Quantifying the Impact of the Built Environment on Neighborhood Crime Rates

Adyasha Maharana [1] **Quynh C. Nguyen** [2] **Elaine O. Nsoesie** [3]

Abstract

The built environment has been postulated to have an impact on neighborhood crime rates; however, measures of the built environment can be subjective and differ across studies, leading to varying observations on its association with crime rates. We illustrate an approach to quantify the impact of the built environment on neighborhood crime rates from high-resolution satellite imagery. We apply convolutional neural networks to 150,000 satellite images for three United States cities to construct an indicator of the built environment. We then assess the association between the built environment indicator, socio-demographic factors and neighborhood crime rates. Our study suggests that the built environment may be a stronger predictor of neighborhood crime rates than socioeconomic and demographic factors. Identification of specific features that are linked to higher crime rates can lead to structural interventions shown to reduce crime incidence in urban settings.

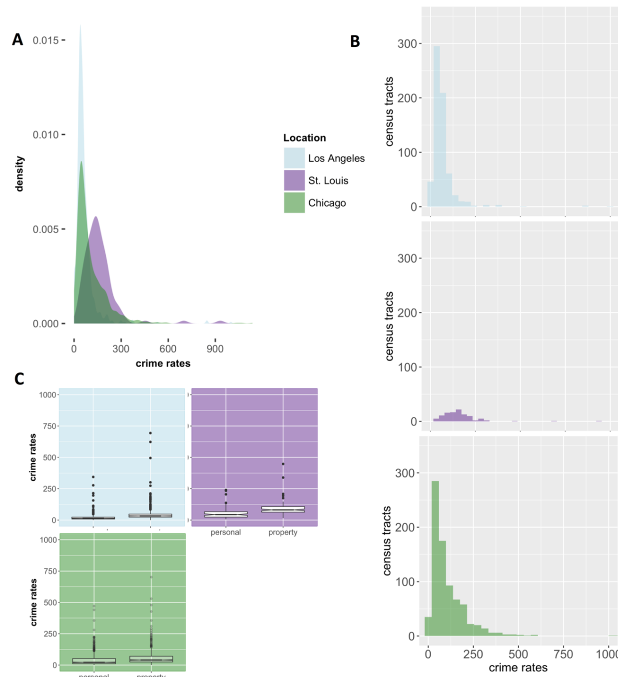

Figure 1. Distribution of crimes per 1,000 persons in census tracts of three metropolitan cities (Los Angeles, St. Louis and Chicago) in the US. As shown in A and B, only a small proportion of census tracts are classified as high crime regions. The boxplots in C show variation in personal and property crime rates at the level of census tracts.

1. Introduction

Aspects of the built environment such as, building design, street layouts, land use, and environmental disrepair and desolation, have been associated with crime incidence (Taylor & Harrell, 1996). However, different built environment characteristics influence particular types of crimes and may work through different mechanisms for crime inducement. For example, the presence of high schools, public parks, vacant lots or buildings can invite gang-related crimes (Loukaitou-Sideris et al., 2001; Spelman, 1993). Neighborhood crime has also been associated with structural aspects of the environment linked to the degree of accessibility, and the ease of entry and exit (Greenberg & Rohe, 1984). In contrast, structural changes in urban neighborhoods have been associated with a reduction in crime rates. For example, a study conducted in London observed that improving lighting in urban streets led to decreases in crime and increases in pedestrian street use after dark (Painter, 1996). In another community intervention in Sarasota, Florida, improvements in city lighting, landscaping, the addition of balconies or porches and residential units to commercial areas combined with new police initiatives for drug dealing and prostitution led to decreases in personal and property crime (Carter et al., 2003).

It is easy to visually identify environmental attributes, however, quantifying the density of these attributes across different geographic regions, populations and over time can

[1]Sciome LLC, Durham, NC, USA [2]University of Maryland, College Park, MD, USA [3]Boston University, Boston, MA, USA. Correspondence to: Adyasha Maharana <adyasha@uw.edu>, Elaine Nsoesie <onelaine@bu.edu>.

Appearing at the International Conference on Machine Learning AI for Social Good Workshop, Long Beach, United States, 2019.

be cumbersome. Studies linking neighborhood crime to features of the physical environment have mostly been conducted using costly and time-consuming onsite visits to count relevant attributes (e.g., the number of liquor stores, vacant lots, and ratings of the level of graffiti or liter in the vicinity of interest) or neighborhood surveys to assess participant perceptions of their neighborhood. The resulting data can therefore be subjective since it relies upon participant or researcher perceptions, and assessment tools that vary across studies. Furthermore, sample sizes for most neighborhood studies tend to be small due to the burden of data collection.

Here, we demonstrate a scalable approach that combines a convolutional neural network (CNN) and satellite imagery to infer characteristics of the built environment to assess the degree to which the built environment can be associated with variations in crime rates. We use the pretrained CNN as a fixed feature extractor, aggregate features for each census tract and train a regression model using those features to predict the annual crime rate in all census tracts. We apply our method to census tracts in three cities (Chicago, Illinois; St. Louis, Missouri; and Los Angeles, California) with high crime rates and available geo-referenced crime data. In contrast to existing methods, our approach is low cost, and can produce fine-grained geographical estimates using publicly available data and software.

2. Related Work

High-resolution satellite imagery are rich and comprehensive repositories of information for a variety of domains, ranging from crop health to the economy (Nsoesie et al., 2015; Karnieli et al., 2008; Gautam et al., 2004; Verbesselt et al., 2012). Recent studies have shown that the application of deep neural networks to satellite images can enable characterization of the physical environment to study poverty, obesity, the economy and the demographic makeup of the United States (Gebru et al., 2017; Jean et al., 2016; Maharana & Nsoesie, 2018).

In this rapidly evolving field, there have been no studies focused on the prediction of population level crime rates using data from satellite images. A related study adopted the Broken Windows theory (Wilson & Kelling, 2003) to identify city landscape features from Google Street View image for crime prediction using support vector regression algorithm (Arietta et al., 2014). This approach was tested for several US cities and achieved more than seventy percent accuracy in binary classification of areas with low and high rates of violent crime. However, Google Street View is only available for select regions, thereby limiting reproducibility. Another study used multi-modal features to classify crime hot-spots in Chicago (Bogomolov et al., 2014). In contrast, our approach is the first comprehensive assessment of the association between the built environment and overall numeric crime rates at the neighborhood level.

3. Methods

Our modeling approach involves three steps: (1) extraction of geocoded crime data from police websites and census tract sociodemographic variables from the American Community Survey (ACS), (2) collection and processing of satellite imagery data, and (3) regression modeling to assess association between crime, the built environment and sociodemographic factors.

3.1. Crime and Sociodemographic Data

We obtained geo-referenced time-stamped 2016 crime records (includes both serious crimes and misdemeanors) provided by law enforcement departments for each of the cities.[1][2][3] The number of crimes were aggregated to the census tracts in accordance with 2010 Census boundaries. As appropriate, some crimes were further separated into categories of personal (e.g., assault, battery, homicide) and property (e.g., robbery, property destruction) crime. We also obtained 2014 5-year estimates of socioeconomic and population characteristics from the ACS. The number of crimes for each census tract was divided by the ACS population estimates to arrive at the number of crime incidents per 1,000 persons (hereafter referred to as crime rates).

3.2. Satellite Image Processing

Next, we collected nearly 150,000 satellite images spanning each census tract from Google Static Maps API at a zoom level of 18. These images were unlabeled. To overcome the challenge of working with unlabeled data, we used a transfer learning framework (Jean et al., 2016). We used the VGG-F network which has been pre-trained on the ImageNet database; a dataset containing approximately 14 million images to differentiate between 1,000 object categories (Simonyan & Zisserman, 2015; Deng et al., 2009). We fine-tuned the network to our specific problem of crime prediction by training the model on data comprised of images from census tracts with high and low crime rates defined as the top and lower fifteen percent of crimes rates. The fine-tuned VGG-F network achieved approximately 80% classification accuracy on the validation set after 30 epochs. The updated model identified features pertinent to

[1]Crime data for Chicago was downloaded from `https://catalog.data.gov/dataset/crimes-2001-to-present-398a4`

[2]Crime reports for St. Louis are updated at `http://www.slmpd.org/Crimereports.shtml`

[3]Data on Los Angeles is available at `https://data.lacity.org/A-Safe-City/Crime-Data-from-2010-to-Present/y8tr-7khq`

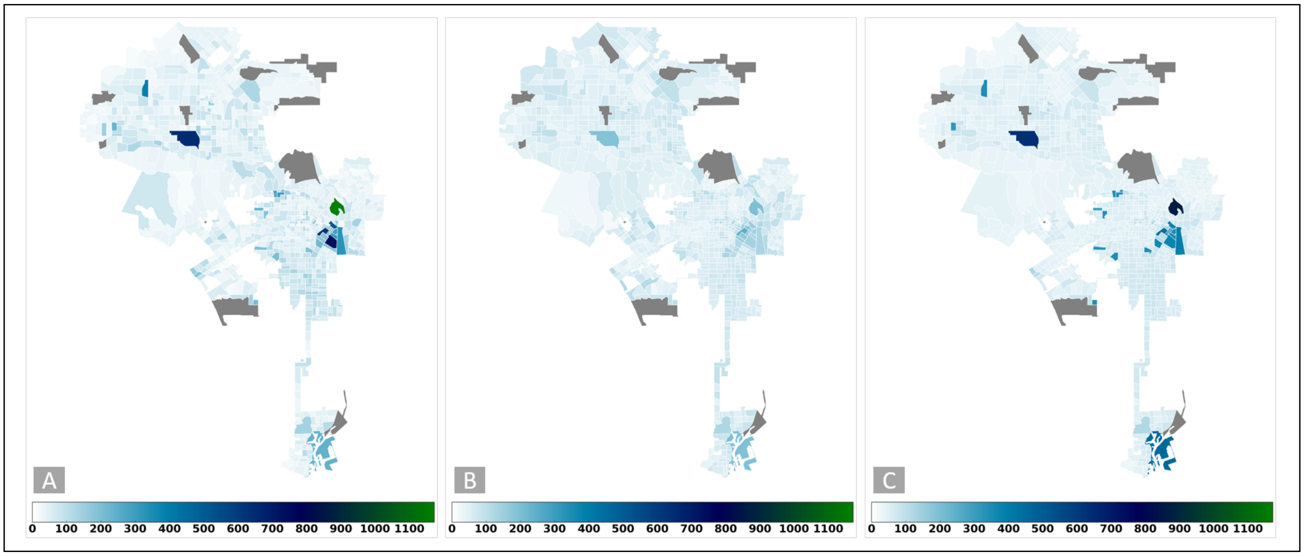

Figure 2. Predictions of crime rates per 1,000 persons across census tracts in Los Angeles, California. (A) the reported crime rates, (B) predictions for a single model fitted to all crime levels, and (C) combined predictions from the separate model without outliers and model for outliers (i.e., census tracts with highest crime rates). The gray shaded regions are census tracts with zero population or no reported crimes.

describing neighborhood structures, which could be useful for making meaningful associations between crime rates and the environment. A total of 4,096 features were extracted from the penultimate layer of the neural network.

3.3. Statistical Analysis

We fitted a regression model to assess the association between the extracted features and neighborhood crime rates. To reduce the dimensionality of the feature matrix and identify relevant predictive features, we used elastic net, which is both a regularization and variable selection technique (Friedman et al., 2001). Our rigorous model fitting approach included a five-fold (for Chicago and Los Angeles) and three-fold (for St. Louis due to small sample size) cross validation process. For each of the three cities, we fitted individual models to predict (predict here refers to out-of-sample predictions, not forecasting of future events) crime rates in each city solely using the features extracted from satellite images. We also assessed how well our approach predicts personal and property crimes. Next, we compared our findings to predictions solely based on socioeconomic and demographic variables, which have been extensively studied and associated with neighborhood crime rates. We developed regression models to predict crime rates based on variables related to unemployment, income, racial demographics, and education (Bogomolov et al., 2014; Raphael & Winter-Ebmer, 2001; Patterson, 1991; Ehrlich, 1975; Fajnzylber et al., 2002). Model estimates were evaluated by comparing the Pearson correlation coefficient (r), and root mean squared error (RMSE).

4. Results

The average crime rates per 1,000 people was 105.8 (95% CI, 99.2-112.4), 60.5 (95% CI, 56.5-64.5) and 165.5 (95% CI, 143.2-187.7) separately for Chicago (n=798 census tracts), Los Angeles (n=991) and St. Louis (n=106). Personal crime rates were higher for Chicago (39.6 [95% CI, 36.5- 42.8]), and St. Louis (53.8 [95% CI, 44.9-62.7]) compared to Los Angeles (17.1 [95% CI, 15.7-18.5]). In contrast, property crime rates were higher for St. Louis (92.1 [95% CI, 81.8-102.4]), compared to Chicago (54.1 [95% CI, 50.8- 57.3]) and Los Angeles (41.1 [95% CI, 38.5-43.8]).

Regression models using the built environment indicator achieved variable predictions of crime rates across census tracts in the three US cities. The association between the estimated and reported crime rates was higher for St. Louis (r=0.76, $p<0.01$) and Chicago (r=0.70, $p<0.01$), compared to Los Angeles (r=0.58, $p<0.01$). These associations were noted both in trend and magnitude (see Figure 1-2). Census tracts with highest crime rates, which can be classified statistically as outliers, were underestimated in some cases, although typically predicted as the highest in each city.

Separate models for each city to investigate these deviations while excluding outliers, showed that our models could reproduce spatial variations in crime rates across census tracts in Chicago (r=0.88, $p<0.01$), Los Angeles (r=0.75,

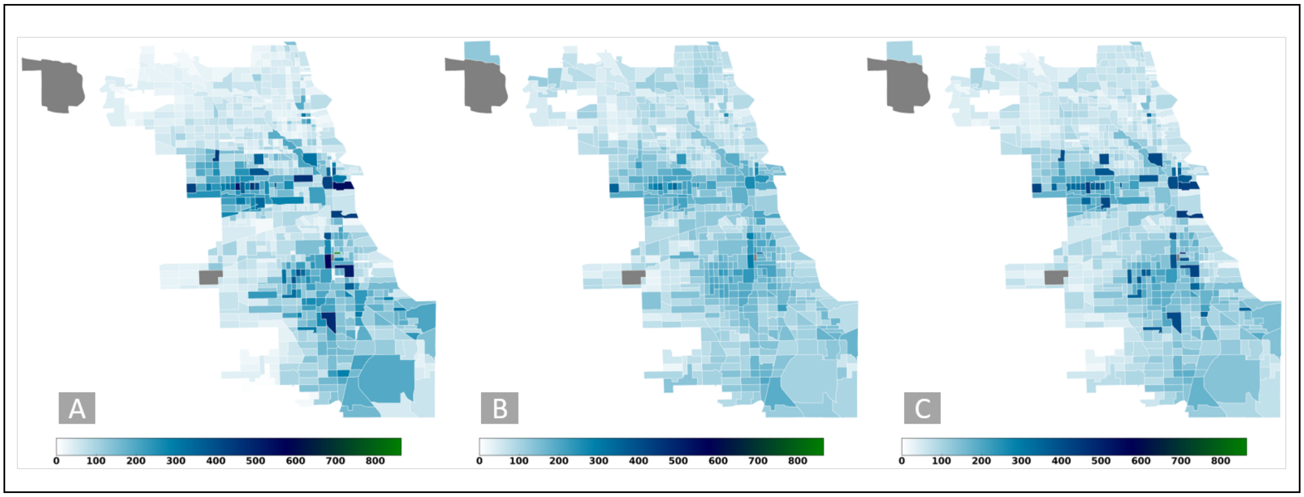

Figure 3. Predictions of crime rates per 1,000 persons across census tracts in Chicago, Illinois. (A) the reported crime rates, (B) predictions for a single model fitted to all crime levels, and (C) combined predictions from the separate model without outliers and model for outliers (i.e., census tracts with highest crime rates). The gray shaded regions are census tracts with zero population or no reported crimes.

$p<0.01$) and St. Louis ($r=0.84$, $p<0.01$).

These associations could be partially explained by socioeconomic status, which studies have demonstrated can be inferred from satellite images (Jean et al., 2016). Although we observed significant associations between some of the sociodemographic variables and personal crime rates for Chicago, these associations were much weaker for the other cities.

The association between overall reported and model estimated crime rates using socioeconomic and demographic variables was consistent across the three cities; 0.61 ($p<0.01$), 0.55 ($p<0.01$) and 0.66 ($p<0.01$) for Chicago, Los Angeles and St. Louis (Figures S3-S5). After the removal of census tracts with highest crime rates, the correlation increased to 0.73 ($p<0.01$), 0.54 ($p<0.01$) and 0.74 ($p<0.01$) for Chicago, Los Angeles and St. Louis. These values were lower than those observed for the models solely based on the built environment indicator inferred from satellite imagery across all cities.

Seventy-one data points were identified as outliers across the three cities. Our models explained approximately 67.6% and 16.6% of the variation in crime rates across census tracts for the models using the built environment and sociodemographic variables, respectively. The population employed was approximately 48.9% (95% CI, 44.5%-53.3%) and percent with income below poverty was 11.8% (95% CI, 10.2%-13.4%) for these census tracts. The tight confidence intervals suggest similarities in sociodemographic factors within high crime census tracts across cities. However, the most significant predictor of crime in these census tracts was education (i.e., percent of population 25 years and over with some college or associates degree). Also, these regions tend to be more easily accessible, with higher population density. Additionally, when shown a specific census tract our approach in many cases can identify regions (i.e. grids) with the highest crime rates based on the built environment indicator.

5. Discussion & Conclusion

Neighborhood crime rates can be explained by a complex interaction of environmental, societal, and individual level factors. While socioeconomic and demographic variables have been presented as predictors of crime, these factors do not completely explain neighborhood crime rates. In this study, we quantified the variation in crime rates at the census tract level across three cities that are explainable by features of the physical environment. Our results suggest that characteristics of the built environment are able to distinguish high and low crime areas above and beyond residential compositional characteristics. We also observe that our models are predictive of overall crime rates and do not favor personal or property crime. The accuracy of predicting personal and property crime rates was comparable when using socioeconomic variables or the built environment indicator. Additionally, census tracts with high crime rates heavily influence predictions when models are fitted to all the data. The differences in predictions across cities might indicate that for some cities the physical environment might explain variations in crime rates better than for others.

However, there are some limitations to our modeling approach. First, we assume that our crime data is accurate. There is extensive criminology research suggesting that crime databases only represent a biased sample of all criminal offence (Langan, 1995; Morrison, 1897; Levitt, 1998). These data are influenced by several factors including, existing police priorities and crime incidence reporting. For example, although drug crimes tend to be widely distributed, police arrests on drug offences tend to be concentrated in lower income and high non-white population neighborhoods (Lum & Isaac, 2016). In our data, we noted higher crime rates were reported in census tracts with lower income and with a higher percentage of blacks. Therefore, although our methods provide some quantitative association between the physical environment and crime rates, it should not be used as the sole predictor of crime rates. Also, processes are needed to address the bias inherent in these data. Second, the various data sources used in our analysis have been collected at different points in time. The 5-year census estimates and crime reports were released in 2014 and 2016 respectively, while the time-stamp of satellite imagery is unknown. We assume that the census estimates aren't drastically different from the true population distribution in 2016, but we acknowledge that the temporal mismatch might have led to small error margins in our analysis. We attempted to collect satellite imagery from the same year as the crime reports, but weren't able to acquire images of appropriate resolution during the time of the study.

Our results are in alignment with empirical research suggesting there is a relationship between physical disorder and fear of crime and crime rates (Wilcox et al., 2004; Sampson & Raudenbush, 2004). Thus, interventions to change the environment may help to prevent future crime, and improve individual-level health and psychological functioning (Hinkle & Weisburd, 2008). One potential approach used for environmental interventions is Crime Prevention Through Environmental Design (CPTED). CPTED is centered around incorporating design features that promote safety and security within a community (Karnieli et al., 2008). Design features include the following: natural surveillance, access control, territorial reinforcement, activity support, and maintenance. These design features help reduce opportunity for crime, increase social control, provide reassurance to community members by signaling to observers that disorder is not tolerated. Examples of specific community design strategies include installing outside lighting to entrances, walkways, and parking lots; decreasing visual barriers and concealed areas (e.g., underpasses); building fencing and walls to demarcate public and private property; designing landscaping with ground cover and tree canopy to allow for visibility and demonstrate ownership; responding to maintenance issues (e.g., graffiti); and providing recreational facilities and structural support for safe activities.

References

Arietta, S. M., Efros, A. A., Ramamoorthi, R., and Agrawala, M. City forensics: Using visual elements to predict non-visual city attributes. *IEEE transactions on visualization and computer graphics*, 20(12):2624–2633, 2014.

Bogomolov, A., Lepri, B., Staiano, J., Oliver, N., Pianesi, F., and Pentland, A. Once upon a crime: towards crime prediction from demographics and mobile data. In *Proceedings of the 16th international conference on multimodal interaction*, pp. 427–434. ACM, 2014.

Carter, S. P., Carter, S. L., and Dannenberg, A. L. Zoning out crime and improving community health in sarasota, florida:crime prevention through environmental design. *American Journal of Public Health*, 93(9):1442–1445, 2003.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009.

Ehrlich, I. On the relation between education and crime. In *Education, income, and human behavior*, pp. 313–338. NBER, 1975.

Fajnzylber, P., Lederman, D., and Loayza, N. Inequality and violent crime. *The journal of Law and Economics*, 45(1):1–39, 2002.

Friedman, J., Hastie, T., and Tibshirani, R. *The elements of statistical learning*. Number 10. Springer series in statistics New York, 2001.

Gautam, A. P., Shivakoti, G. P., and Webb, E. L. Forest cover change, physiography, local economy, and institutions in a mountain watershed in nepal. *Environmental management*, 33(1):48–61, 2004.

Gebru, T., Krause, J., Wang, Y., Chen, D., Deng, J., Aiden, E. L., and Fei-Fei, L. Using deep learning and google street view to estimate the demographic makeup of neighborhoods across the united states. *Proceedings of the National Academy of Sciences*, 114(50):13108–13113, 2017.

Greenberg, S. W. and Rohe, W. M. Neighborhood design and crime a test of two perspectives. *Journal of the American Planning Association*, 50(1):48–61, 1984.

Hinkle, J. C. and Weisburd, D. The irony of broken windows policing: A micro-place study of the relationship between disorder, focused police crackdowns and fear of crime. *Journal of Criminal justice*, 36(6):503–512, 2008.

Jean, N., Burke, M., Xie, M., Davis, W. M., Lobell, D. B., and Ermon, S. Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301):790–794, 2016.

Karnieli, A., Gilad, U., Ponzet, M., Svoray, T., Mirzadinov, R., and Fedorina, O. Assessing land-cover change and degradation in the central asian deserts using satellite image processing and geostatistical methods. *Journal of Arid Environments*, 72(11):2093–2105, 2008.

Langan, P. A. The racial disparity in us drug arrests. *Bureau of Justice Statistics (BJS) and US Dept of Justice and Office of Justice Programs and United States of America*, 1995.

Levitt, S. D. The relationship between crime reporting and police: Implications for the use of uniform crime reports. *Journal of Quantitative Criminology, XIV*, pp. 61–81, 1998.

Loukaitou-Sideris, A., Liggett, R., Iseki, H., and Thurlow, W. Measuring the effects of built environment on bus stop crime. *Environment and Planning B: Planning and Design*, 28(2):255–280, 2001.

Lum, K. and Isaac, W. To predict and serve? *Significance*, pp. 14–19, 2016.

Maharana, A. and Nsoesie, E. O. Use of deep learning to examine the association of the built environment with prevalence of neighborhood adult obesity. *JAMA network open*, 1(4):e181535–e181535, 2018.

Morrison, W. D. The interpretation of criminal statistics. *Journal of the Royal Statistical Society*, pp. 1–32, 1897.

Nsoesie, E. O., Butler, P., Ramakrishnan, N., Mekaru, S. R., and Brownstein, J. S. Monitoring disease trends using hospital traffic data from high resolution satellite imagery: A feasibility study. *Scientific reports*, 5:9112, 2015.

Painter, K. The influence of street lighting improvements on crime, fear and pedestrian street use, after dark. *Landscape and urban planning*, 35(2-3):193–201, 1996.

Patterson, E. B. Poverty, income inequality, and community crime rates. *Criminology*, 29(4):755–776, 1991.

Raphael, S. and Winter-Ebmer, R. Identifying the effect of unemployment on crime. *The Journal of Law and Economics*, 44(1):259–283, 2001.

Sampson, R. J. and Raudenbush, S. W. Seeing disorder: Neighborhood stigma and the social construction of broken windows. *Social psychology quarterly*, 67(4):319–342, 2004.

Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2015.

Spelman, W. Abandoned buildings: magnets for crime? *Journal of Criminal Justice*, 21(5):481–495, 1993.

Taylor, R. B. and Harrell, A. *Physical environment and crime*. US Department of Justice, Office of Justice Programs, National Institute of , 1996.

Verbesselt, J., Zeileis, A., and Herold, M. Near real-time disturbance detection using satellite image time series. *Remote Sensing of Environment*, 123:98–108, 2012.

Wilcox, P., Quisenberry, N., Cabrera, D. T., and Jones, S. Busy places and broken windows? toward defining the role of physical structure and process in community crime models. *Sociological Quarterly*, 45(2):185–207, 2004.

Wilson, J. Q. and Kelling, G. L. Broken windows: The police and neighborhood safety. *Criminological perspectives: essential readings*, 400, 2003.

Supplementary Material: Quantifying the Impact of the Built Environment on Neighborhood Crime Rates

Adyasha Maharana [1] Quynh C. Nguyen [2] Elaine O. Nsoesie [3]

1. Materials and Methods

1.1. Crime Data

Some police departments in metropolitan cities in the United States provide citizens with an up-to-date digital database of crime incidents occurring throughout the city. Every report is associated with a Uniform Crime Reporting (UCR) code, which classifies the offense into one of several categories, such as murder, sexual assault, robbery, aggravated assault etc. Crimes are also tagged with date and time of occurrence, report and geolocation. In this paper, we analyzed crime reports from 2016 for the cities of Chicago, Illinois; St. Louis, Missouri; and Los Angeles, California (Figure 1). In the United States, crimes are categorized into Part I and Part II offenses (also known as index and non-index offenses) according to the UCR program (U.S. Department of Justice, 2004). Part I crimes are considered serious offenses and are further divided into personal and property crimes. These include homicide, rape, aggravated assault, robbery, burglary, motor vehicle theft, larceny-theft and arson. Other crimes such as public peace violation, kidnapping, illegal dumping, liquor law violation, fraud etc. are classified as Part II offenses. Crimes that inflict physical, emotional or psychological harm to the victim are said to be personal crimes. On the other hand, property crimes are committed with the objective of acquiring money or materialistic objects and may or may not be accompanied by the use of force. Our study aims to understand the impact of the built environment on the incidence of both Part I and II crimes. Hence, our definition of personal and property crimes was modified to encompass both Part I and II crimes.

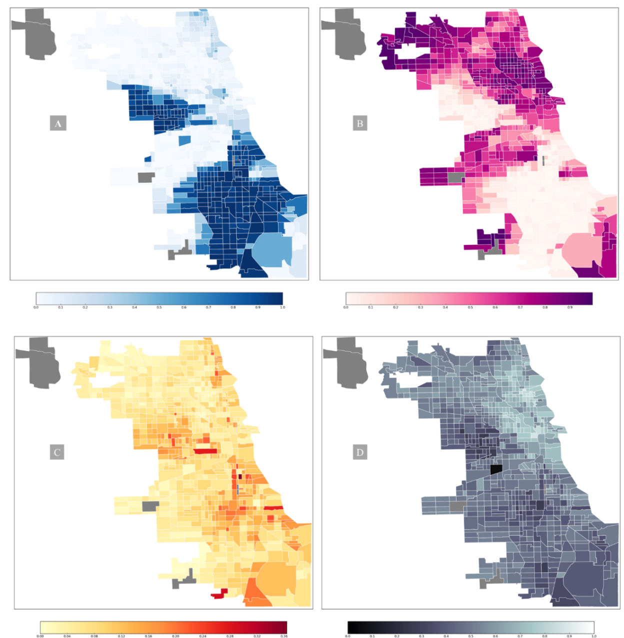

Figure 1. Plots of socioeconomic variables with high significance in prediction of crime rate across census tracts in Chicago, Illinois. The figures show the percentage of white population (A), percentage of black population (B), percentage population with income below poverty (C), and percentage of unemployed population (D). The gray shaded regions are census tracts with zero population or no reported crimes

1.2. Socio-economic Demographic Data

Crime datasets available from respective police departments contained information on both Part I and Part II crimes. In all crime datasets, there were several columns to indicate the crime codes, crime type, description, date and time for occurrence and reporting of crime, several levels of address for the general location of crime, (x, y) coordinates in the State Plane Coordinate System (SPCS) or latitude, longitude information. Often, it was also indicated whether the incident is a new record or an update to a previous record. Some records lacked accurate location information, and addresses were available only at the block level. Due to the uncertainty of geolocation, these records were left out of our analysis. The records were filtered for the year 2016 by the occurrence date. The crime reporting codes vary from one state to another; hence, we examined the description string to select relevant reports. From the cleaned dataset, geocoded crimes were aggregated to the level of census tracts in

[1]Sciome LLC, Durham, NC, USA [2]University of Maryland, College Park, MD, USA [3]Boston University, Boston, MA, USA. Correspondence to: Adyasha Maharana <adyasha@uw.edu>, Elaine Nsoesie <onelaine@bu.edu>.

Appearing at the International Conference on Machine Learning AI for Social Good Workshop, 2019

accordance with the boundaries established by 2010 United States Census. Chicago, Los Angeles and St. Louis have 807, 679 and 106 census tracts respectively. The overall crime counts were normalized with population estimates for each census tract. Our analysis involved first assessing the association between the built environment and overall crime rates. Next, we assessed the association between the built environment and personal crimes and property crimes. Lastly, we evaluated how well socioeconomic and demographic factors could predict crime rates

We used the American Community Survey (ACS) 5-year estimates of the following variables at the census tract level (ACS code: 080): population, poverty status in the last 12 months, race, gross median rent, education, population by age, gender, employment status, and total land area (Figure 2). The specific socioeconomic and demographic variables considered included, percent employed, percent income below poverty level, percent white population, median rent (in dollars), per capita income, percent male, percent female, percent high school graduate (includes equivalency) (25 years and over), percent with some college or associate's degree (25 years and over), percent with bachelor's degree or higher (25 years and over), population density, percent population below 10 years of age, percent population between 10 and 20 years of age, percent population between 20 and 30 years of age, percent population between 30 and 40 years of age, percent population between 40 and 50 years of age, percent population above 50 years of age, percent black population, percent Hispanic population, percent Asian population and percent American-Indian population. Aggregated numbers for crime incidence were normalized with ACS 5-year population estimates to obtain number of crime incidents per 1,000 persons for each tract. Tracts with zero population estimates and more than 100% error margin in population estimate were excluded from our analysis. For census tracts located in Los Angeles city and other county sub-divisions, we use regional data only for the area that lies within Los Angeles.

1.3. Satellite Imagery Data

Each census tract was converted into a square grid where each location is separated by roughly 150 meters from the adjacent location. Satellite images were downloaded from the freely available Google Static Maps API, at a zoom level of 18. Our dataset comprised of nearly 100,000 images for the three metropolitan cities. Using the location information from geo-coded crimes, each satellite image was assigned an annual crime count. This resulted in a dataset consisting of images ranging from very high to zero crime occurrence areas. Areas with zero or low crime occurrences tend to be the ones with city highways or waterways, recreational areas (football fields, parks) and areas with no human habitation. On the other hand, areas with higher crime counts generally belonged to downtown or densely populated neighborhoods. The top and bottom 10% of images in this spectrum of crime counts were selected to prepare a binary classification dataset, which was used for fine-tuning the pre-trained convolutional neural networks. Some images from this dataset are presented in Figure 2.

1.4. Deep Neural Network & Transfer Learning

Convolutional neural networks have been at the helm of breakthrough research in computer vision for several years. However, such architectures need enormous amount of training data and time to achieve this level of proficiency in image understanding. In the recent years, transfer learning has enabled smaller datasets to be used to accomplish complicated tasks such as, training a network to identify high crime regions using satellite images, and in lesser time. Transfer learning is the method of enhancing learning for task B by transferring knowledge from task A which has already been learnt (Pan & Yang, 2010). For transfer learning to be effective, tasks A and B need to be similar to some extent. It has been shown that the features extracted from ImageNet-trained convolutional networks can be successfully deployed as feature vectors in linear classifiers and regression models for various downstream tasks (Yosinski et al., 2014). The transferability of features increases as one moves towards the earlier layers usually the final convolutional layers or the first few fully connected layer tend to perform best in such tasks. Moreover, the pre-trained models can be fine-tuned to a related task with controlled training and small datasets. As before, features also extracted from the fine-tuned models can be utilized in other machine-learning methods. Therefore, we employ fine-tuning and transfer learning to identify features of the built environment from satellite images and associate them with neighborhood crime rates. This approach is similar to that described by Jean et al (Jean et al., 2016); a two-step transfer learning model using convolutional neural networks. We used a pre-trained network VGG-F which has been trained for object recognition on the ImageNet database and was the runner-up in ILSVRC 2014 (Simonyan & Zisserman, 2015). The output layer of this network was originally designed for classifying between 1000 object categories. We modified the output layer for binary classification and fine-tuned the convolutional neural network to distinguish between low and high crime regions from satellite images. Fine-tuning guides the network to learn complex filters which are more relevant to criminal activity. The fine-tuned VGG-F network achieved approximately 80% classification accuracy on the validation set after 30 epochs. Visualization of filters from the convolutional layer revealed that the filters learned to recognize physical attributes such as green cover, buildings and roads. This architecture has a depth of eight of which the final three are fully-connected layers. Feature vectors

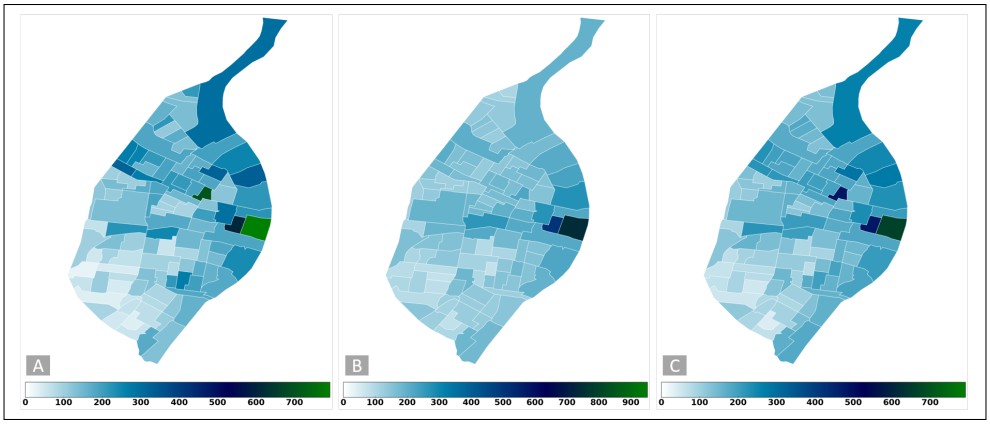

Figure 2. Predictions of crime rates per 1,000 persons across census tracts in St. Louis, Missouri. (A) the reported crime rates, (B) predictions for a single model fitted to all crime levels (C) combined predictions from the separate model without outliers and model for outliers (i.e., census tracts with highest crime rates).

extracted from the seventh layer were one-dimensional arrays of length 4096 and were extracted from a forward pass of each of the satellite images in our dataset through this fine-tuned net. The final predictions for crime rates were performed at the level of census tracts. Hence, features from all images belonging to the same tract were aggregated to produce a single feature vector of dimension 4096 for that tract. These feature vectors were inputs for the next step of transfer learning in our model.

1.5. Elastic Net Regression Model

We used elastic net, a regularization and variable selection technique4 to reduce the dimensionality of our feature vector and identify the most predictive features. Elastic Net is a regularization method that integrates the advantages of Ridge regression and Least Absolute Shrinkage and Selection Operator (LASSO) so that it eliminates variables that are not significant in prediction, while maintaining correlated variables that are significant in the model. For each of the three cities, we fitted an individual model to predict crime rates solely using features extracted from satellite images. We fitted another model separately for low and median crime rates (referred to as low crime), and another for high crime rates based on the assumption that high crime regions might have features that are uniquely different when compared to other census tracts. Additionally, high crime rates heavily influenced the model outcomes when the entire dataset was used. The high crime regions were defined using the statistical definition of outliers as follows, crime rates

$$greater\ than\ max\ \{(\overline{x}\ +\ 2 * \sigma_x)\ and\ (1.5 * IQR(x))\}$$

where x is the mean crime rate, σ_x is the standard deviation and IQR is the interquartile range.

We also assessed cross city performance by predicting crime rates for each of the other cities based on the model fitted to data for one city. For example, the model fit to Chicago was used to predict crime rates in St. Louis. We did this to assess the generalizability of the models developed for each city. Next, we developed separate models for predicting crime rates using demographic and socioeconomic variables that have been used extensively in the crime literature. The crime rates were standardized prior to model fitting. We used a five-fold cross validation approach, which involves splitting the data into five separate groups. Each group is used in the model fitting and the data points (i.e., census tracts) not included in the model are then predicted. The models were fit using the glmnet package in R.

2. Additional Results & Conclusion

The associations observed between the built environment indicator and crime rates could be partially explained by socioeconomic status, which studies have demonstrated can be inferred from satellite images (Figure 3 shows findings for St. Louis). Although we observed significant associations between some of the sociodemographic variables and personal crime rates for Chicago, these associations were much weaker for the other cities (Figure 4). Specifically, variables highly associated with overall crime rates included percent income below poverty (r= 0.42, $p<0.01$), percent black population (r=0.58, $p<0.01$) and percent white population (r= -0.54, $p<0.01$). In addition, percent black population (r=0.67, $p<0.01$), percent income below poverty (r=0.48, $p<0.01$), percent white population (r=-0.63, $p<0.01$) and percent employed (r=-0.52, $p<0.01$) were also strongly associated with personal crime rates. These observations agree with reports on the distribution of crimes in Chicago (chi). Similarly, the strongest positive predictors of overall crime rates in Chicago were poverty, percent black population, and percent population between the ages of ten and twenty. In

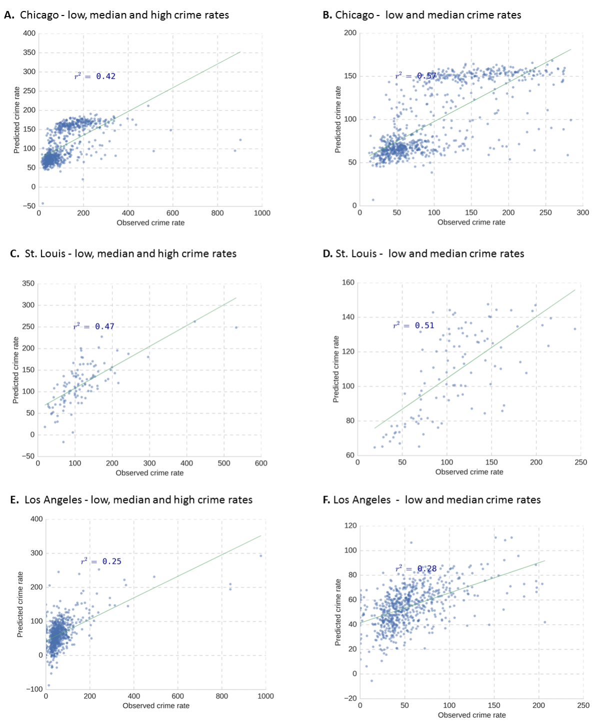

Figure 3. Scatterplots for crime rate predictions using socioeconomic variables. Results are shown for Chicago for the entire dataset (A) and the model without outliers (B). Similar results are presented for St. Louis, (C) and (D), and Los Angeles, (E) and (F). The r^2 values are based on five-fold cross-validation.

contrast, the strongest negative predictors were population density, and employment. We observed similar associations for St. Louis and Los Angeles. Education was also negatively associated with crime rates in Los Angeles.

Furthermore, associations between model estimates and reported personal crime rates were 0.78 (p<0.01), 0.61 (p<0.01) and 0.61 (p<0.01) for Chicago, Los Angeles and St. Louis, respectively. With the exclusion of census tracts with highest crime rates, the correlation increased to 0.88 (p<0.01), 0.76 (p<0.01) and 0.74 (p<0.01) for Chicago, Los Angeles and St. Louis respectively. Similar associations were observed for property crime rates; 0.69 (p<0.01), 0.44 (p<0.01) and 0.84 (p <0.01) for census tracts in Chicago, Los Angeles and St. Louis respectively. After excluding census tracts with highest crime rates, the correlation increased slightly for Chicago (0.71, p<0.01) but dropped for Los Angeles (0.41, p<0.01) and St. Louis (0.72, p<0.01).

Furthermore, the associations between the sociodemographic model estimates and reported personal crime rates were 0.72 (p<0.01), 0.60 (p<0.01) and 0.64 (p<0.01) for Chicago, Los Angeles and St. Louis, respectively. After excluding census tracts with the highest crime rates, the predictions improved slightly for Chicago (0.77, p<0.01), Los Angeles (0.70, p<0.01) and St. Louis (81.3, p<0.01). Compared to estimates using the built environment indicator, these estimates were slightly lower for Chicago and Los Angeles, but comparable for St. Louis.

Additionally, estimates using the socioeconomic data had a 0.51 (p<0.01), 0.50 (p<0.01) and 0.73 (p<0.01) correlation with reported property crime rates for Chicago, Los Angeles, and St. Louis, respectively. After excluding census tracts with the highest crime rates, the associations increased for Chicago (0.65, p<0.01) and St. Louis (0.80, p<0.01), and dropped for Los Angeles (0.48, p<0.01). These observations were comparable to estimates made using the built environment indicator.

The census tracts with some of the highest crime rates in Los Angeles and St. Louis include public parks. In spite of being low-populated areas, they register high incidence of crime which may be attributed to the regular inflow of crowd from nearby areas. Crime rates for such areas can be better predicted by taking the structural features of surrounding census tracts into account. Some census tracts with very high crime rates in Chicago barely span across four to five blocks and also have a large margin of error in their population estimates. Prediction of crime for such areas could also benefit from wider coverage of neighborhood for feature extraction. Advances in deep learning methods will make possible future improvements in feature extraction and predictions using these approaches.

References

Chicago police annual report. `https://home.chicagopolice.org/inside-the-cpd/statistical-reports/annual-reports/`. Accessed: 2017-08-08.

Jean, N., Burke, M., Xie, M., Davis, W. M., Lobell, D. B., and Ermon, S. Combining satellite imagery and machine learning to predict poverty. *Science*, 353:790–794, 2016.

Pan, S. J. and Yang, Q. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22: 1345–1359, 2010.

Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2015.

U.S. Department of Justice, F. B. o. I. *Uniform Crime Reporting Handbook*. U.S. Department of Justice, Federal Bureau of Investigation, 2004.

Yosinski, J., Clune, J., Bengio, Y., and Lipson, H. How transferable are features in deep neural networks? In *NIPS*, 2014.