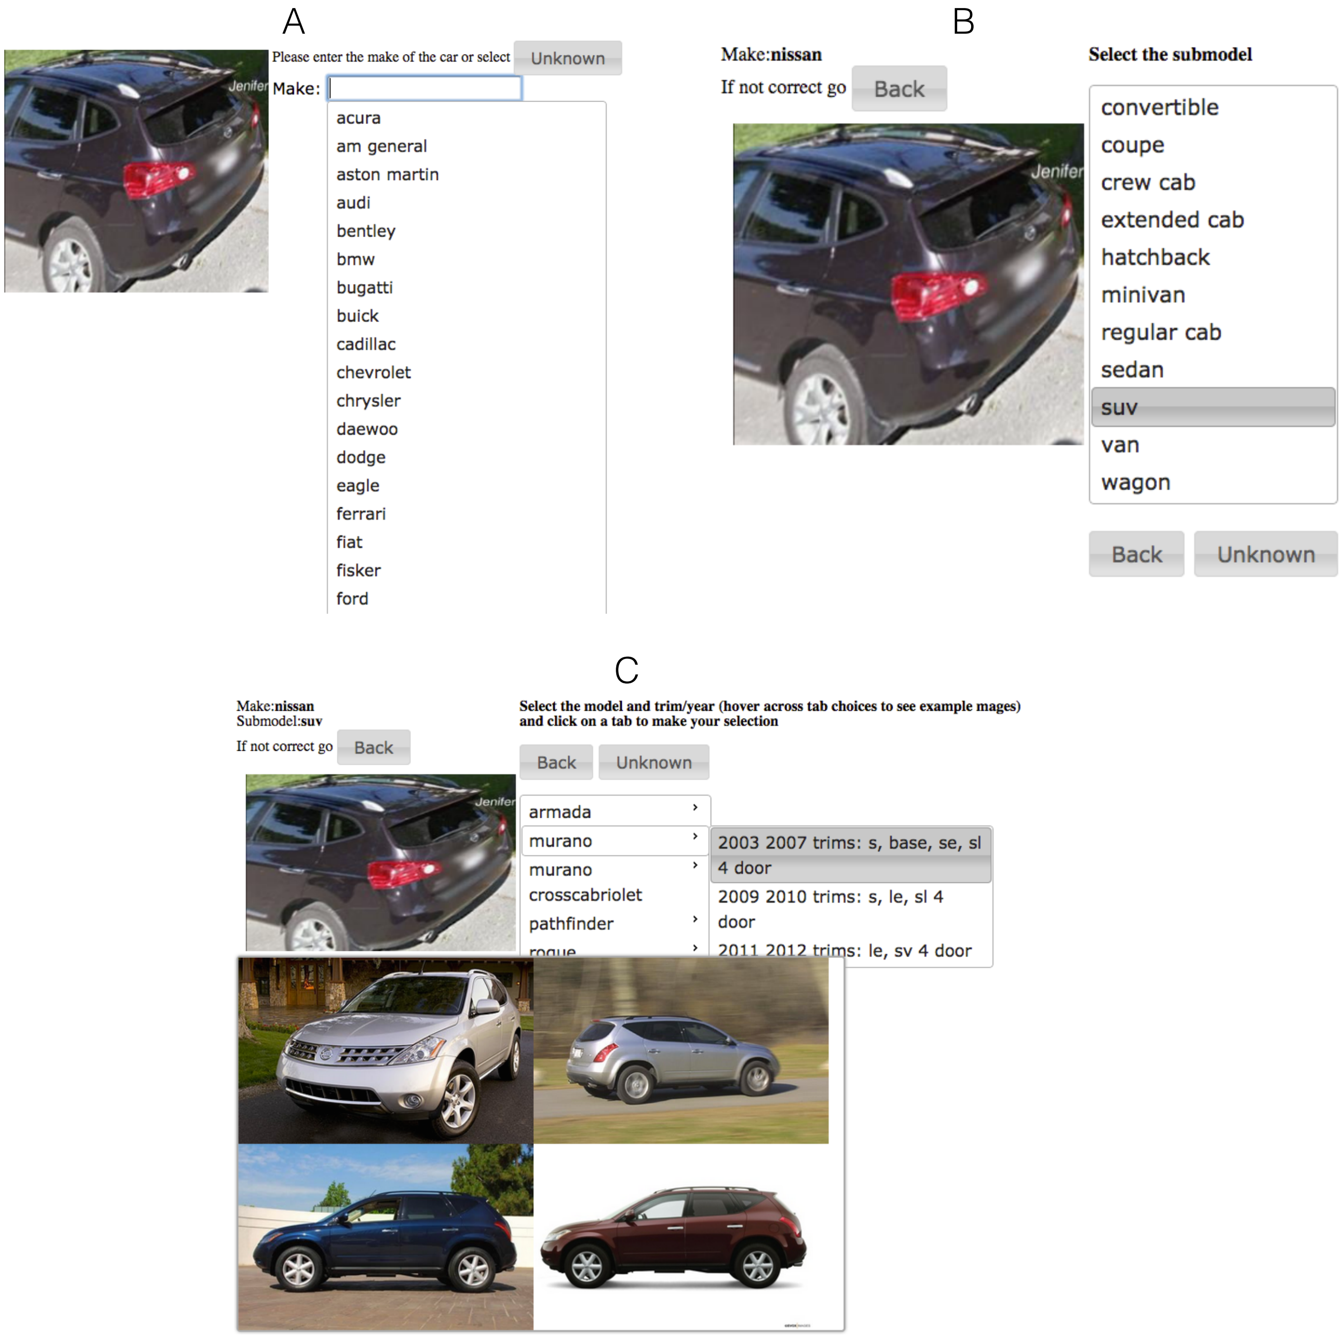

Fig. S15. Screenshots of the user interface for hierarchically annotating Street View images with car categories. A. The expert is first asked to identify the make. B. The next step in the task is to identify the body type of the car which is called submodel in the task. C. Once the body type is identified we provide a list of classes for the selected make and body type. Example images of each class are also shown to aid the user in identification.

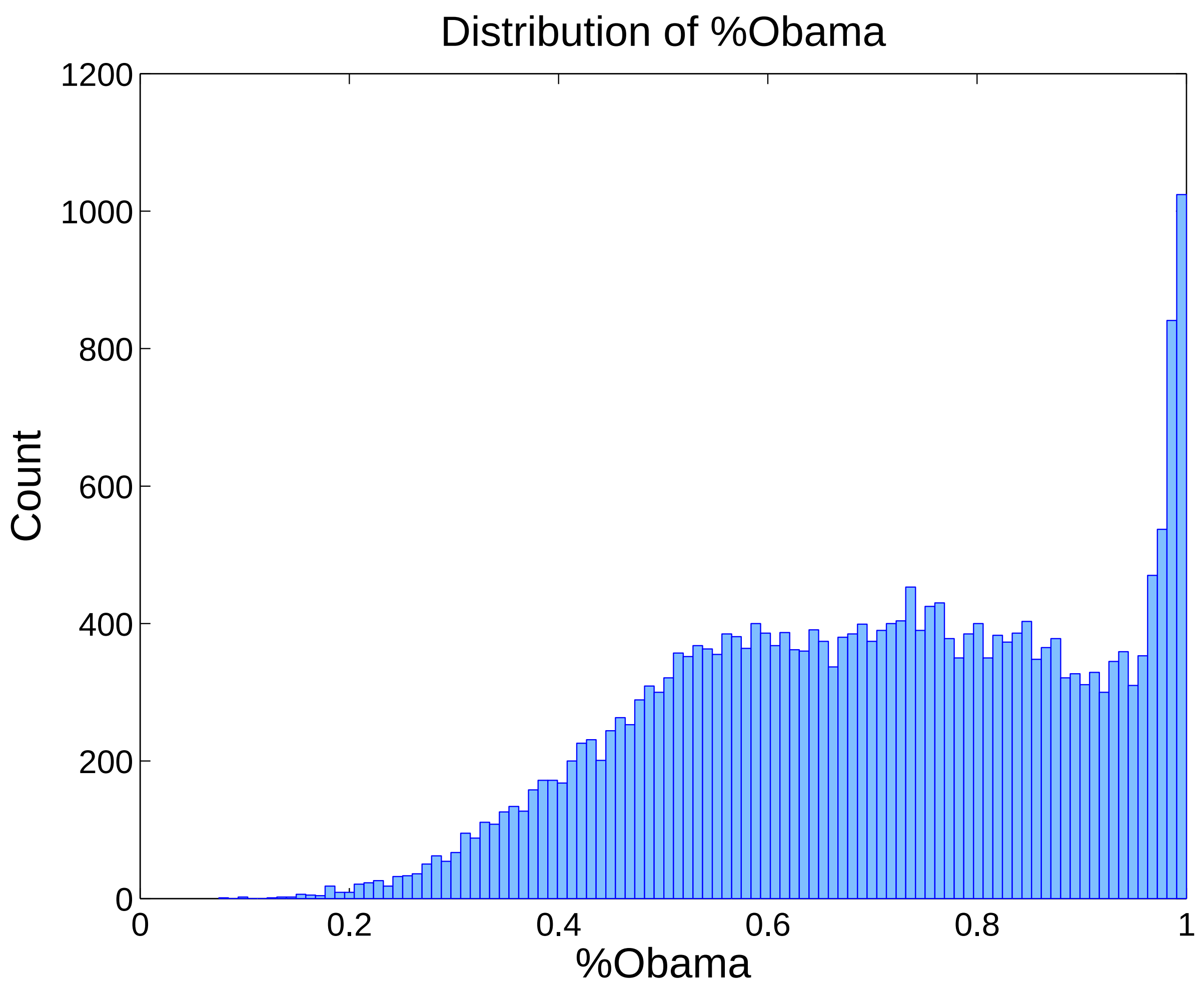

Fig. S16. Histogram of the fraction of votes cast for Barack Obama vs. John McCain in the 2008 presidential election for the precincts in our dataset.

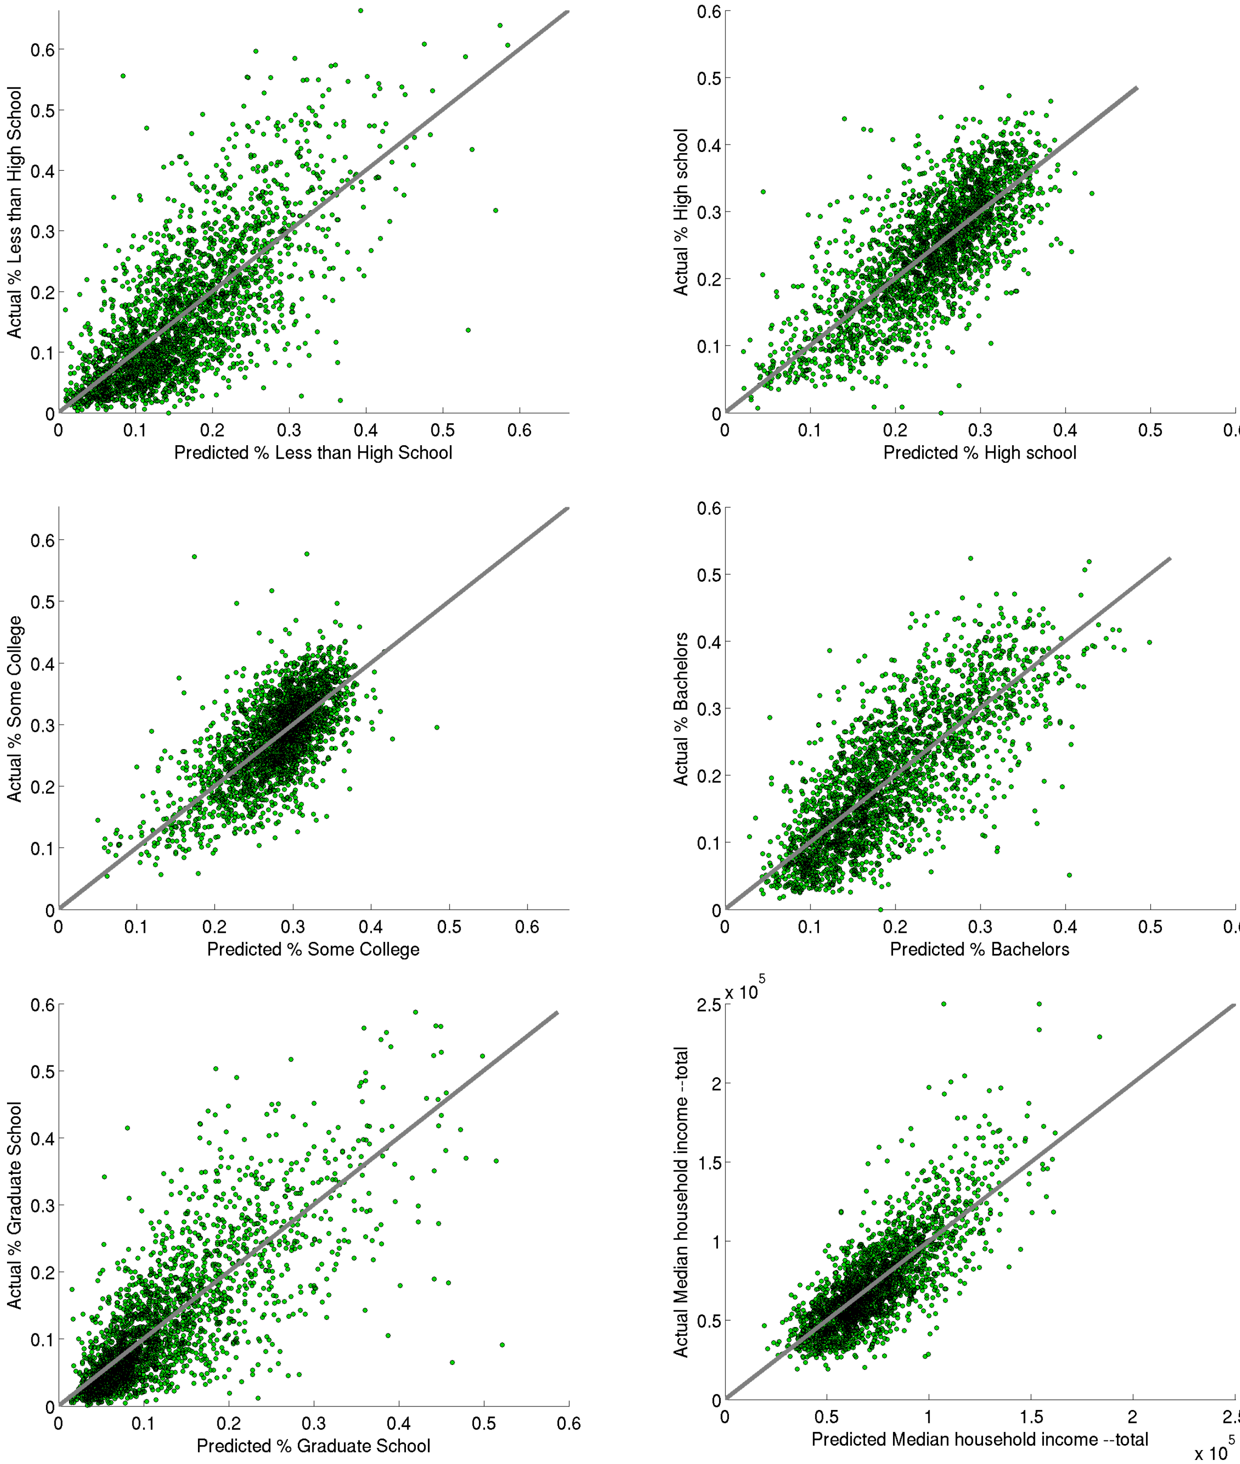

Fig. S17. Scatter plots of cross validated actual versus predicted education and median household income levels. Also shown on each plot is the line $y = x$, which corresponds to a perfect predictor.

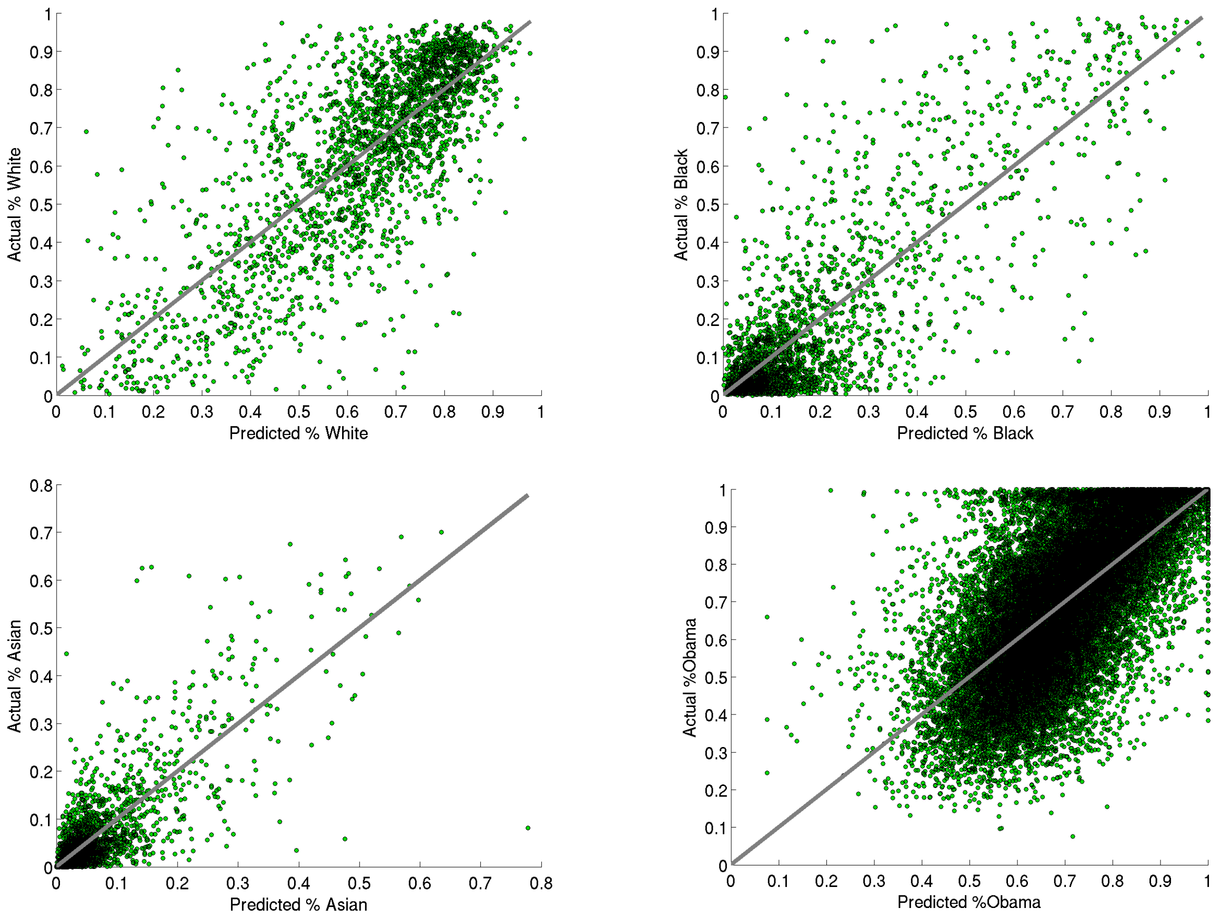

Fig. S18. Scatter plots of crossvalidated actual versus predicted distributions of race. Also shown on each plot is the line $y = x$, which corresponds to a perfect predictor. The last scatter plot shows cross validated actual vs predicted voting results.

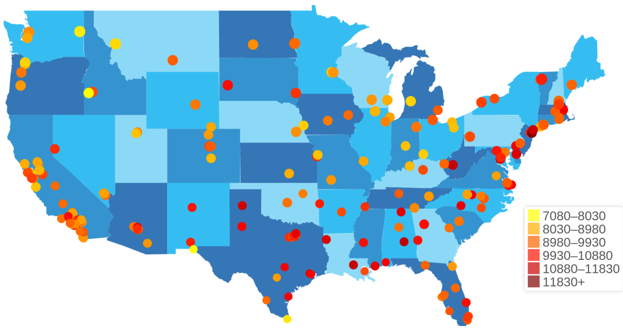

(a) Average car price

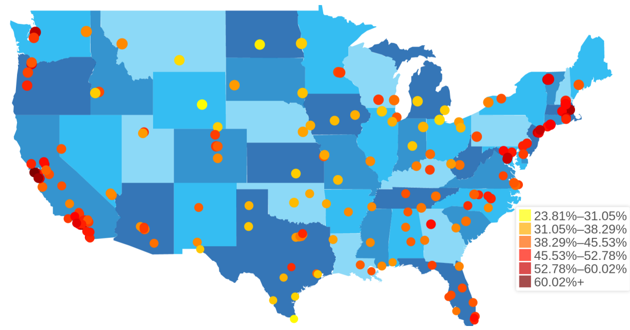

(b) Percentage of foreign cars

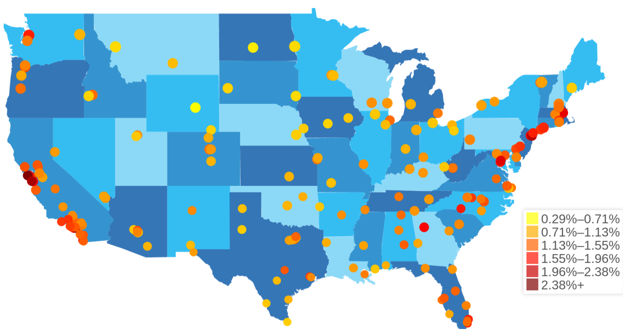

(c) Percentage of BMWs

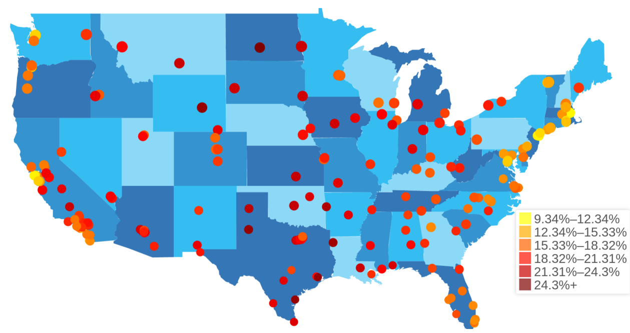

(d) Percentage of Chevrolets

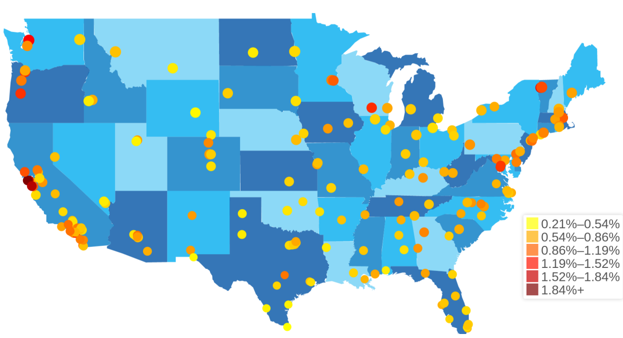

(e) Percentage of Toyota Prius

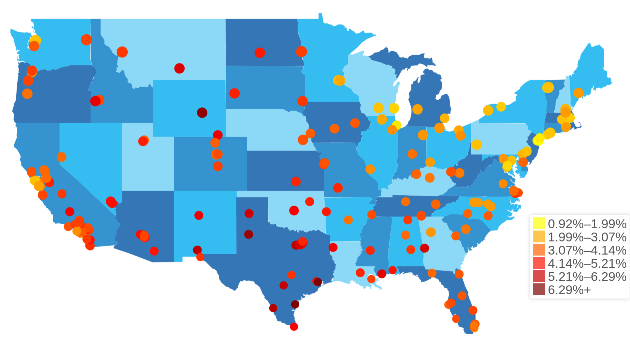

(f) Percentage of Ford F-150

Fig. S19. Maps of a variety of car attributes as measured across the cities in our dataset. Each point corresponds to one city. Not shown: Alaska and Hawaii.

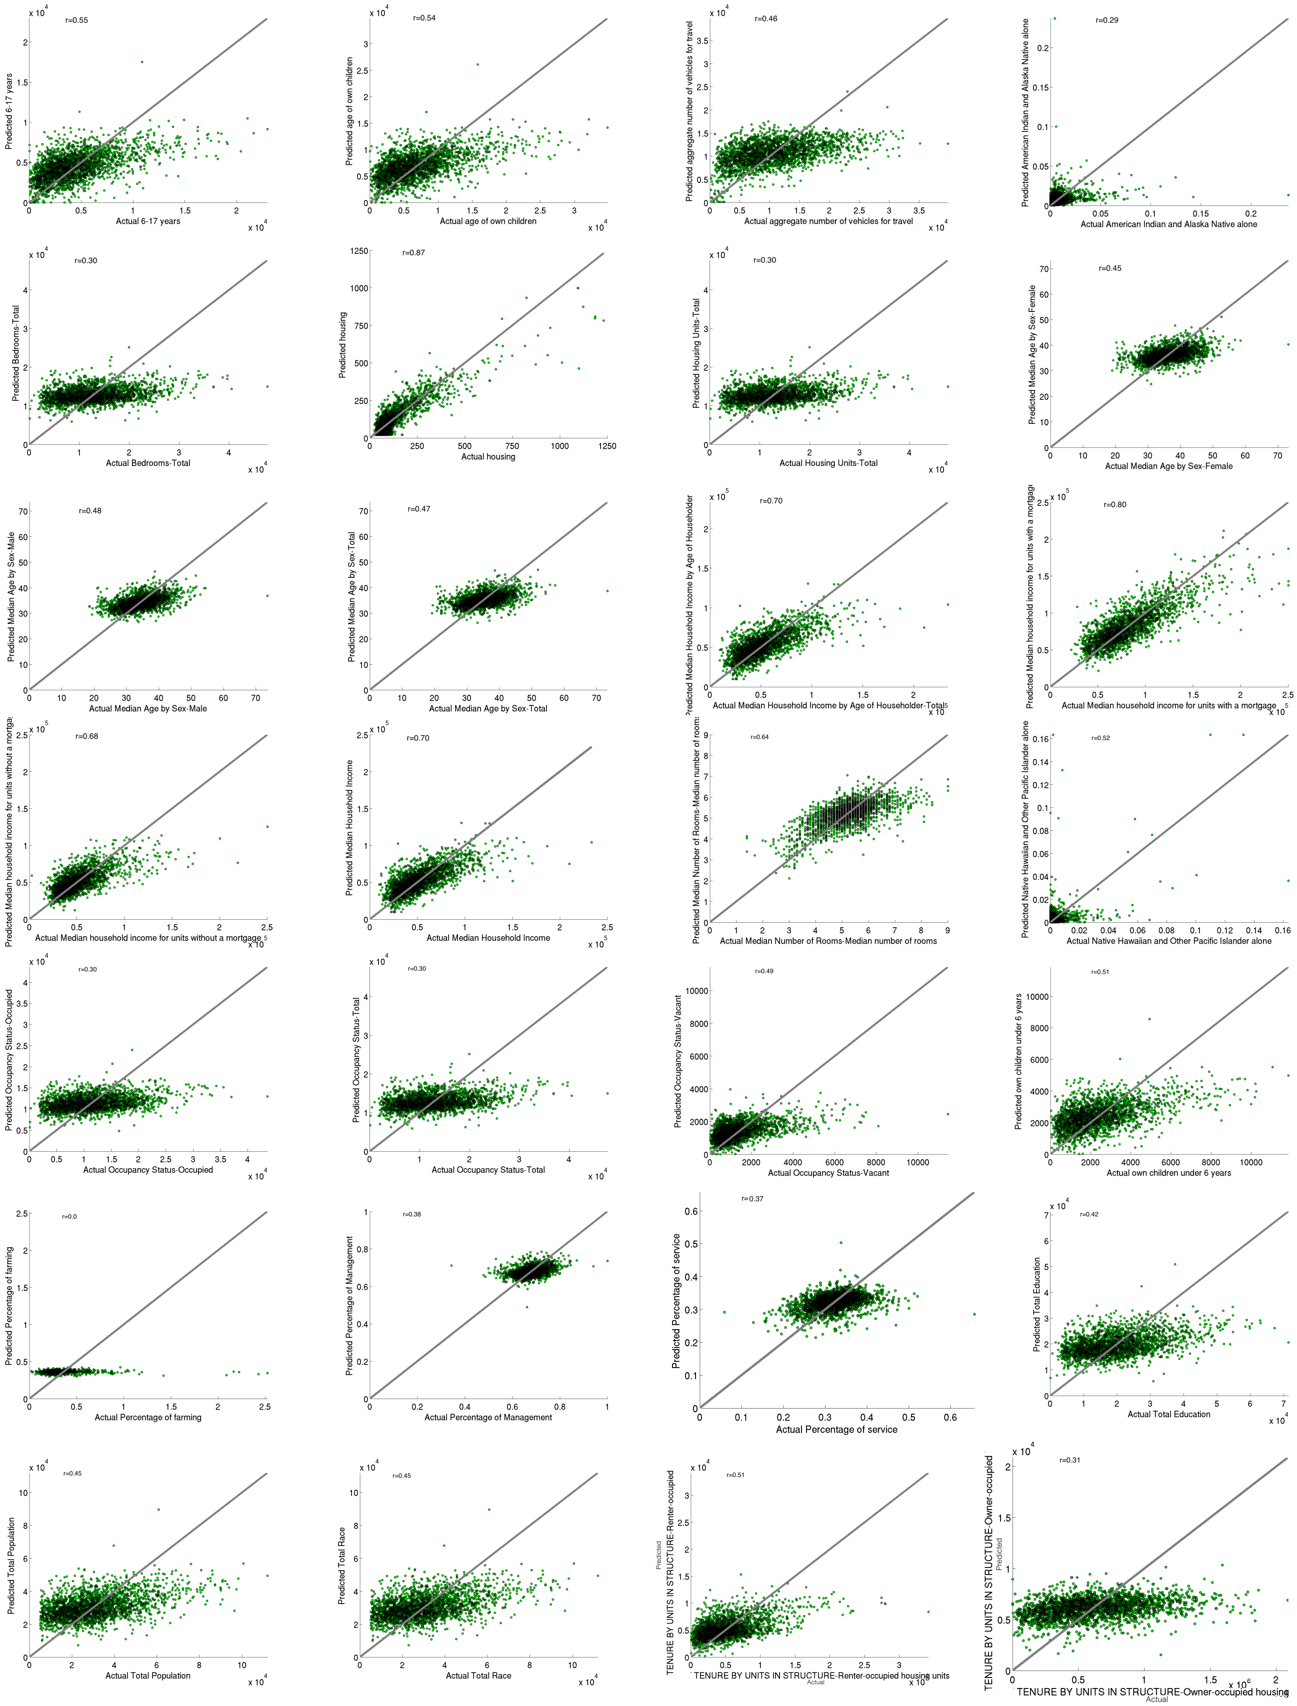

Fig. S20. Our methodology applied to predicting additional ACS attributes not discussed in the main text. Note: this is an application of our methodology to infer variables with no refinement. Best viewed after zooming in.

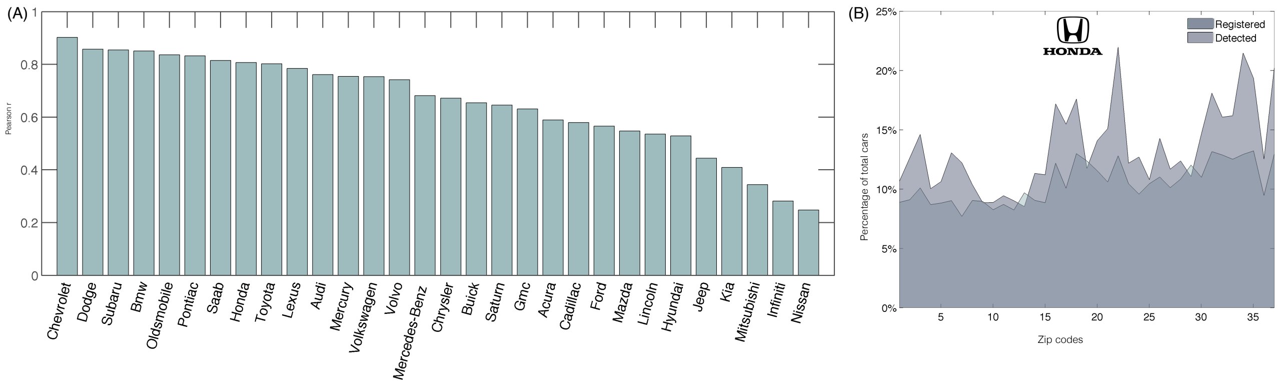

Fig. S21. A. Correlation between our detected makes and Massachusetts DMV data for 30 makes. B. The percentage of registered Hondas vs. Those we detect.

Fig. S22. The percentage of registered makes in each zip code (according to DMV data) vs. those we detect for each make.

Fig. S23. The percentage of registered makes in each zip code (according to DMV data) vs. those we detect for each make (continued from prior page).

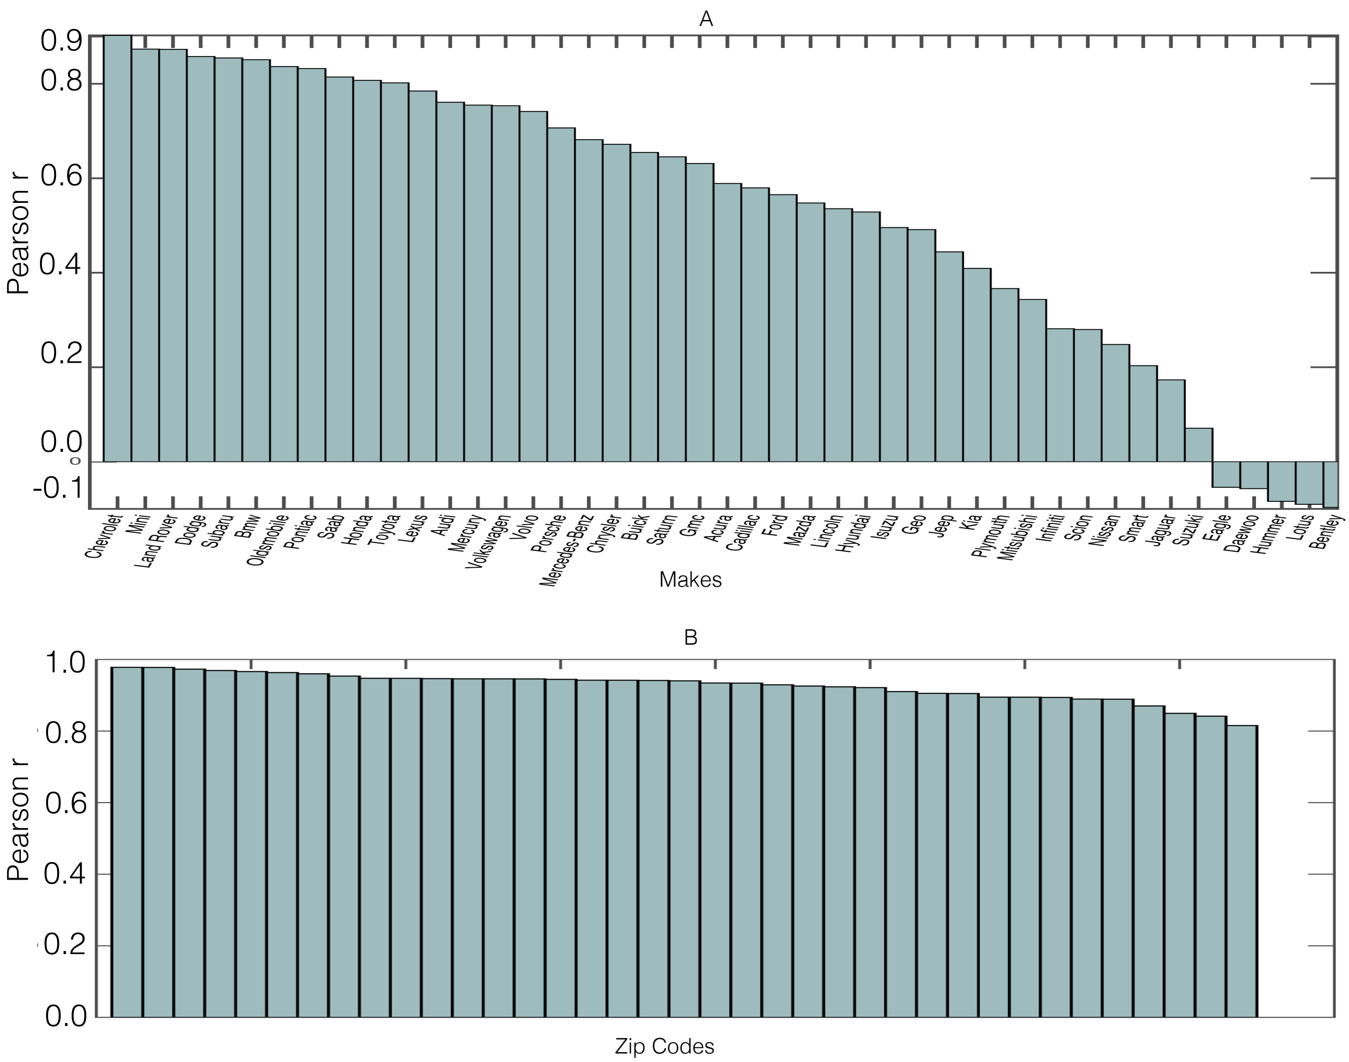

Fig. S24. A. The correlation between the distribution of detected and registered car makes across zip codes in the three cities in Massachusetts, Boston, Springfield, and Worcester. We show results for all 45 makes in the intersection of our and DMV data. B. The correlation between the distribution of detected and registered car makes in each zip code for the three cities in Massachusetts (Boston, Springfield, Worcester). All zip codes have correlation greater than 0.8.

A

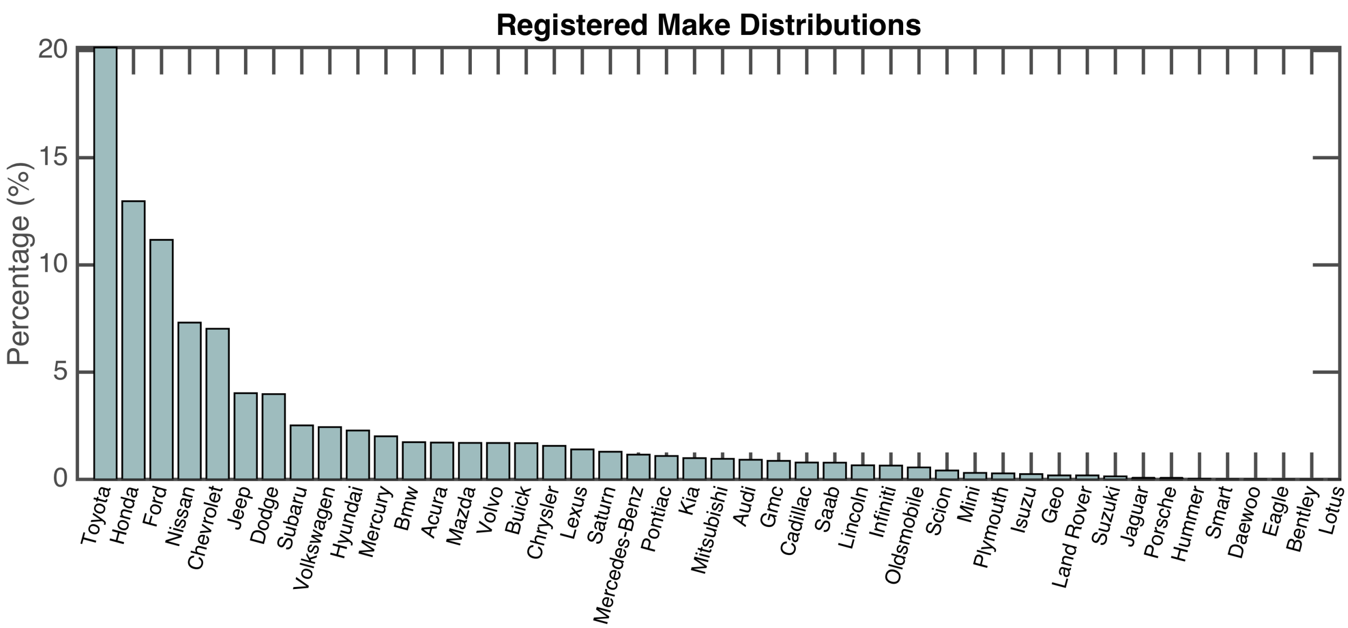

B

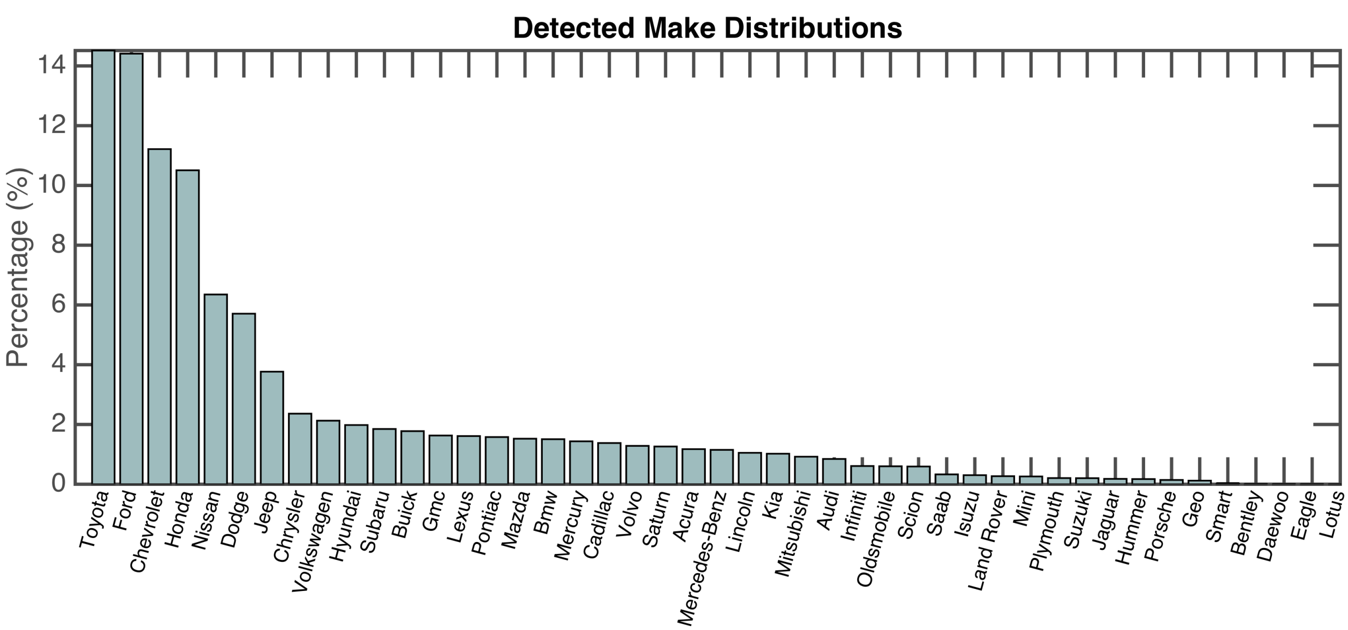

Fig. S25. (A) The Distribution of registered makes in Boston, Springfield, and Worcester Massachusetts. (B) The distribution of detected makes in Boston, Springfield, and Worcester Massachusetts.

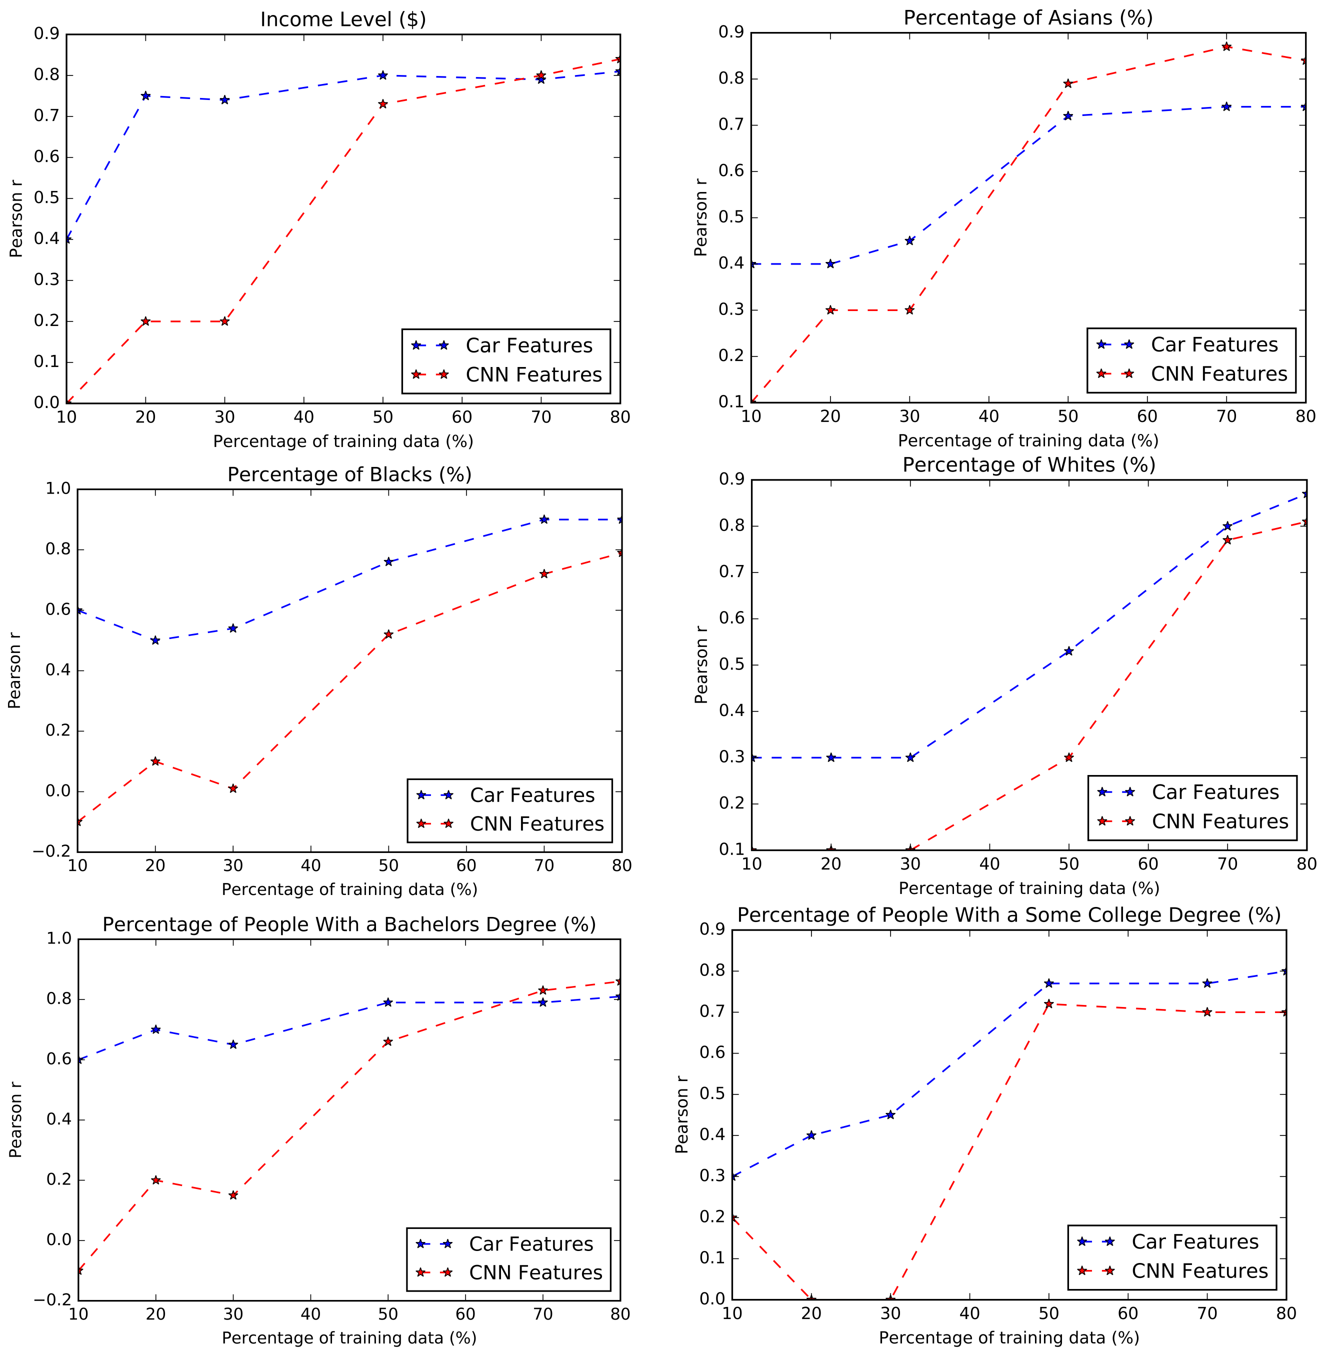

Fig. S26. Comparisons of our methodology with a baseline using features from a convolutional neural network pretrained on ImageNet. This experiment was carried out using 10% of our data (approximately half a million images). The x axis shows the percentage of training data and the y axis shows the Pearson correlation coefficient between actual and predicted values.

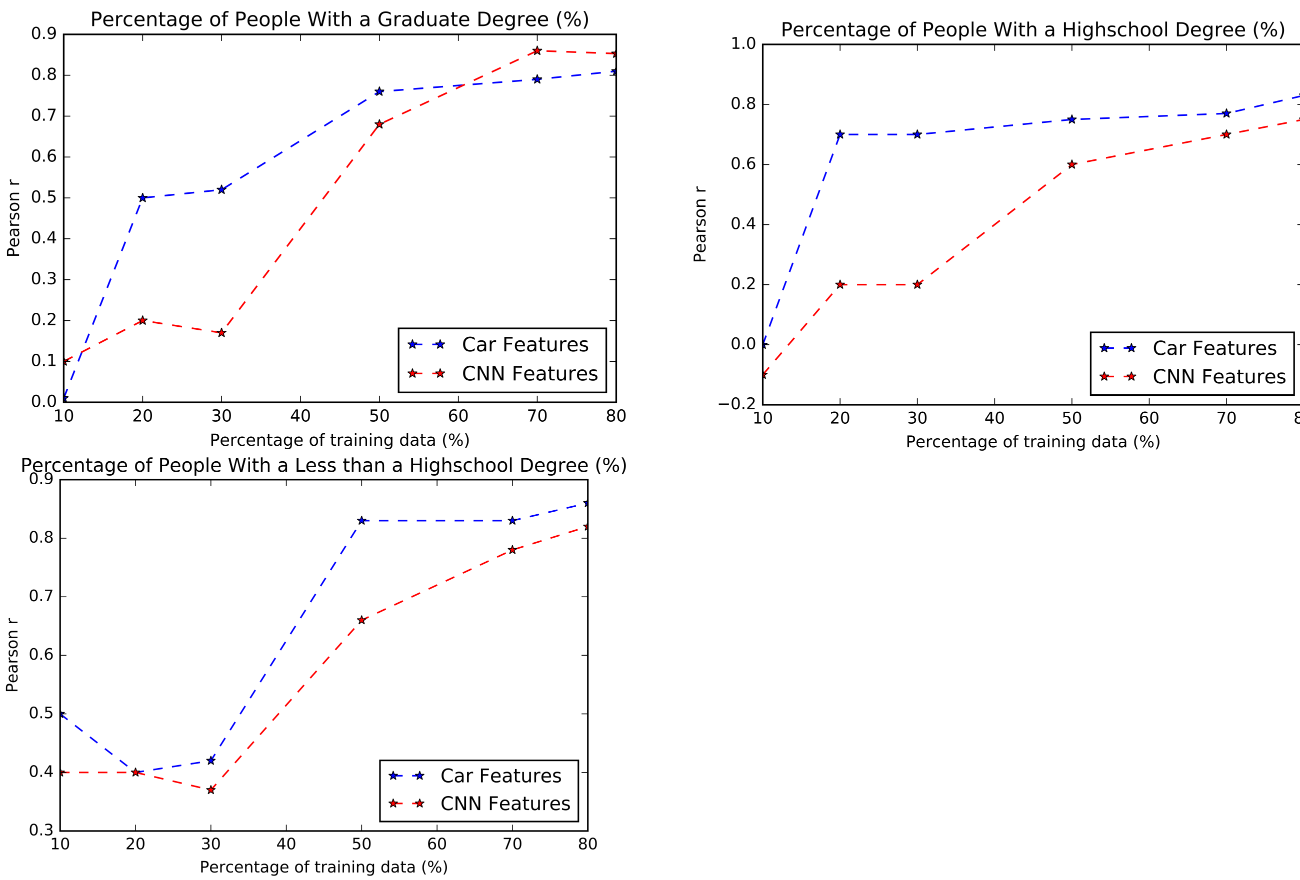

Fig. S27. Comparisons of our methodology with a baseline using features from a convolutional neural network pretrained on ImageNet. This experiment was carried out using 10% of our data (approximately half a million images). The x axis shows the percentage of training data and the y axis shows the Pearson correlation coefficient between actual and predicted values (continued from prior page).

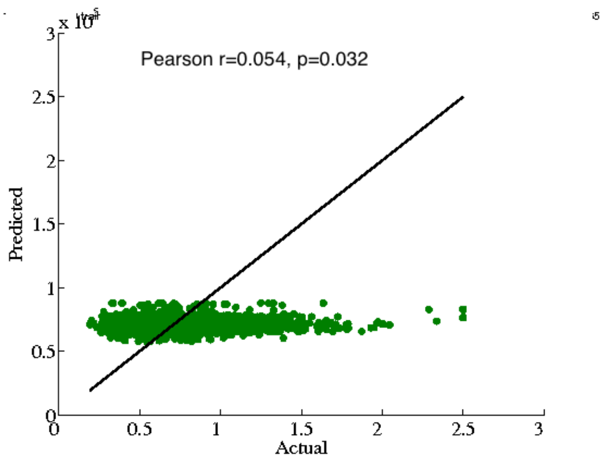

(a) Median Household Income ($)

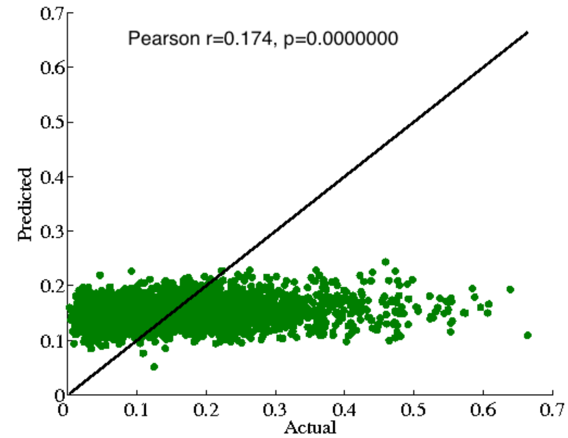

(b) Percentage of People with Less than a High school Education (Ratios between 0 and 1)

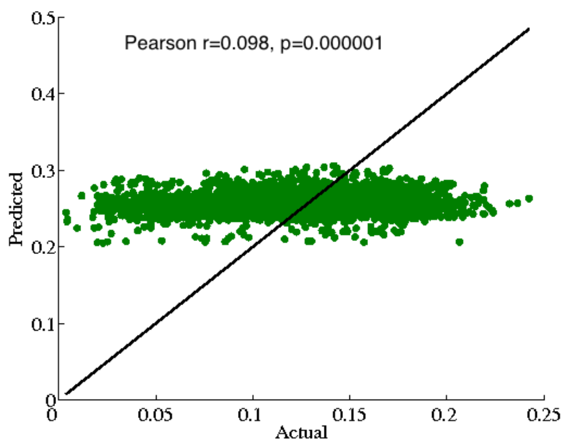

(c) Percentage of People with a High school Education (Ratios between 0 and 1)

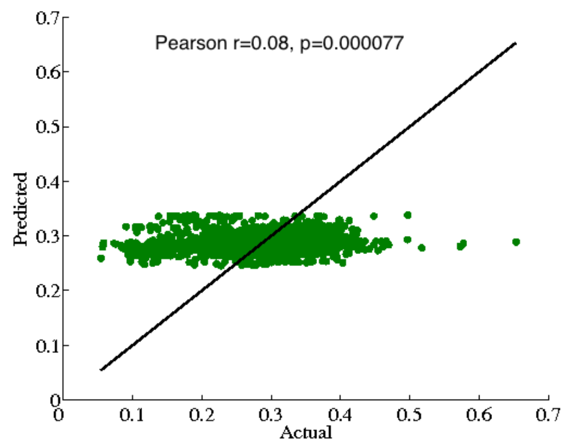

(d) Percentage of People with Some College Education (Ratios between 0 and 1)

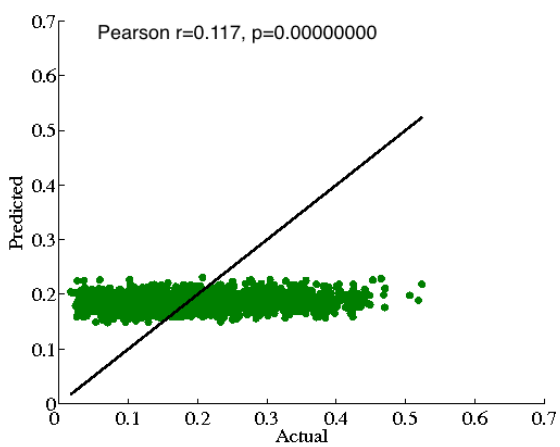

(e) Percentage of People with a Bachelors Degree (Ratios between 0 and 1)

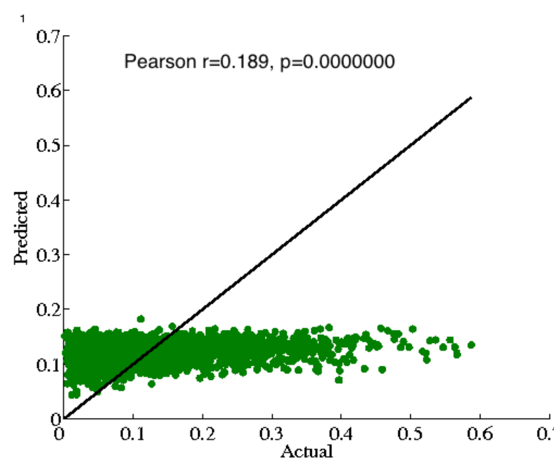

(f) Percentage of People with a Graduate Degree (Ratios between 0 and 1)

Fig. S28. Zip code level income, race and education variables from the ACS inferred using city level data for the same variable and population data for each zip code.

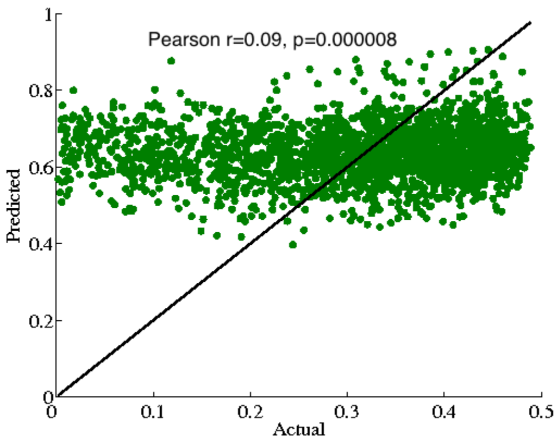

(a) Percentage of Whites (Ratios between 0 and 1)

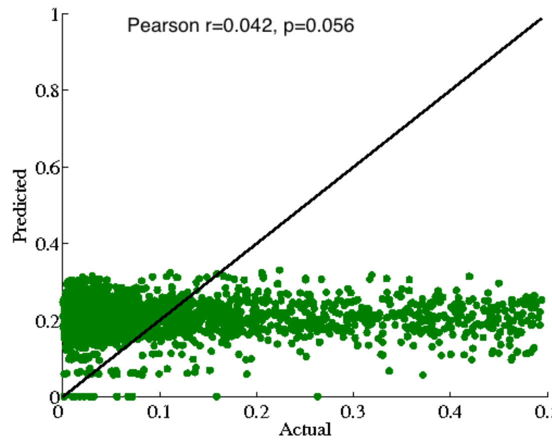

(b) Percentage of Blacks (Ratios between 0 and 1)

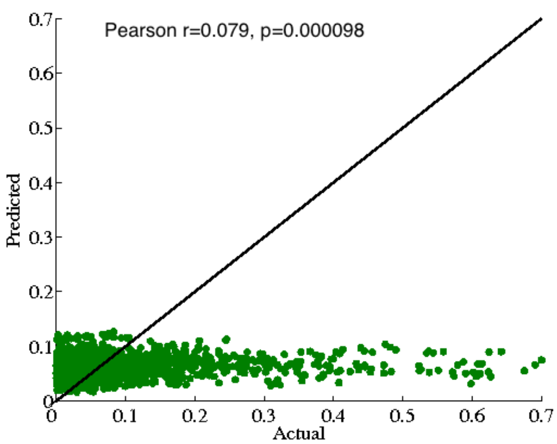

(c) Percentage of Asians (Ratios between 0 and 1)

Fig. S29. Zip code level income, race and education variables from the ACS inferred using city level data for the same variable and population data for each zip code.

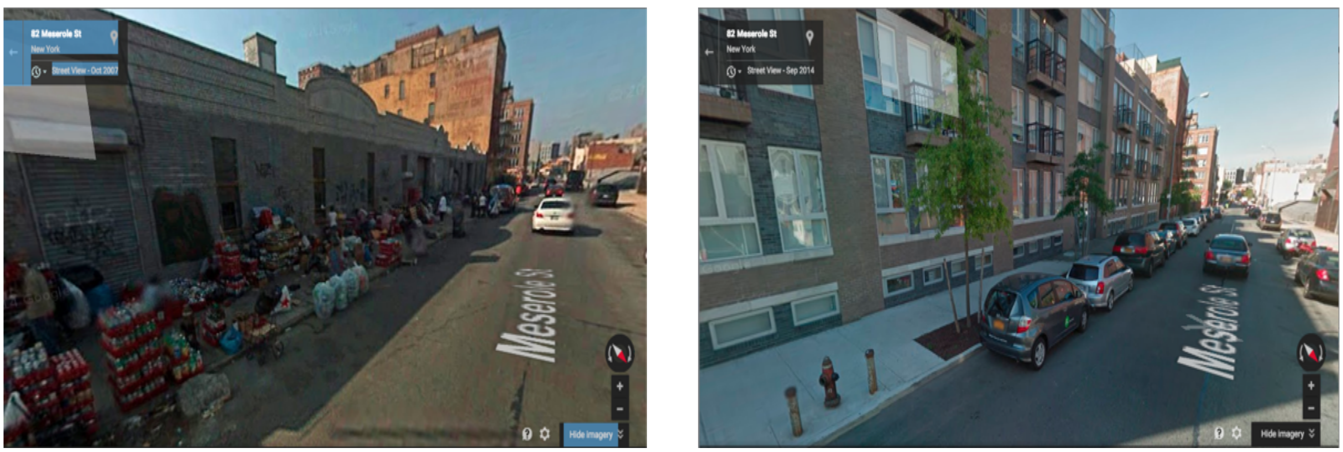

2007

2014

New York City, New York

Fig. S30. Google Street View Timelapse Images of a particular Neighborhood in Brooklyn New York. The economic development of this neighborhood is apparent from its timelapse images in 2007 and 2014 depicting its transformation.

(a) 2007

(b) 2008

(c) 2009

(d) 2010

Fig. S31. Maps showing a random sample of GPS points where Google Street View timelapse images were retrieved for New York city between 2007 and 2014.

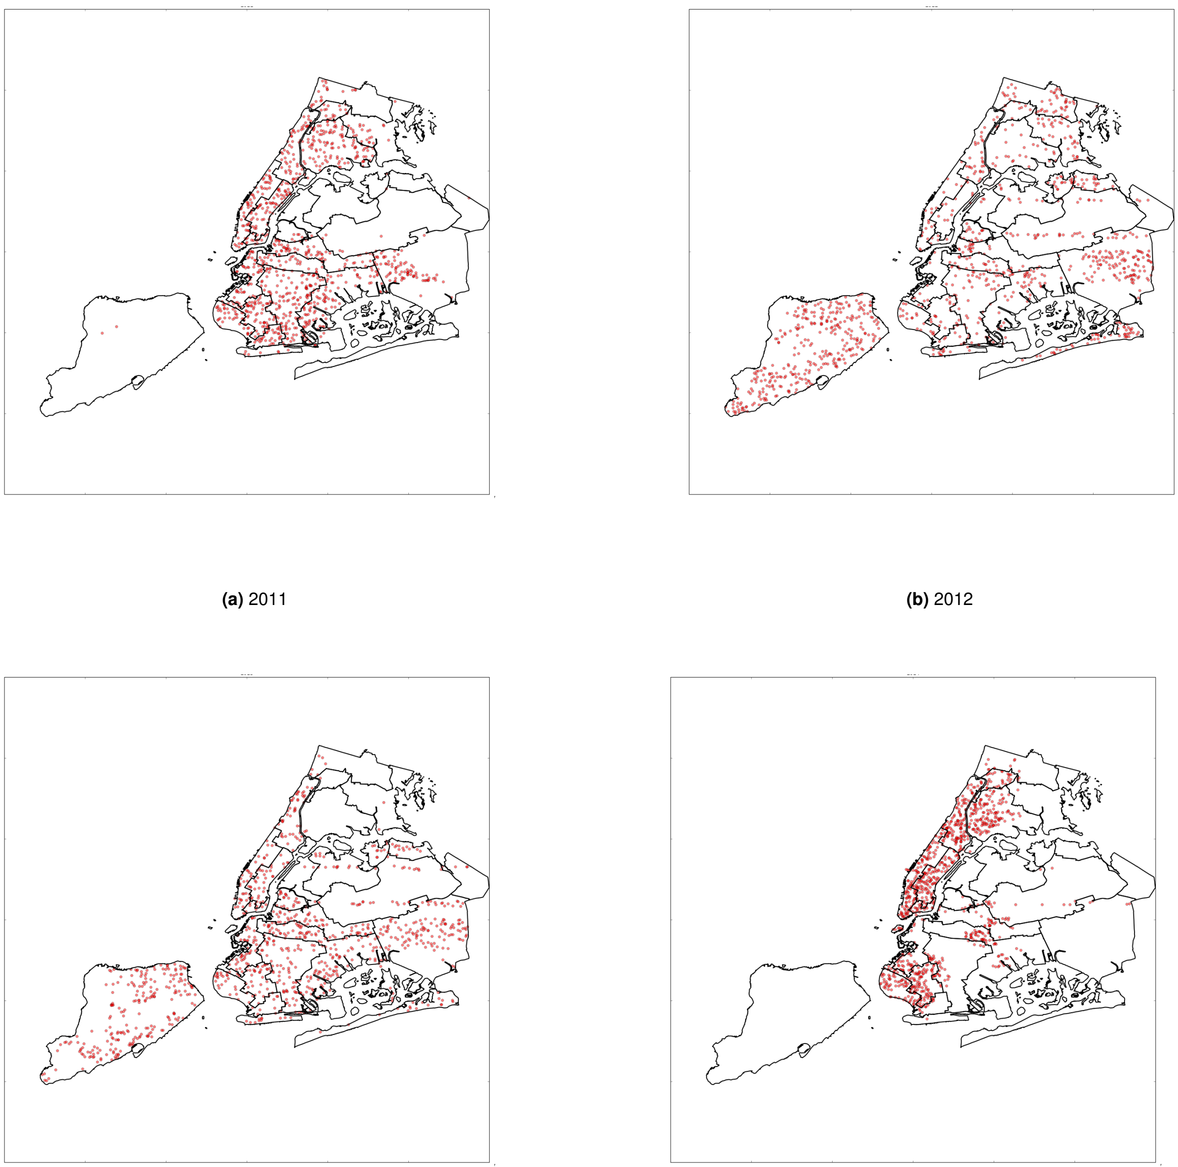

(a) 2011

(b) 2012

(c) 2013

(d) 2014

Fig. S32. Maps showing a random sample of GPS points where Google Street View timelapse images were retrieved for New York city between 2007 and 2014.

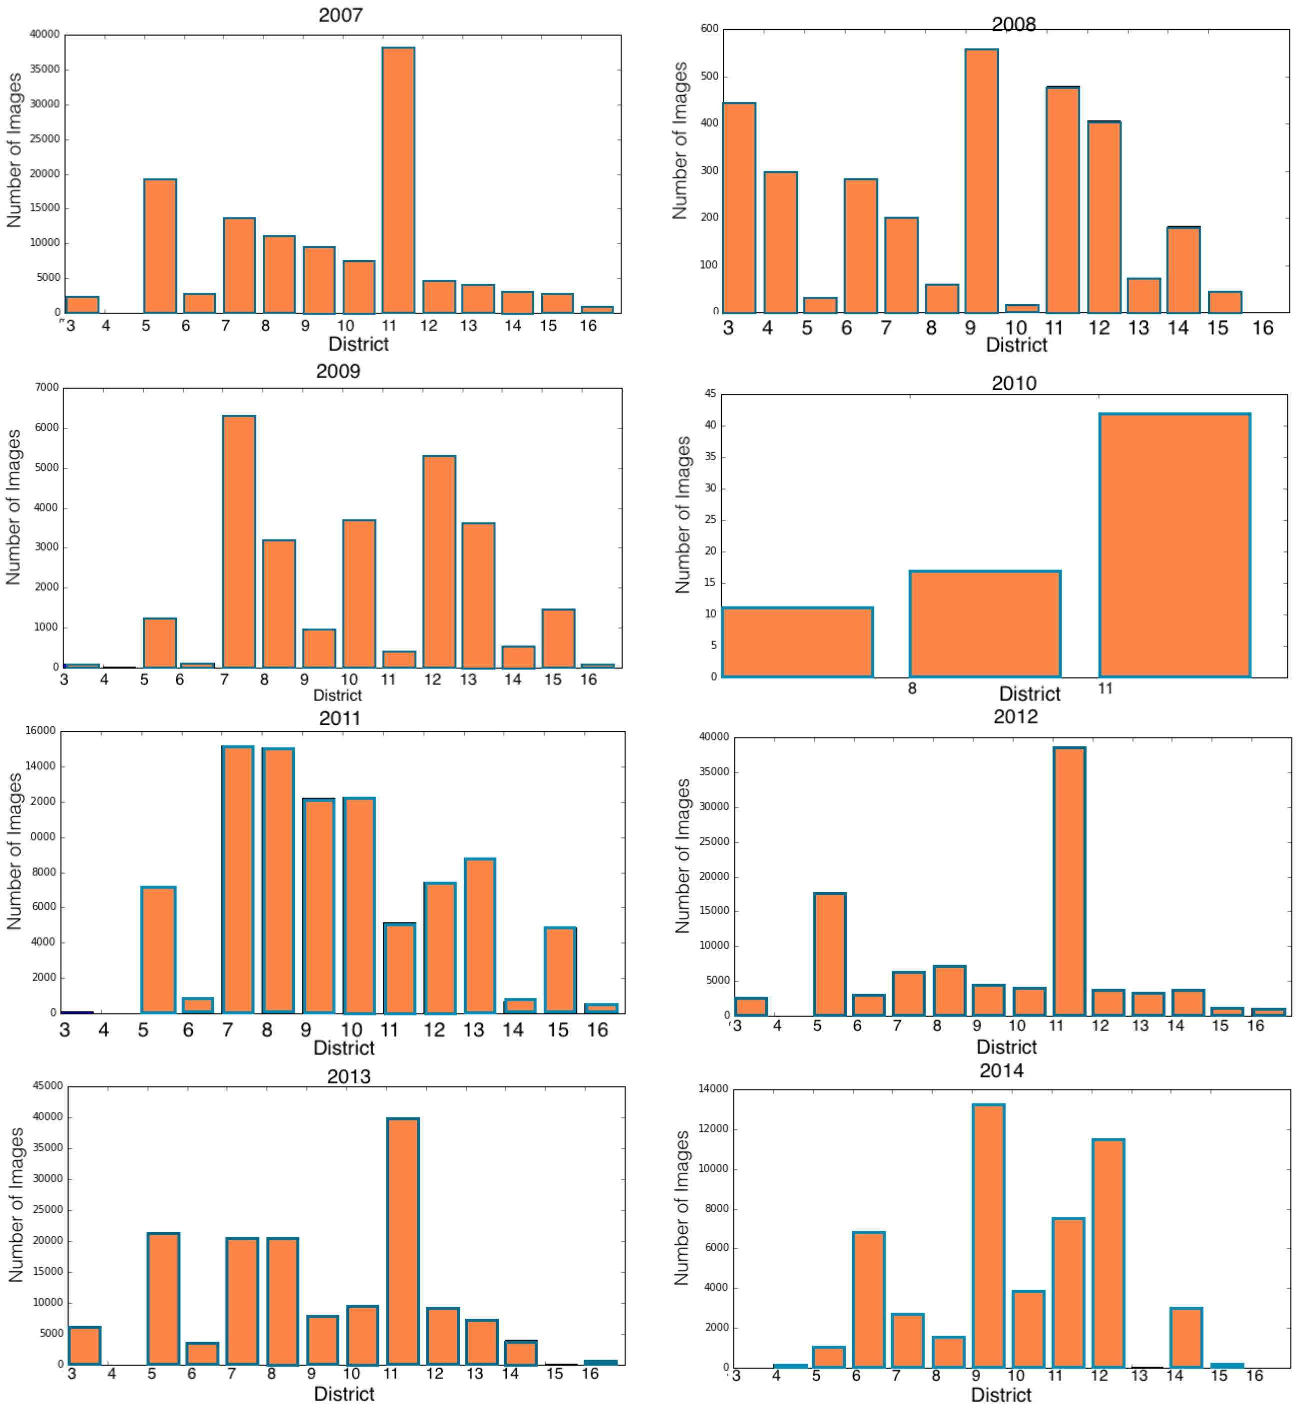

Fig. S33. Bar plots depicting the number of Google Street View timelapse images retrieved for each congressional district in New York city.

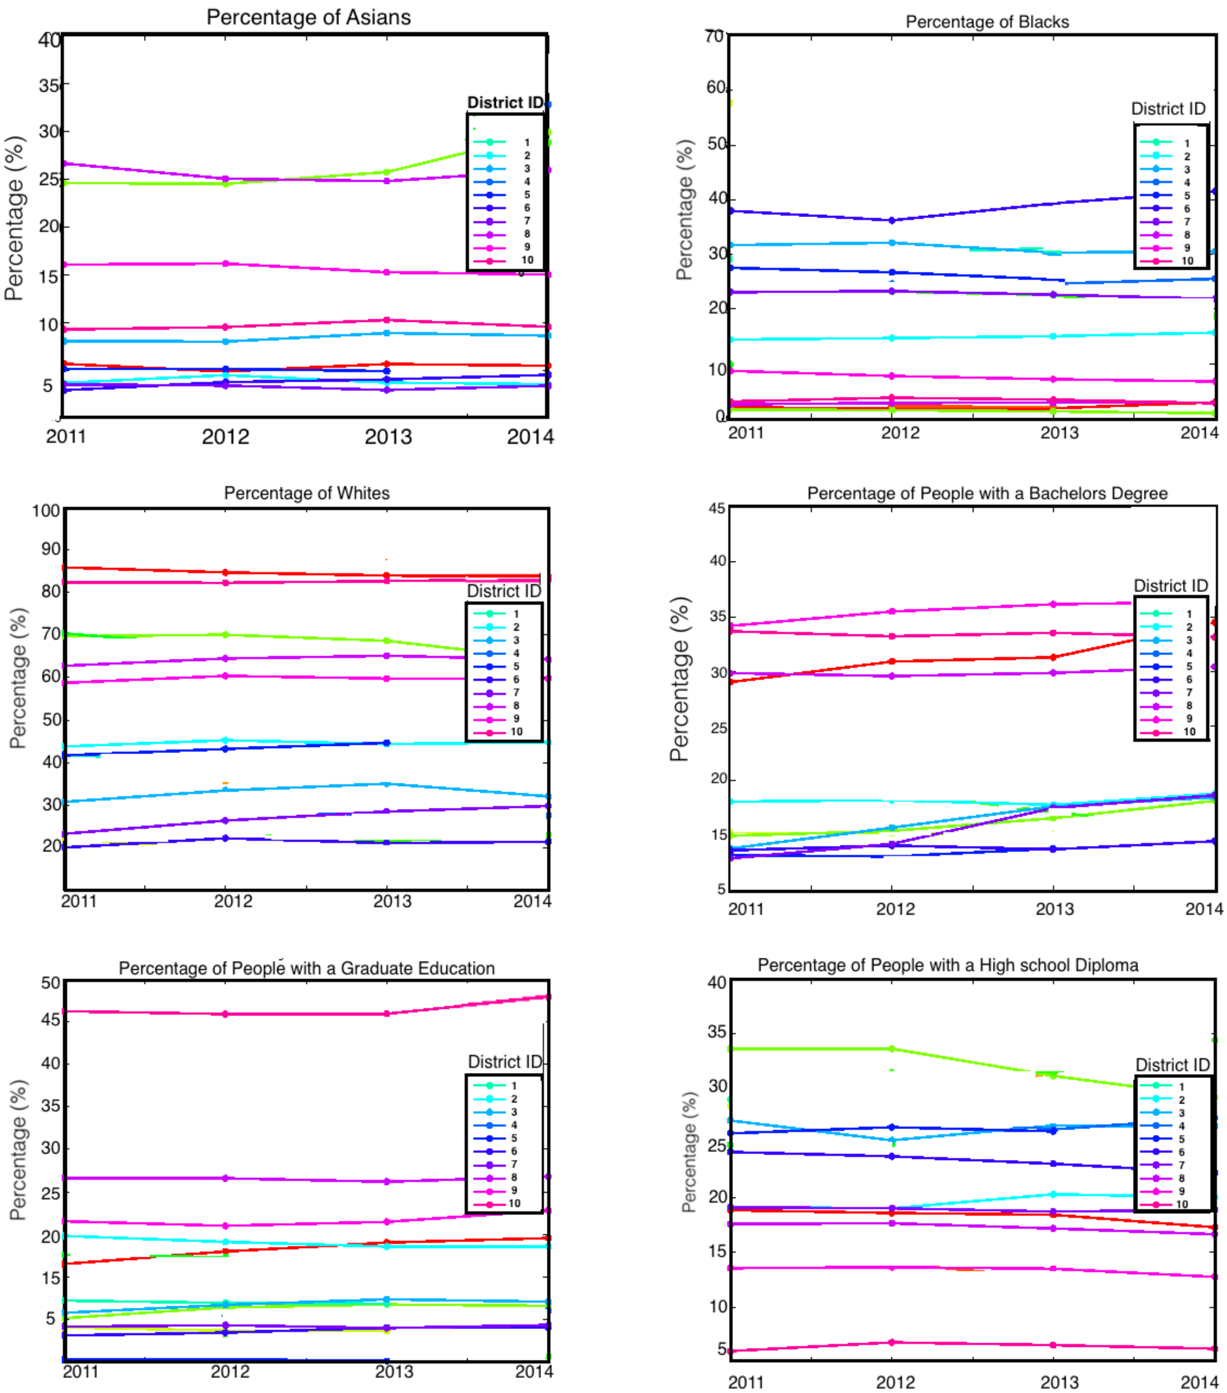

Fig. S34. Plots showing the change in various ACS socioeconomic variables from 2011—2014 for each of the 13 congressional districts in New York city. There is very little change at the district level.

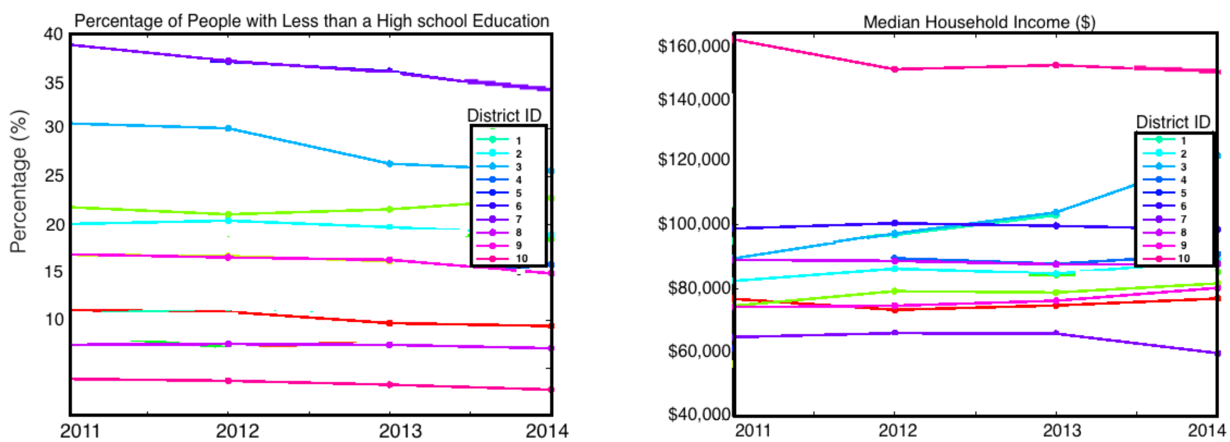

Fig. S35. Plots showing the change in various ACS socioeconomic variables from 2011—2014 for each of the 13 congressional districts in New York city. There is very little change at the district level (continued from prior page).

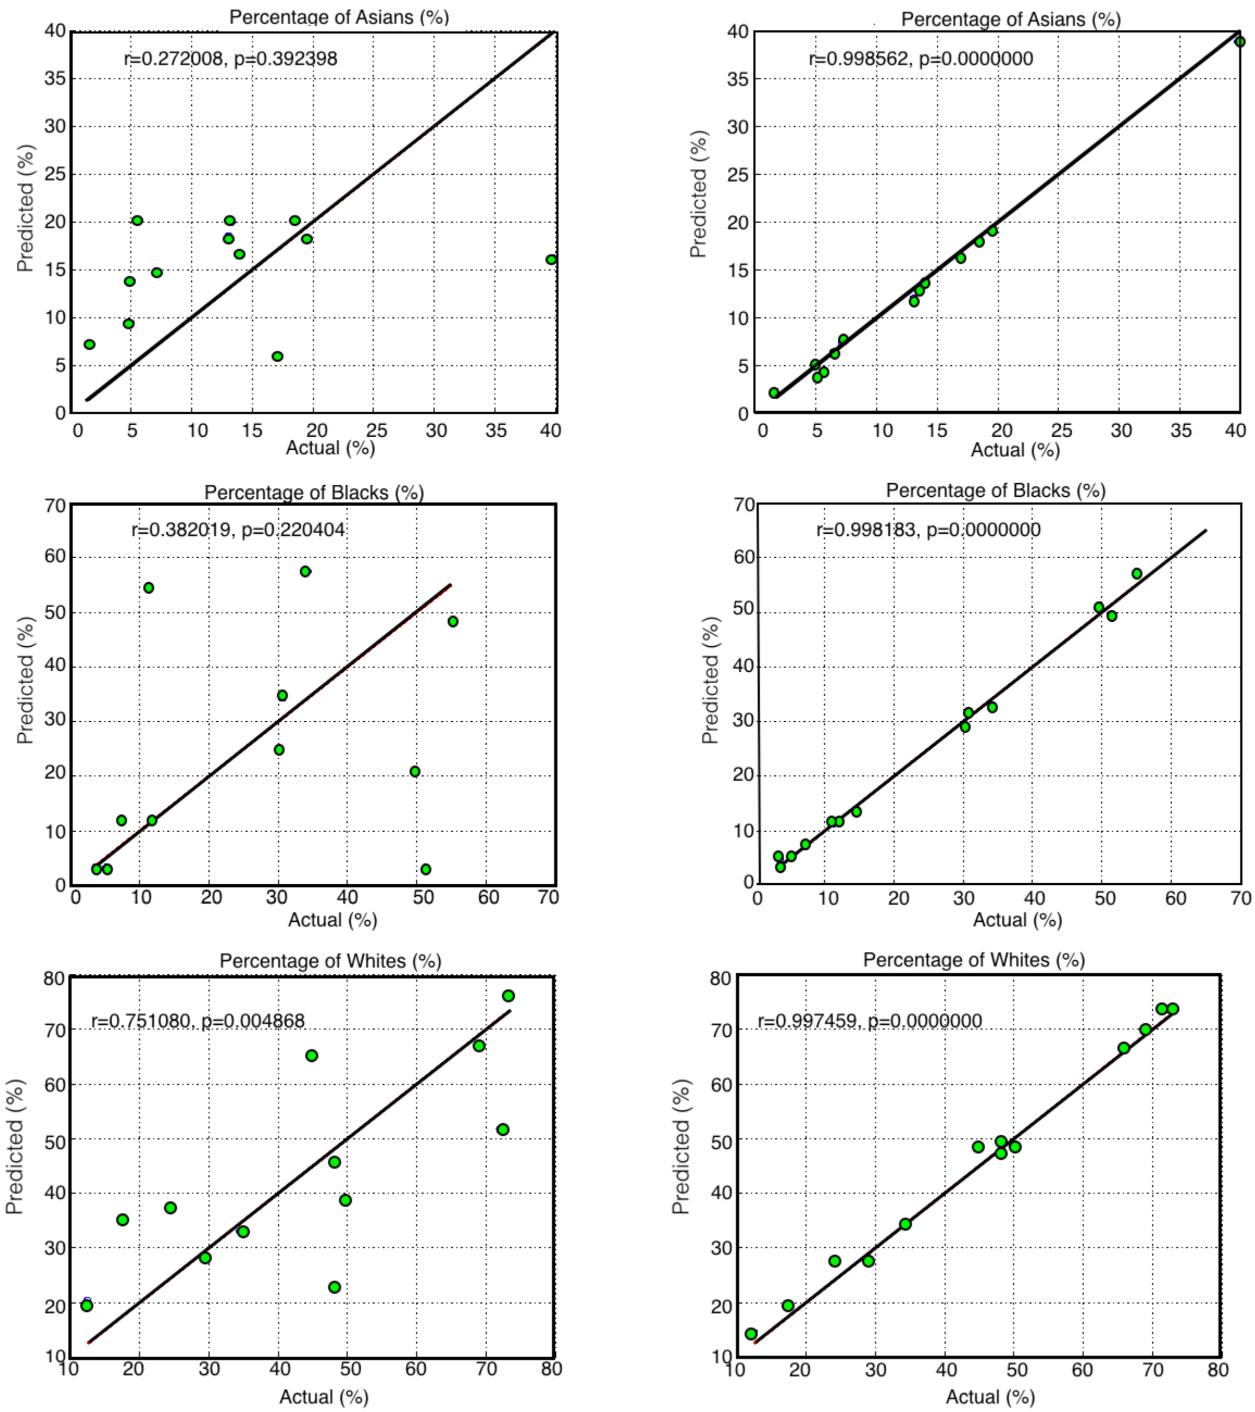

Fig. S36. Results for inferring 2014 demographic variables in New York City's congressional districts using 2013 ACS data for New York city and Google Street View timelapse images for 2014. Left column shows the results applying our methodology. Right column shows a baseline assuming now change in demographic variables. Assuming no change gives better results because we are not specifically training a model to predict changes and there is very little change in NYC ACS data from 2013 to 2014 at the district level.

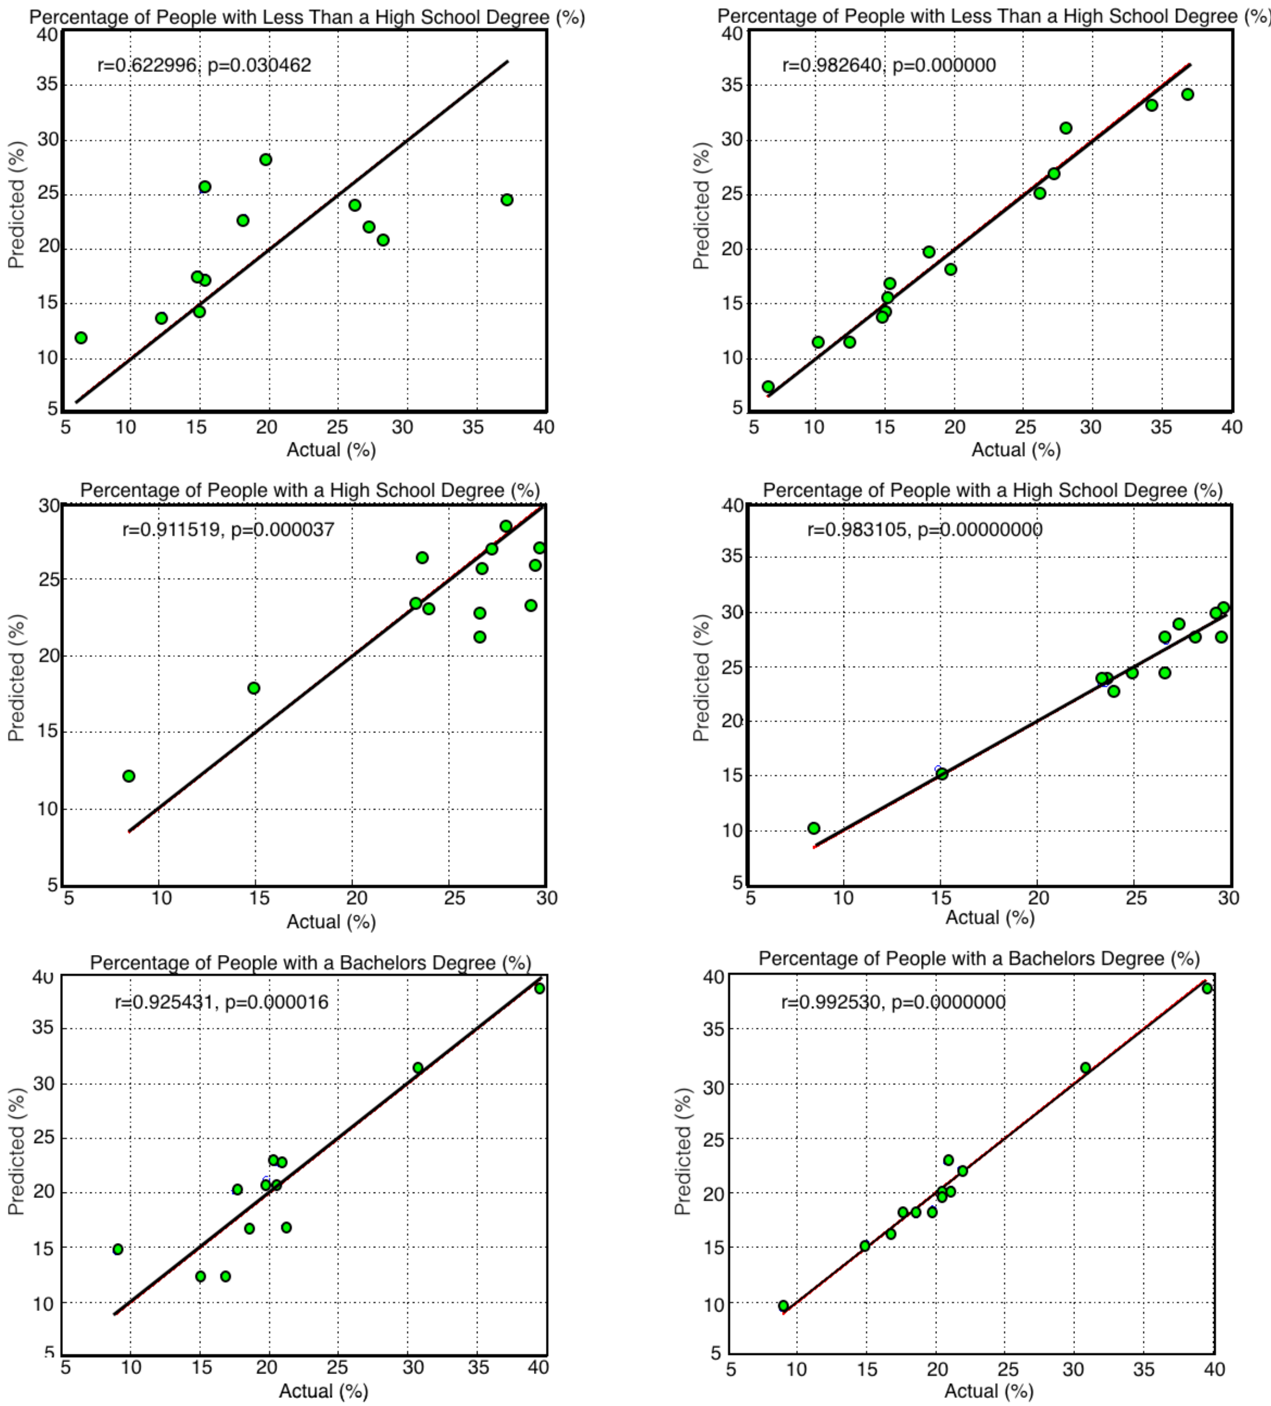

Fig. S37. (Continued from Prior Page) Results for inferring 2014 demographic variables in New York City's congressional districts using 2013 ACS data for New York city and Google Street View timelapse images for 2014. Left column shows the results applying our methodology. Right column shows a baseline assuming now change in demographic variables. Assuming no change gives better results because we are not specifically training a model to predict changes and there is very little change in NYC ACS data from 2013 to 2014 at the district level.

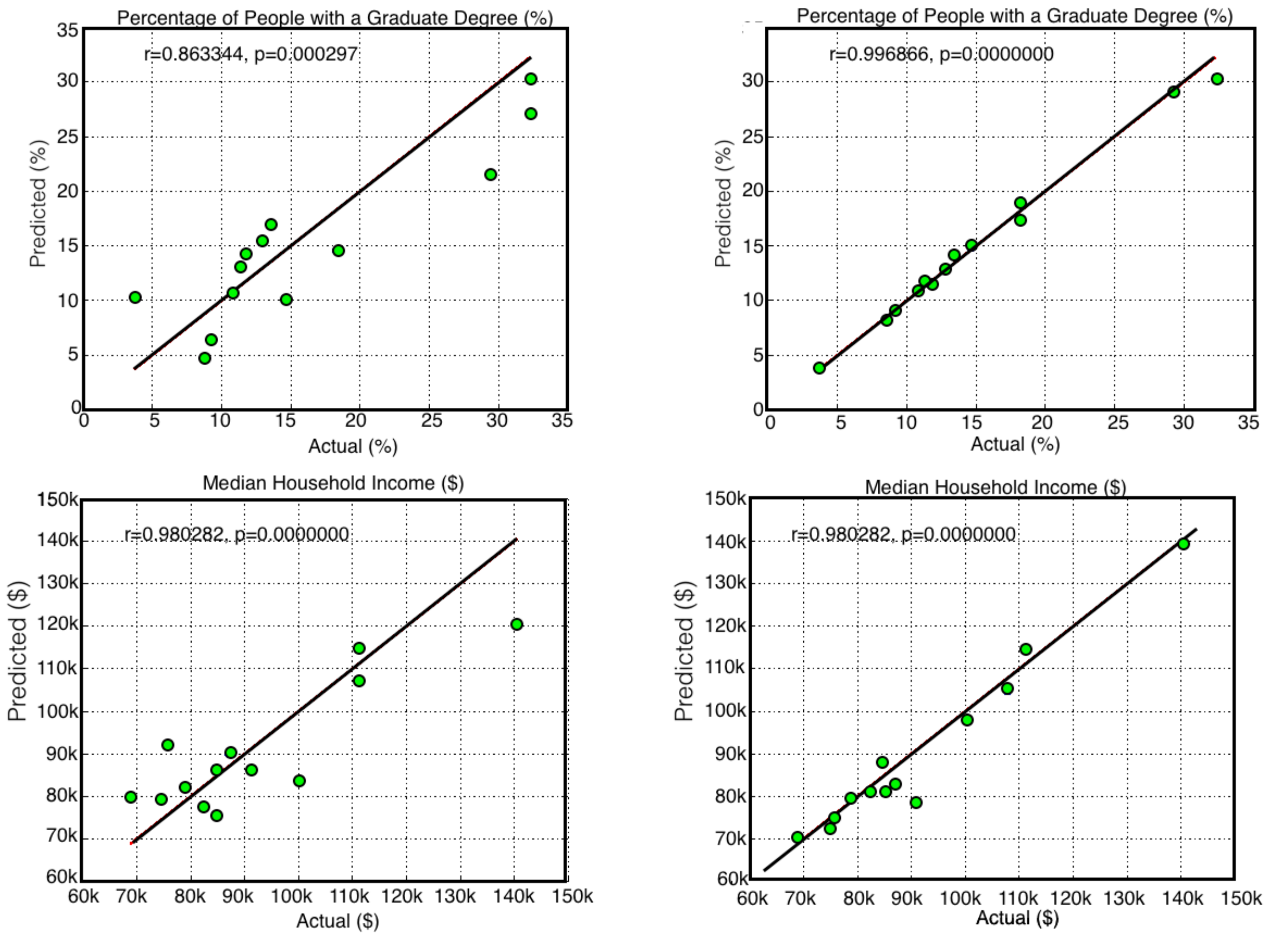

Fig. S38. (Continued from Prior Page) Results for inferring 2014 demographic variables in New York City's congressional districts using 2013 ACS data for New York city and Google Street View timelapse images for 2014. Left column shows the results applying our methodology. Right column shows a baseline assuming now change in demographic variables. Assuming no change gives better results because we are not specifically training a model to predict changes and there is very little change in NYC ACS data from 2013 to 2014 at the district level.

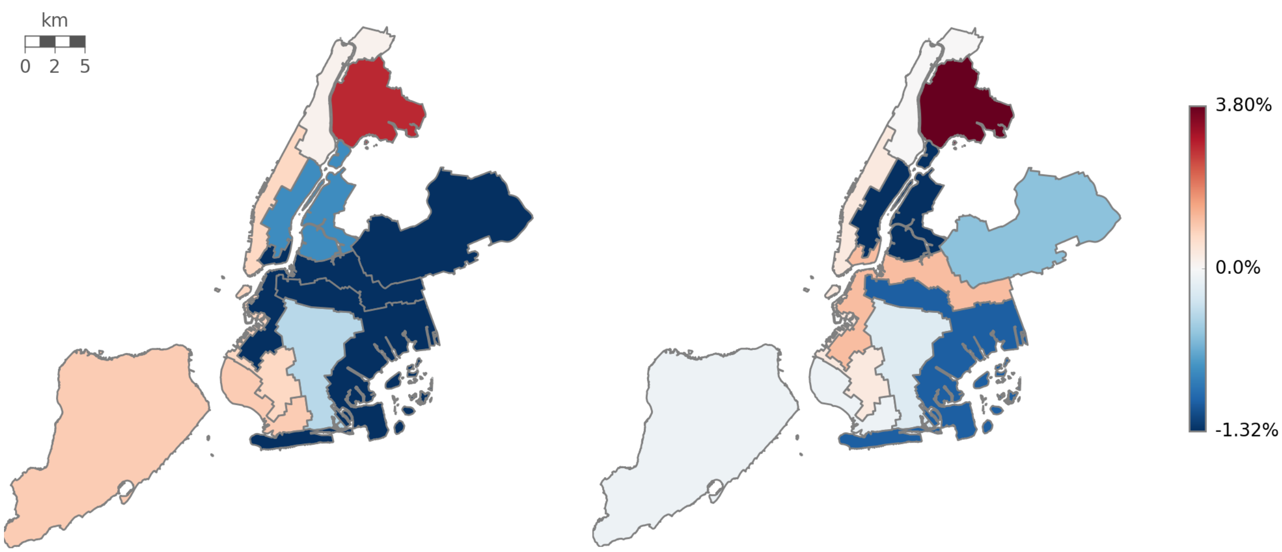

Fig. S39. The change in the percentage of people with less than a high school education in each of the 13 New York city congressional districts between 2013 and 2014. Left is the actual ACS Value and Right is our Prediction. Red signifies a decrease in percentage and blue an increase.

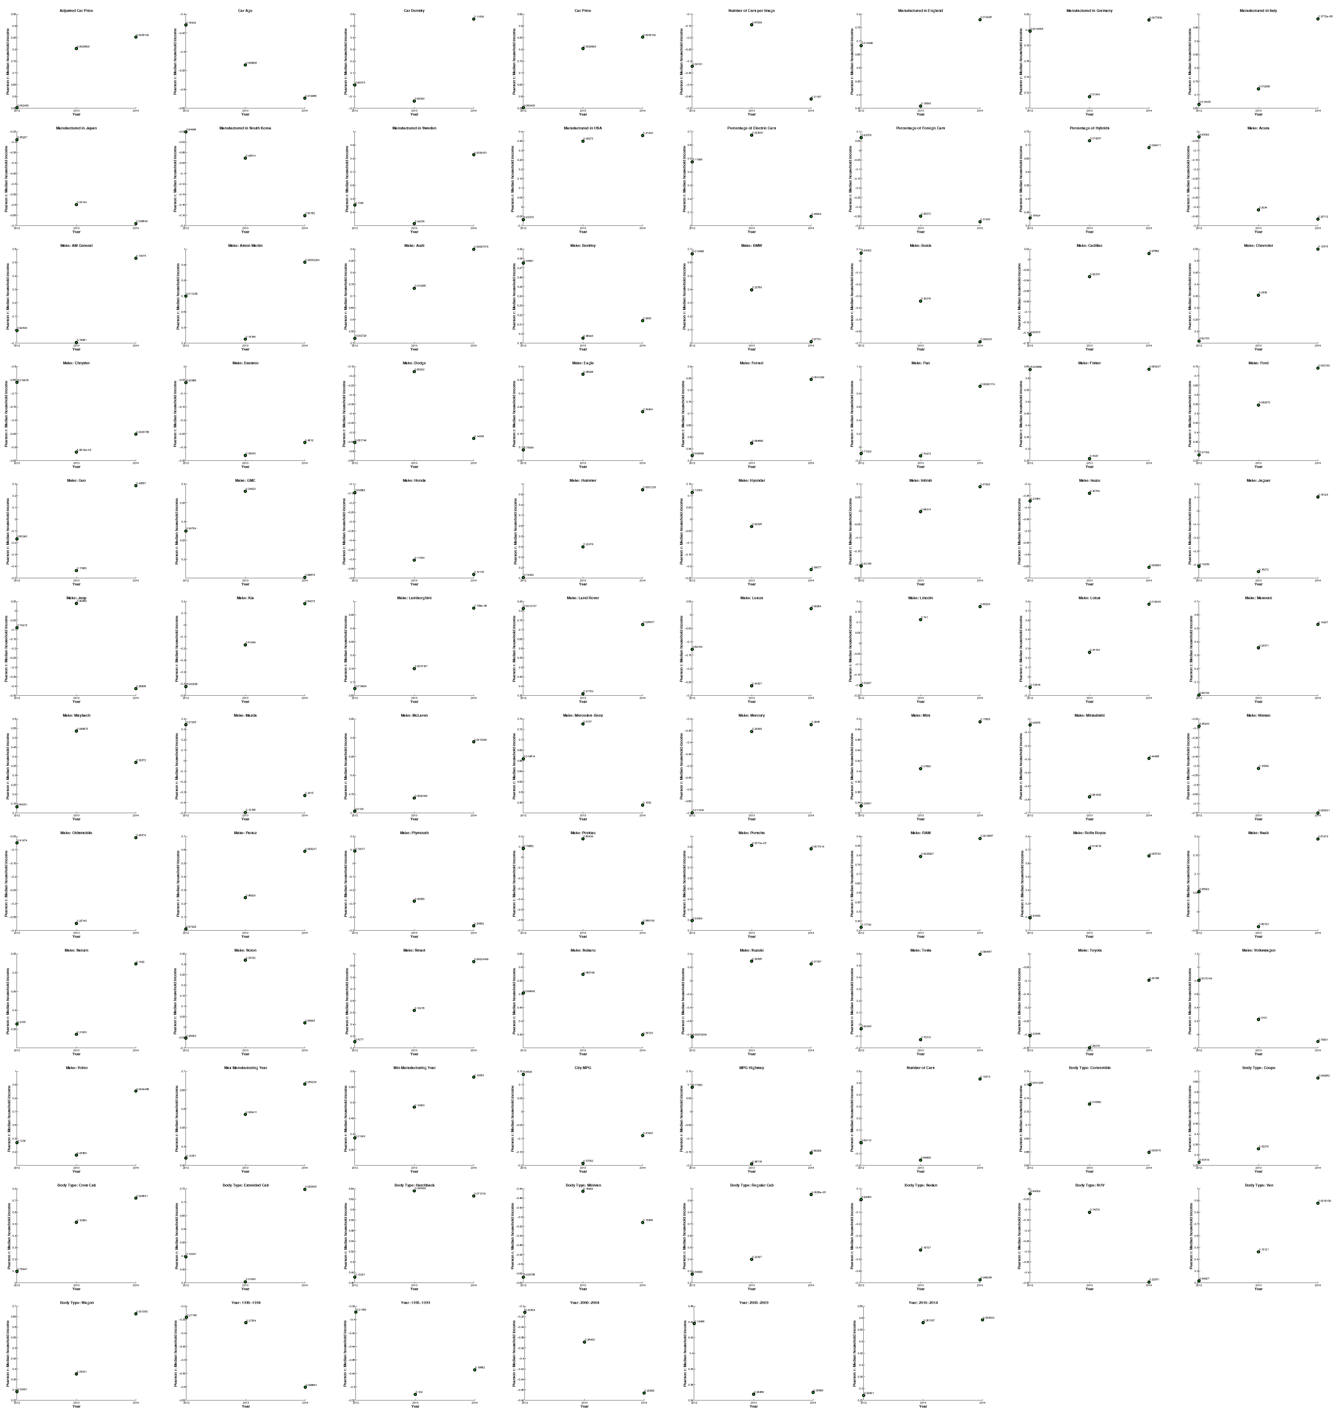

Fig. S40. The Pearson correlation coefficient between various car attributes detected using our method and median household income across time. The p values are also listed next to each point. This tests the stability of the correlations across time. Note that many errors are most probably introduced due to the nonuniform sampling of GPS points across time. Note: zoom in to see a detailed view of each plot.

City	# Im.	City	# Im.	City	# Im.	City	# Im.
Birmingham, AL	484,818	Santa Ana, CA	90,030	Portland, ME	86,874	Salem, OR	102,174
Huntsville, AL	100,410	Santa Clarita, CA	83,298	Baltimore, MD	570,360	Philadelphia, PA	244,194
Mobile, AL	45,114	Santa Rosa, CA	243,324	Frederick, MD	182,388	Pittsburgh, PA	682,728
Montgomery, AL	45,084	Stockton, CA	343,662	Boston, MA	195,864	Providence, RI	130,104
Anchorage, AK	59,484	Sunnyvale, CA	66,318	Springfield, MA	116,928	Warwick, RI	172,092
Fairbanks, AK	42,384	Torrance, CA	136,260	Worcester, MA	197,424	Charleston, SC	56,604
Chandler, AZ	309,414	Aurora, CO	143,508	Detroit, MI	287,736	Columbia, SC	334,914
Gilbert, AZ	175,242	Colorado Springs, CO	492,222	Grand Rapids, MI	202,266	Rapid City, SD	30,954
Glendale, AZ	160,146	Denver, CO	306,990	Minneapolis, MN	654,270	Sioux Falls, SD	74,640
Mesa, AZ	283,620	Fort Collins, CO	307,056	Saint Paul, MN	164,034	Chattanooga, TN	284,214
Peoria, AZ	135,132	Bridgeport, CT	154,092	Gulfport, MS	14,898	Knoxville, TN	457,434
Phoenix, AZ	623,892	New Haven, CT	62,394	Jackson, MS	71,298	Memphis, TN	97,572
Scottsdale, AZ	138,120	Dover, DE	22,134	Kansas City, MO	577,830	Nashville, TN	554,118
Tempe, AZ	302,958	Wilmington, DE	80,754	Springfield, MO	395,502	Amarillo, TX	85,380
Tucson, AZ	634,986	Washington, DC	375,258	St. Louis, MO	426,942	Arlington, TX	509,406
Fort Smith, AR	205,512	Cape Coral, FL	309,102	Billings, MT	54,768	Austin, TX	211,530
Little Rock, AR	398,094	Fort Lauderdale, FL	279,300	Missoula, MT	157,254	Brownsville, TX	284,826
Anaheim, CA	133,098	Hialeah, FL	143,928	Lincoln, NE	444,306	Corpus Christi, TX	61,434
Bakersfield, CA	521,112	Jacksonville, FL	770,016	Omaha, NE	322,602	Dallas, TX	663,006
Chula Vista, CA	189,204	Miami, FL	310,692	Henderson, NV	259,416	El Paso, TX	205,500
Corona, CA	238,932	Orlando, FL	582,018	Las Vegas, NV	521,172	Fort Worth, TX	677,214
Elk Grove, CA	306,600	Pembroke Pines, FL	71,274	North Las Vegas, NV	197,394	Garland, TX	226,140
Escondido, CA	206,550	Port St. Lucie, FL	62,292	Reno, NV	104,328	Grand Prairie, TX	210,198
Fontana, CA	167,604	Saint Petersburg, FL	83,442	Manchester, NH	131,682	Houston, TX	337,830
Fremont, CA	232,608	Tallahassee, FL	419,220	Nashua, NH	139,890	Irving, TX	179,382
Fresno, CA	135,210	Tampa, FL	610,770	Jersey City, NJ	78,036	Laredo, TX	259,878
Garden Grove, CA	77,706	Atlanta, GA	315,336	Newark, NJ	129,948	Lubbock, TX	500,760
Glendale, CA	77,316	Augusta, GA	239,994	Albuquerque, NM	73,746	Pasadena, TX	29,700
Hayward, CA	207,744	Columbus, GA	54,246	Las Cruces, NM	82,098	Plano, TX	330,186
Huntington Beach, CA	101,574	Hilo, HI	14,406	Buffalo, NY	376,806	San Antonio, TX	1,034,358
Irvine, CA	183,474	Honolulu, HI	209,010	New York, NY	508,860	Salt Lake City, UT	272,190
Lancaster, CA	110,550	Boise, ID	42,438	Rochester, NY	391,458	West Valley City, UT	69,432
Long Beach, CA	265,806	Nampa, ID	231,318	Yonkers, NY	27,618	Burlington, VT	31,998
Los Angeles, CA	554,106	Aurora, IL	203,256	Charlotte, NC	111,510	Essex, VT	16,056
Modesto, CA	32,406	Chicago, IL	791,298	Durham, NC	359,592	Alexandria, VA	69,924
Moreno Valley, CA	180,516	Joliet, IL	118,116	Fayetteville, NC	292,296	Chesapeake, VA	38,568
Oakland, CA	326,208	Rockford, IL	372,156	Greensboro, NC	80,730	Newport News, VA	17,862
Oceanside, CA	129,384	Fort Wayne, IN	99,672	Raleigh, NC	409,776	Norfolk, VA	56,688
Ontario, CA	142,230	Indianapolis, IN	468,780	Winston-Salem, NC	457,314	Richmond, VA	504,138
Oxnard, CA	154,074	Cedar Rapids, IA	257,178	Bismarck, ND	156,912	Virginia Beach, VA	40,698
Palmdale, CA	164,064	Des Moines, IA	123,678	Fargo, ND	202,422	Seattle, WA	529,392
Pomona, CA	153,798	Kansas City, KS	577,830	Akron, OH	404,376	Spokane, WA	381,684
Rancho Cucamonga, CA	88,734	Overland Park, KS	9,252	Cincinnati, OH	511,842	Tacoma, WA	331,338
Riverside, CA	446,412	Wichita, KS	569,658	Cleveland, OH	416,142	Vancouver, WA	292,560
Sacramento, CA	525,756	Lexington, KY	345,516	Columbus, OH	568,776	Charleston, WV	38,628
Salinas, CA	175,530	Louisville, KY	419,544	Toledo, OH	51,444	Huntington, WV	42,144
San Bernardino, CA	124,002	Baton Rouge, LA	65,592	Oklahoma City, OK	687,234	Madison, WI	218,580
San Diego, CA	472,872	New Orleans, LA	456,042	Tulsa, OK	541,458	Milwaukee, WI	446,172
San Francisco, CA	215,298	Shreveport, LA	100,662	Eugene, OR	108,582	Casper, WY	43,542
San Jose, CA	274,848	Lewiston, ME	50,562	Portland, OR	548,334	Cheyenne, WY	211,668

Table S1. List of cities we collected Street View images for and the number of Street View images we collected for each city.

Attribute	Training	Validation	Test
Street View Images	199,666	39,933	159,732
Product Shot Images	313,099	-	-
Total Images	512,765	39,933	159,732
Street View Bounding Boxes	272,142	54,691	216,808
Product Shot Bounding Boxes	313,099	-	-
Total Bounding Boxes	585,241	54,691	216,808
Street View Category Labels	34,753	6,921	27,888
Product Shot Category Labels	313,099	-	-
Total Category Labels	347,852	6,921	27,888

Table S2. Dataset statistics for our training, validation, and test splits, separated into Street View and product shot images.

Comp.	Parts	AP	Time
1	0	52.3	2.27
1	4	63.2	3.48
1	8	64.2	4.84
3	0	62.9	6.48
3	4	66.7	12.20
3	8	68.4	16.47
5	0	64.8	10.25
5	4	67.3	16.33
5	8	68.7	22.07
6	0	65.2	10.48
8	0	66.0	11.17

Table S3. Average Precision (AP) on the Street View validation set for various DPM configurations. Time is measured in seconds per image. Comp. is the number of DPM components, and Parts indicates the number of parts in the model.

Variable	Pearson's r	p-value	Variable	Pearson's r	p-value
Price	0.4435	1.16e-117	Make: Fiat	0.0747	0.000226
Cars/Image	0.0235	0.247	Make: Fisker	0.0751	0.000213
MPG Highway	0.1642	3.72e-16	Make: Ford	-0.2697	9.17e-42
MPG City	0.2565	8.08e-38	Make: Geo	-0.2051	1.73e-24
Hybrid	0.1169	7.58e-09	Make: GMC	-0.1627	7.08e-16
Electric	0.1589	3.35e-15	Make: Honda	0.4234	2.67e-106
Foreign	0.4672	5.26e-132	Make: Hummer	0.0799	7.99e-05
Country: England	0.3126	3.15e-56	Make: Hyundai	0.1201	2.91e-09
Country: Germany	0.4335	6.3e-112	Make: Infiniti	0.2649	2.58e-40
Country: Italy	0.2167	3.18e-27	Make: Isuzu	-0.1252	5.91e-10
Country: Japan	0.4471	1.01e-119	Make: Jaguar	-0.0397	0.0504
Country: South Korea	0.0834	3.83e-05	Make: Jeep	0.0153	0.452
Country: Sweden	0.2553	1.8e-37	Make: Kia	-0.0239	0.238
Country: USA	-0.4672	5.26e-132	Make: Lamborghini	0.1999	2.56e-23
Body Type: Convertible	0.1484	1.95e-13	Make: Land Rover	0.3000	1e-51
Body Type: Coupe	-0.2211	2.69e-28	Make: Lexus	0.4432	1.73e-117
Body Type: Crew Cab	-0.0002	0.993	Make: Lincoln	-0.1652	2.54e-16
Body Type: Extended Cab	-0.0943	3.19e-06	Make: Lotus	0.1232	1.11e-09
Body Type: Hatchback	0.3352	7.16e-65	Make: Maserati	0.1096	6.05e-08
Body Type: Minivan	0.0833	3.94e-05	Make: Maybach	0.0570	0.00494
Body Type: Regular Cab	-0.2179	1.7e-27	Make: Mazda	0.2094	1.76e-25
Body Type: Sedan	-0.1537	2.62e-14	Make: Mclaren	0.1002	7.54e-07
Body Type: SUV	0.3136	1.34e-56	Make: Mercedes-Benz	0.3873	8.94e-88
Body Type: Van	-0.0391	0.0542	Make: Mercury	-0.3367	1.71e-65
Body Type: Wagon	0.1776	1.14e-18	Make: Mini	0.2749	2.21e-43
Year: 1990-1994	-0.3230	4.21e-60	Make: Mitsubishi	-0.0739	0.000269
Year: 1995-1999	-0.4202	1.5e-104	Make: Nissan	0.0894	1.01e-05
Year: 2000-2004	0.2043	2.56e-24	Make: Oldsmobile	-0.3964	3.12e-92
Year: 2005-2009	0.3864	2.38e-87	Make: Panoz	0.0507	0.0124
Year: 2010-2014	0.3694	2.01e-79	Make: Plymouth	-0.2496	7.85e-36
Make: Acura	0.3528	3.78e-72	Make: Pontiac	-0.3805	1.46e-84
Make: AM General	0.0008	0.969	Make: Porsche	0.2967	1.45e-50
Make: Aston Martin	0.0934	4e-06	Make: Ram	-0.0513	0.0114
Make: Audi	0.3420	1.19e-67	Make: Rolls-Royce	0.0843	3.18e-05
Make: Bentley	-0.0319	0.115	Make: Saab	0.2215	2.14e-28
Make: BMW	0.3939	5.53e-91	Make: Saturn	-0.0793	9.11e-05
Make: Buick	-0.3975	8.43e-93	Make: Scion	0.2463	6.47e-35
Make: Cadillac	-0.3248	8.39e-61	Make: Smart	0.1464	4.15e-13
Make: Chevrolet	-0.3553	3.08e-73	Make: Subaru	0.1727	1.03e-17
Make: Chrysler	-0.2720	1.81e-42	Make: Suzuki	-0.0679	0.000817
Make: Daewoo	-0.0214	0.293	Make: Tesla	0.0860	2.19e-05
Make: Dodge	-0.3807	1.22e-84	Make: Toyota	0.4239	1.43e-106
Make: Eagle	-0.2009	1.55e-23	Make: Volkswagen	0.3014	3.24e-52
Make: Ferrari	0.0694	0.000619	Make: Volvo	0.2398	3.9e-33

Table S4. Pearson r correlation coefficients and their associated p-values for each car attribute between median household income and each car attribute, at the zip code level. p-values are with respect to the null hypothesis of no correlation.

Variable	Pearson's r	p-value	Variable	Pearson's r	p-value
Price	-0.3333	4e-64	Make: Fiat	-0.0721	0.000372
Cars/Image	0.1246	7.06e-10	Make: Fisker	-0.1153	1.19e-08
MPG Highway	-0.2920	5.69e-49	Make: Ford	0.2991	2.1e-51
MPG City	-0.3182	2.57e-58	Make: Geo	0.1532	3.18e-14
Hybrid	-0.1331	4.43e-11	Make: GMC	0.1745	4.48e-18
Electric	-0.1508	7.82e-14	Make: Honda	-0.2530	8.23e-37
Foreign	-0.2955	3.64e-50	Make: Hummer	0.0199	0.328
Country: England	-0.2434	4.13e-34	Make: Hyundai	-0.3468	1.26e-69
Country: Germany	-0.3029	1.01e-52	Make: Infiniti	-0.1423	1.8e-12
Country: Italy	-0.1587	3.6e-15	Make: Isuzu	0.1501	1.03e-13
Country: Japan	-0.2301	1.45e-30	Make: Jaguar	-0.0182	0.37
Country: South Korea	-0.3456	4.03e-69	Make: Jeep	-0.2030	5.22e-24
Country: Sweden	-0.3541	1.04e-72	Make: Kia	-0.1942	4.4e-22
Country: USA	0.2955	3.64e-50	Make: Lamborghini	-0.1403	3.68e-12
Body Type: Convertible	-0.2246	3.67e-29	Make: Land Rover	-0.1698	3.55e-17
Body Type: Coupe	0.0892	1.06e-05	Make: Lexus	-0.2421	9.4e-34
Body Type: Crew Cab	0.1080	9.47e-08	Make: Lincoln	0.1483	2.05e-13
Body Type: Extended Cab	0.2204	3.98e-28	Make: Lotus	-0.0684	0.000745
Body Type: Hatchback	-0.3445	1.13e-68	Make: Maserati	-0.0745	0.000239
Body Type: Minivan	-0.0206	0.31	Make: Maybach	-0.0537	0.00806
Body Type: Regular Cab	0.2732	7.55e-43	Make: Mazda	-0.2595	1.11e-38
Body Type: Sedan	-0.0173	0.395	Make: Mclaren	-0.0764	0.000164
Body Type: SUV	-0.2361	3.88e-32	Make: Mercedes-Benz	-0.1895	4.47e-21
Body Type: Van	0.2208	3.32e-28	Make: Mercury	0.0984	1.16e-06
Body Type: Wagon	-0.3888	1.55e-88	Make: Mini	-0.2471	3.94e-35
Year: 1990-1994	0.3230	3.87e-60	Make: Mitsubishi	0.1061	1.58e-07
Year: 1995-1999	0.3715	2.12e-80	Make: Nissan	0.1217	1.78e-09
Year: 2000-2004	-0.1652	2.5e-16	Make: Oldsmobile	0.1131	2.24e-08
Year: 2005-2009	-0.3831	8.51e-86	Make: Panoz	-0.1770	1.47e-18
Year: 2010-2014	-0.3296	1.17e-62	Make: Plymouth	0.1237	9.6e-10
Make: Acura	-0.1895	4.45e-21	Make: Pontiac	0.0520	0.0103
Make: AM General	0.0676	0.00085	Make: Porsche	-0.2387	8.08e-33
Make: Aston Martin	-0.1157	1.07e-08	Make: Ram	-0.0817	5.54e-05
Make: Audi	-0.3176	4.47e-58	Make: Rolls-Royce	-0.1050	2.12e-07
Make: Bentley	0.0356	0.0792	Make: Saab	-0.2844	1.93e-46
Make: BMW	-0.1956	2.24e-22	Make: Saturn	-0.1775	1.19e-18
Make: Buick	0.0353	0.0822	Make: Scion	-0.1481	2.14e-13
Make: Cadillac	0.1866	1.81e-20	Make: Smart	-0.1571	6.68e-15
Make: Chevrolet	0.3183	2.46e-58	Make: Subaru	-0.3597	3.82e-75
Make: Chrysler	0.0264	0.194	Make: Suzuki	-0.0054	0.79
Make: Daewoo	0.0122	0.548	Make: Tesla	-0.0661	0.00112
Make: Dodge	0.2208	3.24e-28	Make: Toyota	-0.1686	6.02e-17
Make: Eagle	0.0612	0.00254	Make: Volkswagen	-0.2975	7.55e-51
Make: Ferrari	-0.0584	0.00396	Make: Volvo	-0.3376	7.58e-66

Table S5. Pearson r correlation coefficients and their associated p-values for each car attribute between the percentage of residents who did not graduate high school and each car attribute, at the zip code level.

Variable	Pearson's r	p-value	Variable	Pearson's r	p-value
Price	-0.4345	1.91e-112	Make: Fiat	-0.0925	4.9e-06
Cars/Image	-0.2684	2.34e-41	Make: Fisker	-0.1364	1.45e-11
MPG Highway	-0.3294	1.33e-62	Make: Ford	0.3735	2.78e-81
MPG City	-0.4373	4.66e-114	Make: Geo	0.1799	4.06e-19
Hybrid	-0.1288	1.88e-10	Make: GMC	0.2760	1.01e-43
Electric	-0.2963	1.92e-50	Make: Honda	-0.5371	1.12e-181
Foreign	-0.6048	2.04e-242	Make: Hummer	-0.0502	0.0134
Country: England	-0.4811	5.74e-141	Make: Hyundai	-0.0521	0.0102
Country: Germany	-0.6142	5.35e-252	Make: Infiniti	-0.3755	3.16e-82
Country: Italy	-0.2462	6.92e-35	Make: Isuzu	0.0364	0.0731
Country: Japan	-0.5670	9.84e-207	Make: Jaguar	-0.0201	0.322
Country: South Korea	-0.0235	0.247	Make: Jeep	-0.0247	0.223
Country: Sweden	-0.4033	9.8e-96	Make: Kia	0.0442	0.0292
Country: USA	0.6048	2.04e-242	Make: Lamborghini	-0.1902	3.14e-21
Body Type: Convertible	-0.2258	1.84e-29	Make: Land Rover	-0.4198	2.38e-104
Body Type: Coupe	0.1453	6.19e-13	Make: Lexus	-0.4841	5.05e-143
Body Type: Crew Cab	0.1141	1.69e-08	Make: Lincoln	0.2109	7.74e-26
Body Type: Extended Cab	0.2023	7.48e-24	Make: Lotus	-0.0922	5.33e-06
Body Type: Hatchback	-0.5576	1.38e-198	Make: Maserati	-0.1612	1.3e-15
Body Type: Minivan	0.0811	6.31e-05	Make: Maybach	-0.0328	0.106
Body Type: Regular Cab	0.3076	2.01e-54	Make: Mazda	-0.3406	4.81e-67
Body Type: Sedan	0.0382	0.0601	Make: Mclaren	-0.1054	1.89e-07
Body Type: SUV	-0.2750	1.97e-43	Make: Mercedes-Benz	-0.4409	3.81e-116
Body Type: Van	0.0753	0.000204	Make: Mercury	0.3778	2.66e-83
Body Type: Wagon	-0.3653	1.41e-77	Make: Mini	-0.4305	3.09e-110
Year: 1990-1994	0.3364	2.25e-65	Make: Mitsubishi	0.0650	0.00135
Year: 1995-1999	0.4220	1.51e-105	Make: Nissan	-0.1273	3.02e-10
Year: 2000-2004	-0.2078	4.08e-25	Make: Oldsmobile	0.4061	3.97e-97
Year: 2005-2009	-0.4040	4.68e-96	Make: Panoz	-0.0190	0.349
Year: 2010-2014	-0.3541	1.06e-72	Make: Plymouth	0.3036	5.54e-53
Make: Acura	-0.3951	1.3e-91	Make: Pontiac	0.3608	1.34e-75
Make: AM General	0.0806	6.91e-05	Make: Porsche	-0.3736	2.44e-81
Make: Aston Martin	-0.1640	4.18e-16	Make: Ram	0.1413	2.58e-12
Make: Audi	-0.4816	2.64e-141	Make: Rolls-Royce	-0.1119	3.19e-08
Make: Bentley	0.0386	0.0569	Make: Saab	-0.3086	8.85e-55
Make: BMW	-0.5318	1.81e-177	Make: Saturn	0.0705	0.000509
Make: Buick	0.4488	9.49e-121	Make: Scion	-0.2673	5e-41
Make: Cadillac	0.3722	1.01e-80	Make: Smart	-0.2644	3.66e-40
Make: Chevrolet	0.4747	8.44e-137	Make: Subaru	-0.3302	6.37e-63
Make: Chrysler	0.3358	3.92e-65	Make: Suzuki	-0.0052	0.8
Make: Daewoo	-0.0043	0.832	Make: Tesla	-0.1519	5.15e-14
Make: Dodge	0.5527	1.99e-194	Make: Toyota	-0.5068	9.3e-159
Make: Eagle	0.1834	7.98e-20	Make: Volkswagen	-0.5290	2.39e-175
Make: Ferrari	-0.1523	4.5e-14	Make: Volvo	-0.3879	4.55e-88

Table S6. Pearson r correlation coefficients and their associated p-values for each car attribute between the percentage of residents with a high school degree and each car attribute, at the zip code level.

Variable	Pearson's r	p-value	Variable	Pearson's r	p-value
Price	-0.2577	3.62e-38	Make: Fiat	-0.0947	2.92e-06
Cars/Image	-0.4257	1.43e-107	Make: Fisker	-0.1472	3.03e-13
MPG Highway	-0.3847	1.5e-86	Make: Ford	0.2781	2.18e-44
MPG City	-0.3933	1.01e-90	Make: Geo	0.0910	7e-06
Hybrid	-0.1792	5.51e-19	Make: GMC	0.2901	2.46e-48
Electric	-0.2099	1.35e-25	Make: Honda	-0.3382	4.53e-66
Foreign	-0.3834	6.37e-86	Make: Hummer	-0.0548	0.00693
Country: England	-0.4480	2.85e-120	Make: Hyundai	-0.1058	1.72e-07
Country: Germany	-0.4791	1.16e-139	Make: Infiniti	-0.3544	7.84e-73
Country: Italy	-0.2939	1.25e-49	Make: Isuzu	0.0305	0.133
Country: Japan	-0.3221	8.96e-60	Make: Jaguar	-0.2449	1.61e-34
Country: South Korea	-0.0805	7.07e-05	Make: Jeep	-0.0230	0.257
Country: Sweden	-0.2889	6.45e-48	Make: Kia	0.0052	0.799
Country: USA	0.3834	6.37e-86	Make: Lamborghini	-0.2414	1.48e-33
Body Type: Convertible	-0.0883	1.29e-05	Make: Land Rover	-0.3614	7.2e-76
Body Type: Coupe	0.0791	9.54e-05	Make: Lexus	-0.2781	2.03e-44
Body Type: Crew Cab	0.3601	2.67e-75	Make: Lincoln	-0.0891	1.08e-05
Body Type: Extended Cab	0.4391	4.61e-115	Make: Lotus	-0.1283	2.21e-10
Body Type: Hatchback	-0.3453	5.36e-69	Make: Maserati	-0.1971	1.06e-22
Body Type: Minivan	-0.0567	0.00514	Make: Maybach	-0.0698	0.000579
Body Type: Regular Cab	0.4194	3.99e-104	Make: Mazda	-0.2630	1.02e-39
Body Type: Sedan	-0.2780	2.2e-44	Make: Mclaren	-0.1225	1.37e-09
Body Type: SUV	-0.1038	2.89e-07	Make: Mercedes-Benz	-0.3848	1.29e-86
Body Type: Van	-0.1936	5.98e-22	Make: Mercury	0.0195	0.337
Body Type: Wagon	-0.1893	4.75e-21	Make: Mini	-0.3150	4.19e-57
Year: 1990-1994	0.2961	2.2e-50	Make: Mitsubishi	0.0614	0.00248
Year: 1995-1999	0.1242	8.16e-10	Make: Nissan	-0.1723	1.17e-17
Year: 2000-2004	-0.0270	0.183	Make: Oldsmobile	0.1694	4.19e-17
Year: 2005-2009	-0.2729	9.06e-43	Make: Panoz	-0.0402	0.0474
Year: 2010-2014	-0.2570	5.71e-38	Make: Plymouth	0.0452	0.0259
Make: Acura	-0.3662	5.09e-78	Make: Pontiac	0.1203	2.66e-09
Make: AM General	0.1799	3.94e-19	Make: Porsche	-0.3172	6.45e-58
Make: Aston Martin	-0.1173	6.6e-09	Make: Ram	0.1174	6.41e-09
Make: Audi	-0.4422	6.62e-117	Make: Rolls-Royce	-0.2056	1.33e-24
Make: Bentley	-0.0002	0.992	Make: Saab	-0.3364	2.3e-65
Make: BMW	-0.4488	9.54e-121	Make: Saturn	0.1309	9.42e-11
Make: Buick	0.1495	1.28e-13	Make: Scion	-0.0187	0.358
Make: Cadillac	0.1507	8.26e-14	Make: Smart	-0.2235	6.88e-29
Make: Chevrolet	0.3650	1.81e-77	Make: Subaru	-0.1721	1.32e-17
Make: Chrysler	0.0675	0.000878	Make: Suzuki	-0.0502	0.0133
Make: Daewoo	-0.0027	0.893	Make: Tesla	-0.0891	1.1e-05
Make: Dodge	0.4096	5.97e-99	Make: Toyota	-0.2077	4.24e-25
Make: Eagle	0.0409	0.0438	Make: Volkswagen	-0.3294	1.31e-62
Make: Ferrari	-0.1439	1.02e-12	Make: Volvo	-0.2526	1.08e-36

Table S7. Pearson r correlation coefficients and their associated p-values for each car attribute between the percentage of residents with a some amount of college-level education and each car attribute, at the zip code level.

Variable	Pearson's r	p-value	Variable	Pearson's r	p-value
Price	0.4762	9.32e-138	Make: Fiat	0.1001	7.63e-07
Cars/Image	0.1769	1.53e-18	Make: Fisker	0.1594	2.75e-15
MPG Highway	0.4188	8.61e-104	Make: Ford	-0.4310	1.68e-110
MPG City	0.4928	6.96e-149	Make: Geo	-0.1954	2.48e-22
Hybrid	0.1709	2.16e-17	Make: GMC	-0.3045	2.61e-53
Electric	0.2644	3.8e-40	Make: Honda	0.5044	5.55e-157
Foreign	0.5761	7.19e-215	Make: Hummer	0.0296	0.145
Country: England	0.4858	3.69e-144	Make: Hyundai	0.2421	9.39e-34
Country: Germany	0.5990	1.41e-236	Make: Infiniti	0.3671	2.2e-78
Country: Italy	0.2677	3.81e-41	Make: Isuzu	-0.1013	5.61e-07
Country: Japan	0.5101	4.13e-161	Make: Jaguar	0.0730	0.000314
Country: South Korea	0.2224	1.28e-28	Make: Jeep	0.1450	6.77e-13
Country: Sweden	0.4482	2.03e-120	Make: Kia	0.0877	1.51e-05
Country: USA	-0.5761	7.19e-215	Make: Lamborghini	0.2240	5.22e-29
Body Type: Convertible	0.2604	5.8e-39	Make: Land Rover	0.3980	4.81e-93
Body Type: Coupe	-0.1435	1.17e-12	Make: Lexus	0.4678	2.14e-132
Body Type: Crew Cab	-0.1952	2.66e-22	Make: Lincoln	-0.1939	5.15e-22
Body Type: Extended Cab	-0.3285	3e-62	Make: Lotus	0.1168	7.61e-09
Body Type: Hatchback	0.5563	1.63e-197	Make: Maserati	0.1667	1.33e-16
Body Type: Minivan	-0.0054	0.789	Make: Maybach	0.0604	0.00292
Body Type: Regular Cab	-0.4164	1.62e-102	Make: Mazda	0.3960	4.65e-92
Body Type: Sedan	0.0435	0.0322	Make: Mclaren	0.1201	2.86e-09
Body Type: SUV	0.3200	5.67e-59	Make: Mercedes-Benz	0.4196	2.85e-104
Body Type: Van	-0.1251	6.18e-10	Make: Mercury	-0.2760	9.64e-44
Body Type: Wagon	0.4356	4.55e-113	Make: Mini	0.4333	8.48e-112
Year: 1990-1994	-0.4318	5.94e-111	Make: Mitsubishi	-0.0962	2.04e-06
Year: 1995-1999	-0.4575	5.71e-126	Make: Nissan	0.0588	0.00375
Year: 2000-2004	0.2168	3.01e-27	Make: Oldsmobile	-0.3223	7.67e-60
Year: 2005-2009	0.4910	1.18e-147	Make: Panoz	0.1190	4.04e-09
Year: 2010-2014	0.4297	8.57e-110	Make: Plymouth	-0.2432	4.71e-34
Make: Acura	0.3970	1.45e-92	Make: Pontiac	-0.2527	1.01e-36
Make: AM General	-0.1202	2.76e-09	Make: Porsche	0.3995	8.61e-94
Make: Aston Martin	0.1597	2.35e-15	Make: Ram	-0.0559	0.00582
Make: Audi	0.5247	4.77e-172	Make: Rolls-Royce	0.1492	1.43e-13
Make: Bentley	-0.0555	0.00625	Make: Saab	0.3839	3.61e-86
Make: BMW	0.4894	1.41e-146	Make: Saturn	0.0338	0.0959
Make: Buick	-0.3060	7.86e-54	Make: Scion	0.2404	2.71e-33
Make: Cadillac	-0.3432	3.96e-68	Make: Smart	0.2721	1.68e-42
Make: Chevrolet	-0.5134	1.48e-163	Make: Subaru	0.3973	1.04e-92
Make: Chrysler	-0.2066	7.9e-25	Make: Suzuki	0.0255	0.209
Make: Daewoo	-0.0122	0.548	Make: Tesla	0.1318	7.03e-11
Make: Dodge	-0.5082	8.88e-160	Make: Toyota	0.4168	9.85e-103
Make: Eagle	-0.1532	3.08e-14	Make: Volkswagen	0.5222	4.34e-170
Make: Ferrari	0.1161	9.6e-09	Make: Volvo	0.4221	1.32e-105

Table S8. Pearson r correlation coefficients and their associated p-values for each car attribute between the percentage of residents with a bachelor's degree and each car attribute, at the zip code level.

Variable	Pearson's r	p-value	Variable	Pearson's r	p-value
Price	0.4799	3.15e-140	Make: Fiat	0.1333	4.16e-11
Cars/Image	0.2212	2.65e-28	Make: Fisker	0.2006	1.81e-23
MPG Highway	0.4867	9.11e-145	Make: Ford	-0.4457	6.03e-119
MPG City	0.5451	3.44e-188	Make: Geo	-0.2069	6.79e-25
Hybrid	0.2256	2.05e-29	Make: GMC	-0.3500	6.06e-71
Electric	0.3236	2.45e-60	Make: Honda	0.5075	3.17e-159
Foreign	0.5775	3.65e-216	Make: Hummer	0.0311	0.125
Country: England	0.5434	8.86e-187	Make: Hyundai	0.2812	2.12e-45
Country: Germany	0.6397	8.72e-280	Make: Infiniti	0.3845	1.91e-86
Country: Italy	0.3448	8.8e-69	Make: Isuzu	-0.1266	3.78e-10
Country: Japan	0.4903	3.53e-147	Make: Jaguar	0.1406	3.33e-12
Country: South Korea	0.2560	1.16e-37	Make: Jeep	0.1283	2.21e-10
Country: Sweden	0.5266	1.77e-173	Make: Kia	0.0941	3.34e-06
Country: USA	-0.5775	3.65e-216	Make: Lamborghini	0.2799	5.56e-45
Body Type: Convertible	0.2634	7.39e-40	Make: Land Rover	0.4297	8.43e-110
Body Type: Coupe	-0.1454	6.02e-13	Make: Lexus	0.4415	1.86e-116
Body Type: Crew Cab	-0.2882	1.05e-47	Make: Lincoln	-0.1016	5.15e-07
Body Type: Extended Cab	-0.4198	2.35e-104	Make: Lotus	0.1358	1.8e-11
Body Type: Hatchback	0.5848	6.04e-223	Make: Maserati	0.2035	3.99e-24
Body Type: Minivan	-0.0024	0.905	Make: Maybach	0.0808	6.63e-05
Body Type: Regular Cab	-0.4727	1.7e-135	Make: Mazda	0.3948	1.91e-91
Body Type: Sedan	0.1410	2.89e-12	Make: Mclaren	0.1495	1.27e-13
Body Type: SUV	0.2720	1.79e-42	Make: Mercedes-Benz	0.4664	1.53e-131
Body Type: Van	-0.0593	0.00345	Make: Mercury	-0.1878	1.01e-20
Body Type: Wagon	0.4746	9.77e-137	Make: Mini	0.4604	9.23e-128
Year: 1990-1994	-0.4526	4.82e-123	Make: Mitsubishi	-0.1283	2.21e-10
Year: 1995-1999	-0.4365	1.25e-113	Make: Nissan	0.0365	0.0721
Year: 2000-2004	0.1785	7.62e-19	Make: Oldsmobile	-0.2907	1.59e-48
Year: 2005-2009	0.5065	1.59e-158	Make: Panoz	0.1316	7.47e-11
Year: 2010-2014	0.4501	1.58e-121	Make: Plymouth	-0.2022	7.88e-24
Make: Acura	0.4350	9.08e-113	Make: Pontiac	-0.2142	1.3e-26
Make: AM General	-0.1582	4.44e-15	Make: Porsche	0.4353	6.14e-113
Make: Aston Martin	0.2038	3.39e-24	Make: Ram	-0.0590	0.00363
Make: Audi	0.5867	8.92e-225	Make: Rolls-Royce	0.2186	1.11e-27
Make: Bentley	-0.0199	0.326	Make: Saab	0.4599	1.93e-127
Make: BMW	0.5301	3.5e-176	Make: Saturn	0.0157	0.441
Make: Buick	-0.2406	2.42e-33	Make: Scion	0.1821	1.49e-19
Make: Cadillac	-0.3113	9.66e-56	Make: Smart	0.3045	2.64e-53
Make: Chevrolet	-0.5376	4.31e-182	Make: Subaru	0.4357	4.02e-113
Make: Chrysler	-0.1708	2.34e-17	Make: Suzuki	0.0211	0.298
Make: Daewoo	0.0039	0.849	Make: Tesla	0.1434	1.22e-12
Make: Dodge	-0.5311	5.83e-177	Make: Toyota	0.3773	4.59e-83
Make: Eagle	-0.1099	5.66e-08	Make: Volkswagen	0.5278	2.34e-174
Make: Ferrari	0.1898	3.87e-21	Make: Volvo	0.4940	1.02e-149

Table S9. Pearson r correlation coefficients and their associated p-values for each car attribute between the percentage of residents with a graduate or professional degree and each car attribute, at the zip code level.

	%Less Than High School		%High School		%Some College	
Rank	Variable	Pearson's r	Variable	Pearson's r	Variable	Pearson's r
1	Year: 1995-1999	0.3715	Country: USA	0.6048	Body Type: Extended Cab	0.4391
2	Year: 1990-1994	0.3230	Make: Dodge	0.5527	Body Type: Regular Cab	0.4194
3	Make: Chevrolet	0.3183	Make: Chevrolet	0.4747	Make: Dodge	0.4096
4	Make: Ford	0.2991	Make: Buick	0.4488	Country: USA	0.3834
5	Country: USA	0.2955	Year: 1995-1999	0.4220	Make: Chevrolet	0.3650
84	Make: Hyundai	-0.3468	Make: Honda	-0.5371	Cars/Image	-0.4257
85	Country: Sweden	-0.3541	Body Type: Hatchback	-0.5576	Make: Audi	-0.4422
86	Make: Subaru	-0.3597	Country: Japan	-0.5670	Country: England	-0.4480
87	Year: 2005-2009	-0.3831	Foreign	-0.6048	Make: BMW	-0.4488
88	Body Type: Wagon	-0.3888	Country: Germany	-0.6142	Country: Germany	-0.4791

	%Bachelor's Degree		%Graduate Degree	
Rank	Variable	Pearson's r	Variable	Pearson's r
1	Country: Germany	0.5990	Country: Germany	0.6397
2	Foreign	0.5761	Make: Audi	0.5867
3	Body Type: Hatchback	0.5563	Body Type: Hatchback	0.5848
4	Make: Audi	0.5247	Foreign	0.5775
5	Make: Volkswagen	0.5222	MPG City	0.5451
84	Year: 1990-1994	-0.4318	Year: 1990-1994	-0.4526
85	Year: 1995-1999	-0.4575	Body Type: Regular Cab	-0.4727
86	Make: Dodge	-0.5082	Make: Dodge	-0.5311
87	Make: Chevrolet	-0.5134	Make: Chevrolet	-0.5376
88	Country: USA	-0.5761	Country: USA	-0.5775

Table S10. The five car attributes that correlate most positively and most negatively with the percentage of each education level in zip code.

Variable	Pearson's r	p-value	Variable	Pearson's r	p-value
Price	0.2182	1.39e-27	Make: Fiat	0.0370	0.0679
Cars/Image	-0.1478	2.42e-13	Make: Fisker	0.0113	0.579
MPG Highway	-0.0753	0.000205	Make: Ford	0.0559	0.00582
MPG City	-0.0008	0.967	Make: Geo	-0.0557	0.00606
Hybrid	-0.0068	0.739	Make: GMC	0.0536	0.00819
Electric	0.0560	0.00575	Make: Honda	0.0045	0.823
Foreign	0.0358	0.0778	Make: Hummer	0.0877	1.49e-05
Country: England	0.0422	0.0376	Make: Hyundai	0.1538	2.46e-14
Country: Germany	0.0572	0.00478	Make: Infiniti	-0.1273	3.03e-10
Country: Italy	0.0290	0.153	Make: Isuzu	-0.0218	0.283
Country: Japan	0.0075	0.712	Make: Jaguar	-0.2120	4.21e-26
Country: South Korea	0.1635	4.96e-16	Make: Jeep	0.2893	4.6e-48
Country: Sweden	0.0749	0.000219	Make: Kia	0.1149	1.34e-08
Country: USA	-0.0358	0.0778	Make: Lamborghini	0.0407	0.0451
Body Type: Convertible	0.0727	0.000332	Make: Land Rover	0.0686	0.000718
Body Type: Coupe	-0.1948	3.29e-22	Make: Lexus	-0.0263	0.194
Body Type: Crew Cab	0.1986	4.98e-23	Make: Lincoln	-0.3003	7.7e-52
Body Type: Extended Cab	0.2041	2.97e-24	Make: Lotus	0.0438	0.0308
Body Type: Hatchback	0.1702	2.95e-17	Make: Maserati	-0.0308	0.129
Body Type: Minivan	-0.0093	0.647	Make: Maybach	-0.0355	0.08
Body Type: Regular Cab	0.1237	9.42e-10	Make: Mazda	0.0734	0.000295
Body Type: Sedan	-0.4181	1.84e-103	Make: Mclaren	0.0066	0.745
Body Type: SUV	0.3053	1.37e-53	Make: Mercedes-Benz	-0.1011	5.86e-07
Body Type: Van	-0.1390	6.01e-12	Make: Mercury	-0.2581	2.76e-38
Body Type: Wagon	0.2153	7.21e-27	Make: Mini	0.1182	5.13e-09
Year: 1990-1994	-0.1668	1.28e-16	Make: Mitsubishi	-0.0663	0.00107
Year: 1995-1999	-0.2599	8.08e-39	Make: Nissan	-0.1289	1.8e-10
Year: 2000-2004	0.1514	6.23e-14	Make: Oldsmobile	-0.2065	8.01e-25
Year: 2005-2009	0.2104	1e-25	Make: Panoz	0.0313	0.123
Year: 2010-2014	0.1864	1.97e-20	Make: Plymouth	-0.1291	1.69e-10
Make: Acura	-0.0497	0.0142	Make: Pontiac	-0.1455	5.67e-13
Make: AM General	0.0649	0.00137	Make: Porsche	0.0839	3.48e-05
Make: Aston Martin	0.0098	0.629	Make: Ram	0.0635	0.00174
Make: Audi	0.1198	3.18e-09	Make: Rolls-Royce	-0.0234	0.25
Make: Bentley	-0.1299	1.29e-10	Make: Saab	0.1443	8.86e-13
Make: BMW	-0.0357	0.0789	Make: Saturn	0.0590	0.00359
Make: Buick	-0.2529	8.89e-37	Make: Scion	0.1415	2.42e-12
Make: Cadillac	-0.3535	1.93e-72	Make: Smart	0.1069	1.29e-07
Make: Chevrolet	0.0097	0.632	Make: Subaru	0.2397	4.34e-33
Make: Chrysler	-0.2134	1.96e-26	Make: Suzuki	0.0349	0.0855
Make: Daewoo	-0.0324	0.111	Make: Tesla	0.0339	0.0946
Make: Dodge	0.0648	0.00139	Make: Toyota	0.0034	0.867
Make: Eagle	-0.1144	1.55e-08	Make: Volkswagen	0.1827	1.09e-19
Make: Ferrari	-0.0134	0.509	Make: Volvo	0.0531	0.00887

Table S11. Pearson r correlation coefficients and their associated p-values for each car attribute between the percentage of white residents and each car attribute, at the zip code level.

Variable	Pearson's r	p-value	Variable	Pearson's r	p-value
Price	-0.1895	4.4e-21	Make: Fiat	-0.0250	0.217
Cars/Image	-0.0710	0.000459	Make: Fisker	0.0304	0.134
MPG Highway	0.0465	0.0218	Make: Ford	0.0432	0.033
MPG City	-0.0832	4.05e-05	Make: Geo	0.0196	0.335
Hybrid	0.0466	0.0217	Make: GMC	0.0281	0.167
Electric	-0.0942	3.32e-06	Make: Honda	-0.2304	1.24e-30
Foreign	-0.2580	2.97e-38	Make: Hummer	-0.0968	1.76e-06
Country: England	-0.0640	0.00161	Make: Hyundai	0.0248	0.221
Country: Germany	-0.1650	2.75e-16	Make: Infiniti	0.0334	0.0998
Country: Italy	-0.0033	0.87	Make: Isuzu	-0.0659	0.00116
Country: Japan	-0.2855	8.43e-47	Make: Jaguar	0.2494	8.76e-36
Country: South Korea	0.0273	0.178	Make: Jeep	-0.1016	5.19e-07
Country: Sweden	-0.0272	0.181	Make: Kia	0.0188	0.355
Country: USA	0.2580	2.97e-38	Make: Lamborghini	-0.0171	0.399
Body Type: Convertible	-0.0265	0.191	Make: Land Rover	-0.0994	9.16e-07
Body Type: Coupe	0.2151	7.87e-27	Make: Lexus	-0.1338	3.58e-11
Body Type: Crew Cab	-0.1996	2.93e-23	Make: Lincoln	0.4088	1.57e-98
Body Type: Extended Cab	-0.2583	2.52e-38	Make: Lotus	-0.0515	0.0111
Body Type: Hatchback	-0.2692	1.32e-41	Make: Maserati	0.0424	0.0366
Body Type: Minivan	-0.0707	0.00049	Make: Maybach	0.0712	0.000443
Body Type: Regular Cab	-0.1403	3.68e-12	Make: Mazda	-0.1301	1.22e-10
Body Type: Sedan	0.4421	7.52e-117	Make: Mclaren	-0.0252	0.215
Body Type: SUV	-0.2267	1.05e-29	Make: Mercedes-Benz	-0.0204	0.314
Body Type: Van	0.0799	7.98e-05	Make: Mercury	0.4479	3.12e-120
Body Type: Wagon	-0.1386	6.71e-12	Make: Mini	-0.1564	9.09e-15
Year: 1990-1994	0.0795	8.65e-05	Make: Mitsubishi	0.0176	0.386
Year: 1995-1999	0.2483	1.86e-35	Make: Nissan	-0.0285	0.16
Year: 2000-2004	-0.1324	5.73e-11	Make: Oldsmobile	0.3670	2.24e-78
Year: 2005-2009	-0.1417	2.25e-12	Make: Panoz	0.1011	5.96e-07
Year: 2010-2014	-0.1408	3.11e-12	Make: Plymouth	0.2056	1.3e-24
Make: Acura	-0.0751	0.000212	Make: Pontiac	0.3529	3.4e-72
Make: AM General	-0.0647	0.00142	Make: Porsche	-0.0925	4.9e-06
Make: Aston Martin	0.0086	0.671	Make: Ram	0.0667	0.000998
Make: Audi	-0.1189	4.17e-09	Make: Rolls-Royce	0.0955	2.39e-06
Make: Bentley	0.1577	5.27e-15	Make: Saab	-0.0704	0.000516
Make: BMW	-0.1151	1.29e-08	Make: Saturn	0.0510	0.0119
Make: Buick	0.4922	1.73e-148	Make: Scion	-0.2472	3.83e-35
Make: Cadillac	0.5015	6.37e-155	Make: Smart	-0.1127	2.58e-08
Make: Chevrolet	0.1058	1.71e-07	Make: Subaru	-0.1821	1.46e-19
Make: Chrysler	0.4137	4.29e-101	Make: Suzuki	-0.0578	0.00436
Make: Daewoo	0.0187	0.358	Make: Tesla	-0.0672	0.000924
Make: Dodge	0.1010	6.02e-07	Make: Toyota	-0.3335	3.34e-64
Make: Eagle	0.2003	2.04e-23	Make: Volkswagen	-0.2372	1.97e-32
Make: Ferrari	0.0199	0.326	Make: Volvo	-0.0153	0.451

Table S12. Pearson r correlation coefficients and their associated p-values for each car attribute between the percentage of black residents and each car attribute, at the zip code level.

Variable	Pearson's r	p-value	Variable	Pearson's r	p-value
Price	0.1130	2.32e-08	Make: Fiat	-0.0092	0.651
Cars/Image	0.3179	3.54e-58	Make: Fisker	-0.0110	0.587
MPG Highway	0.2580	2.93e-38	Make: Ford	-0.3324	8.99e-64
MPG City	0.3379	5.95e-66	Make: Geo	-0.0159	0.433
Hybrid	0.0140	0.49	Make: GMC	-0.2607	4.7e-39
Electric	0.1512	6.73e-14	Make: Honda	0.5174	1.56e-166
Foreign	0.5162	1.34e-165	Make: Hummer	-0.0058	0.775
Country: England	0.1797	4.47e-19	Make: Hyundai	-0.1209	2.28e-09
Country: Germany	0.3445	1.14e-68	Make: Infiniti	0.2498	6.76e-36
Country: Italy	0.0566	0.00525	Make: Isuzu	0.0369	0.0687
Country: Japan	0.5727	7.39e-212	Make: Jaguar	0.0279	0.169
Country: South Korea	-0.1476	2.67e-13	Make: Jeep	-0.2754	1.51e-43
Country: Sweden	0.0857	2.33e-05	Make: Kia	-0.1391	5.78e-12
Country: USA	-0.5162	1.34e-165	Make: Lamborghini	0.0505	0.0127
Body Type: Convertible	0.0381	0.0602	Make: Land Rover	0.1286	1.98e-10
Body Type: Coupe	-0.0233	0.251	Make: Lexus	0.4142	2.16e-101
Body Type: Crew Cab	-0.1247	6.92e-10	Make: Lincoln	-0.1779	1e-18
Body Type: Extended Cab	-0.1183	4.99e-09	Make: Lotus	0.0609	0.00265
Body Type: Hatchback	0.3293	1.54e-62	Make: Maserati	0.0355	0.0804
Body Type: Minivan	0.1799	4.09e-19	Make: Maybach	-0.0151	0.458
Body Type: Regular Cab	-0.1631	5.93e-16	Make: Mazda	0.2303	1.33e-30
Body Type: Sedan	0.0961	2.05e-06	Make: Mclaren	0.0742	0.00025
Body Type: SUV	-0.1002	7.39e-07	Make: Mercedes-Benz	0.3549	4.99e-73
Body Type: Van	0.0378	0.0628	Make: Mercury	-0.2835	3.67e-46
Body Type: Wagon	0.0187	0.356	Make: Mini	0.1919	1.41e-21
Year: 1990-1994	-0.0363	0.0739	Make: Mitsubishi	0.0137	0.5
Year: 1995-1999	-0.1500	1.08e-13	Make: Nissan	0.1730	8.91e-18
Year: 2000-2004	0.0434	0.0324	Make: Oldsmobile	-0.2771	4.28e-44
Year: 2005-2009	0.0888	1.17e-05	Make: Panoz	-0.1239	8.9e-10
Year: 2010-2014	0.0879	1.43e-05	Make: Plymouth	-0.1663	1.56e-16
Make: Acura	0.3251	6.24e-61	Make: Pontiac	-0.3100	2.91e-55
Make: AM General	-0.0515	0.0111	Make: Porsche	0.1269	3.48e-10
Make: Aston Martin	0.0266	0.19	Make: Ram	-0.1782	8.81e-19
Make: Audi	0.1654	2.31e-16	Make: Rolls-Royce	-0.0467	0.0213
Make: Bentley	-0.0346	0.0878	Make: Saab	0.0064	0.753
Make: BMW	0.3801	2.18e-84	Make: Saturn	-0.1075	1.08e-07
Make: Buick	-0.3356	4.73e-65	Make: Scion	0.2085	2.85e-25
Make: Cadillac	-0.2621	1.86e-39	Make: Smart	0.0740	0.000261
Make: Chevrolet	-0.3734	2.81e-81	Make: Subaru	0.0245	0.226
Make: Chrysler	-0.2971	1.06e-50	Make: Suzuki	0.0179	0.378
Make: Daewoo	0.0249	0.22	Make: Tesla	0.0979	1.33e-06
Make: Dodge	-0.4053	9.68e-97	Make: Toyota	0.6340	2.59e-273
Make: Eagle	-0.1243	7.72e-10	Make: Volkswagen	0.2052	1.66e-24
Make: Ferrari	0.0504	0.0129	Make: Volvo	0.0953	2.51e-06

Table S13. Pearson r correlation coefficients and their associated p-values for each car attribute between the percentage of Asian residents and each car attribute, at the zip code level.

	%White		%Black		%Asian	
Rank	Variable	Pearson's r	Variable	Pearson's r	Variable	Pearson's r
1	Body Type: SUV	0.3053	Make: Cadillac	0.5015	Make: Toyota	0.6340
2	Make: Jeep	0.2893	Make: Buick	0.4922	Country: Japan	0.5727
3	Make: Subaru	0.2397	Make: Mercury	0.4479	Make: Honda	0.5174
4	Price	0.2182	Body Type: Sedan	0.4421	Foreign	0.5162
5	Body Type: Wagon	0.2153	Make: Chrysler	0.4137	Make: Lexus	0.4142
84	Make: Mercury	-0.2581	Foreign	-0.2580	Make: Ford	-0.3324
85	Year: 1995-1999	-0.2599	Body Type: Extended Cab	-0.2583	Make: Buick	-0.3356
86	Make: Lincoln	-0.3003	Body Type: Hatchback	-0.2692	Make: Chevrolet	-0.3734
87	Make: Cadillac	-0.3535	Country: Japan	-0.2855	Make: Dodge	-0.4053
88	Body Type: Sedan	-0.4181	Make: Toyota	-0.3335	Country: USA	-0.5162

Table S14. The five car attributes that correlate most positively and most negatively with the percentage of each race in a zip code.

Variable	Pearson's r	p-value	Variable	Pearson's r	p-value
Price	-0.2768	$\leq 10^{-300}$	Make: Fiat	0.0165	0.00815
Cars/Image	0.3718	$\leq 10^{-300}$	Make: Fisker	0.0173	0.00543
MPG Highway	0.3307	$\leq 10^{-300}$	Make: Ford	-0.1746	7.41e-176
MPG City	0.2597	$\leq 10^{-300}$	Make: Geo	0.0681	6.63e-28
Hybrid	0.0318	3.1e-07	Make: GMC	-0.1675	8.16e-162
Electric	0.0347	2.35e-08	Make: Honda	0.0705	8.59e-30
Foreign	0.0743	5.93e-33	Make: Hummer	-0.0587	3.72e-21
Country: England	0.1159	6.69e-78	Make: Hyundai	-0.0162	0.00938
Country: Germany	0.1665	7.7e-160	Make: Infiniti	0.0772	2.08e-35
Country: Italy	0.0563	1.32e-19	Make: Isuzu	0.0369	3.12e-09
Country: Japan	0.0297	1.83e-06	Make: Jaguar	0.1180	1.09e-80
Country: South Korea	-0.0150	0.0158	Make: Jeep	-0.0563	1.39e-19
Country: Sweden	0.1509	1.96e-131	Make: Kia	-0.0076	0.22
Country: USA	-0.0743	5.93e-33	Make: Lamborghini	0.0366	4.03e-09
Body Type: Convertible	0.0144	0.0206	Make: Land Rover	0.0438	1.87e-12
Body Type: Coupe	0.1426	2.52e-117	Make: Lexus	-0.0578	1.5e-20
Body Type: Crew Cab	-0.4799	$\leq 10^{-300}$	Make: Lincoln	0.1387	3.71e-111
Body Type: Extended Cab	-0.4266	$\leq 10^{-300}$	Make: Lotus	0.0147	0.0178
Body Type: Hatchback	0.1193	1.6e-82	Make: Maserati	0.0291	2.94e-06
Body Type: Minivan	0.0524	3.38e-17	Make: Maybach	0.0160	0.01
Body Type: Regular Cab	-0.3047	$\leq 10^{-300}$	Make: Mazda	0.0801	5.49e-38
Body Type: Sedan	0.4829	$\leq 10^{-300}$	Make: Mclaren	0.0301	1.3e-06
Body Type: SUV	-0.2246	1e-292	Make: Mercedes-Benz	0.0830	1.08e-40
Body Type: Van	0.1154	2.73e-77	Make: Mercury	0.1830	2.67e-193
Body Type: Wagon	0.1463	1.97e-123	Make: Mini	0.0700	2.19e-29
Year: 1990-1994	0.1396	1.89e-112	Make: Mitsubishi	0.0101	0.104
Year: 1995-1999	0.2950	$\leq 10^{-300}$	Make: Nissan	0.0212	0.000668
Year: 2000-2004	-0.1728	3.4e-172	Make: Oldsmobile	0.2250	9.64e-294
Year: 2005-2009	-0.1951	7.58e-220	Make: Panoz	0.0128	0.0392
Year: 2010-2014	-0.1651	2.71e-157	Make: Plymouth	0.1696	7.47e-166
Make: Acura	0.1032	3.93e-62	Make: Pontiac	0.2191	3.61e-278
Make: AM General	-0.1182	6.02e-81	Make: Porsche	0.0089	0.153
Make: Aston Martin	0.0005	0.931	Make: Ram	-0.0855	4.51e-43
Make: Audi	0.0951	6.14e-53	Make: Rolls-Royce	0.0537	6.09e-18
Make: Bentley	0.0481	1e-14	Make: Saab	0.0784	1.86e-36
Make: BMW	0.1203	8.72e-84	Make: Saturn	0.0358	8.69e-09
Make: Buick	0.2177	1.57e-274	Make: Scion	-0.1093	1.66e-69
Make: Cadillac	0.1144	6.54e-76	Make: Smart	0.0141	0.0232
Make: Chevrolet	-0.1842	8.44e-196	Make: Subaru	0.1414	1.72e-115
Make: Chrysler	0.1708	2.51e-168	Make: Suzuki	0.0396	1.98e-10
Make: Daewoo	0.0453	3.09e-13	Make: Tesla	-0.0059	0.347
Make: Dodge	-0.1010	1.43e-59	Make: Toyota	-0.0545	1.97e-18
Make: Eagle	0.0776	9.26e-36	Make: Volkswagen	0.1529	8.78e-135
Make: Ferrari	0.0463	1.01e-13	Make: Volvo	0.1442	5.31e-120

Table S15. Pearson r correlation coefficients and their associated p-values between each car attribute and %Obama. The variables "Price", "MPG City", and "MPG Highway" are calculated as expected values for each precinct, and all other variables are expressed as a percent of all cars observed in each precinct.