Deep Learning Visual Methods for Identifying Abandoned Houses

Jim DeLisle
University of Missouri–Kansas City

Hye-Sung Han
University of Kansas

Duy H. Ho
Yugyung Lee
Brent Never
Ye Wang
University of Missouri–Kansas City

Authors are listed alphabetically, as each contributed equally to the development of this report.

The opinions expressed in this article are those of the authors; they do not necessarily reflect the views and policies of HUD or the U.S. Government.

Abstract

Housing abandonment contributes to neighborhood decline and disinvestment. Abandonment has plagued large metropolitan areas for decades, yet quantifying the scope and impact of abandonment has proven costly and elusive. This study introduces an innovative approach to detect and measure abandoned houses using technology innovations without requiring significant resource commitments. It presents a system of detecting abandoned houses leveraging deep learning models for image classification, building an ensemble model that considers both global and local contexts to identify abandoned structures. This study takes imagery and structure data from multiple sources and uses transfer learning in a three-stage ensemble approach to identify abandoned houses. Four deep learning models are constructed for this study: the ResNet-50 model, an incremental knowledge model, a hybrid approach, and a check model. Results from the different models are compared and analyzed to identify the visual characteristics of houses that improve or degrade each method's accuracy. The methodology presented herein is scalable and could be applied in other neighborhoods and communities. The data generated by this method will empower communities and cities to design more effective strategies to address housing abandonment.

Introduction

Housing abandonment contributes to neighborhood decline and disinvestment. Abandonment has plagued large metropolitan areas for decades (HUD PD&R, 2014). Quantifying the scope and the impact of abandonment has proven elusive in part because there is no agreed-upon way to identify and track abandonment (Bieretz and Schilling, 2019). This void has forced governments and scholars to rely on indirect measures, including vacancy data compiled by the U.S. Census Bureau or the United States Postal Service (USPS). These data do not provide an accurate measure of abandonment, failing to identify functionally abandoned houses unfit for human habitation (Schilling, 2002; U.S. Government Accountability Office [GAO], 2011). While some field-based methods have been applied (e.g., foot surveys), they require substantial resources and cannot keep up with abandoned properties' rapid pace of turnover (Morckel, 2012; Pagano and Bowman, 2000). The absence of accurate and cost-effective methods has contributed to the inability to address the abandonment crisis. The application of artificial intelligence (AI) and other emerging technologies presented in this study presents an innovative method for identifying the dynamic stock of abandoned houses. The data generated by this method will empower communities and cities to design more effective strategies to address housing abandonment.

Problem Statement

Housing Abandonment Issue in the United States

The interaction between blight, abandoned housing, and the adverse impacts on American cities have been recognized since the early 20th century. Robbins (1935) advocated that cities implement demolition programs to remove neglected and unfit housing to reverse the adverse impacts on residents and communities. The United States Housing Act of 1937 provided support for local demolition programs and mandated that cities demolish unsafe housing—both abandoned and occupied—making it a requirement for funding for new low-rent housing (Woodbury, 1937). Over the next two decades, the abandoned and substandard housing problem continued to plague cities, despite federal funding for demolition programs and local efforts to eradicate the problem, with a growing recognition that in many cases the problem was endemic and had sociological roots (Rosenthal, 1953; Schneider, 1941; Sjoberg, 1955).

During the early 1970s, renewed attention was paid to the housing abandonment problem, which had become a national phenomenon, rising to crisis proportions in some cities. At the same time, there was a shift from demolitions to the rehabilitation of abandoned housing. This shift received additional momentum through a change in the Community Development Grant Program that allowed cities to use neighborhood preservation and renewal funds outside of blighted areas (Cannon, Lachman, and Bernhard, 1977). In the 1980s, researchers looked at the efficacy of efforts to address the abandoned housing problem, including HUD's "Urban Homesteading Program" (Varady, 1984). Researchers also began advocating for additional programs to fund the rehabilitation of abandoned housing to supplant the prior emphasis on demolitions (Margulis and Sheets, 1985). Efforts were also made to develop a more proactive approach to the challenge of abandoned housing, including the application of discriminant analysis to determine which houses were more vulnerable to demolition, thus creating an early warning system that would reflect the

relative vulnerability of buildings that would render them "endangered" buildings that should be protected from demolition (Bell and Kelso, 1986).

The inability to resolve the abandonment challenge was attributed to the tangled web of causes and effects of urban decline, such as economic decline, job loss, quality of life decline, tax/mortgage delinquency, eviction/foreclosure, market distortions, and development barriers (Goldstein, Jensen, and Reiskin, 2001). The 2008 housing crisis exacerbated the abandoned housing problem with a tsunami of foreclosures leading to a surge in abandonment. The federally funded Neighborhood Stabilization Program generated significant data to support research into the efficacy of initiatives launched to resolve the foreclosure and abandonment problem stemming from the housing crisis (Bak and Hewings, 2017; Fraser and Oakley, 2015; Leonard, Jha, and Zhang, 2017; Schuetz, Spader, and Cortes, 2016).

State and local efforts were launched to help address the abandonment problem in distressed or blighted neighborhoods, but resource constraints and lack of access to timely data placed a damper on such initiatives and failed to address abandonment associated with systemic issues. To help address the abandonment problem, a growing number of states passed legislation enabling the creation of land banks and community land trusts (Decker, 2018; Fujii, 2016; Martin et al., 2020; Whitaker and Fitzpatrick, 2016). While land bank programs achieved some success in resolving the challenge of abandoned housing, these efforts struggled due to an adverse selection process whereby the properties transferred to land banks suffered from significant physical deterioration, due in part to the time lags between abandonment and transfer. Furthermore, land banks served as intermediaries and did not have the mission, budget, or staff to maintain or improve abandoned houses. In some markets, houses held by the land banks were exempt from property violations, leading to further deterioration; this circumstance made it challenging and expensive for private market efforts to remediate the problem and return abandoned housing to productive use, and it heightened interest in demolition programs, which merely shifted the problem to abandoned lots.

The abandoned housing problem has continued to plague many cities for over 80 years, especially following cyclical downturns. While some progress has been made, the efforts of researchers and advocates have been thwarted by the complexity and interactive nature of forces that have sustained the long-term abandonment problem (Foster and Hipp, 2011; Grinstein-Weiss et al., 2013; Keating, 2010; Mennis, Dayanim, and Grunwald, 2013). Since abandonment is a process rather than a discrete event, it is imperative that this study develop the ability to predict abandonment so that steps can be taken to reverse the forces or to develop rapid response models (Williams, Galster, and Verma, 2013). A prerequisite to such work, however, is the ability to detect properties that are clearly abandoned. While detection might appear to be relatively straightforward, the absence of an unambiguous definition of abandoned property complicates the process. That is, unlike foreclosed properties that go through a predetermined process that is documented along the way, the road to abandonment can take many twists and turns. For example, water shut-offs, postal records, tax delinquency, property violations, and 311 reports can be used, but such indicators might merely flag vacant properties rather than identify abandoned properties. Even if these trails can be mapped out along the journey, there are no consistent checkpoints along the route, and the destination (i.e., abandonment) is not even clear. Morckel (2014) tested the

premise that operational definitions of housing abandonment matter, which pointed to additional informational inefficiencies that had plagued researchers (e.g., tax delinquency as a proxy for abandonment, abandonment as identified by foot survey, and abandonment defined as "other vacant" by the Census Bureau). Depending on the definition, the application of spatial and other analytical models generates statistically significant outcomes and weightings that vary substantially from model to model. Immergluck and Smith (2006) recognized the lack of good data related to long-term vacancy problems after the housing crisis, utilizing data from the USPS to explore changes in vacancy and noting the difference in vacancy and abandonment across neighborhoods, with more persistence in poorer neighborhoods. Despite these efforts and those of other researchers, the inability to detect and predict abandonment frustrates the efforts to understand and thus develop valid and reliable intervention programs.

Challenges to Measure Housing Abandonment: Lack of Standardized Definition

One major challenge to measuring housing abandonment is the lack of a formal or standardized definition of abandoned houses. Wachsmuth (2008) states there are as many definitions of abandonment as municipal governments addressing abandonment and scholars writing about it. Among scholars, Mallach (2006) of the Brookings Institution considers a property abandoned if the owner has stopped carrying out at least one of the significant responsibilities of property ownership, causing a property to be vacant or likely to become vacant. Many scholars consider neglected property ownership duties—e.g., delinquent property taxes or noncompliance with relevant codes—as indicators of abandonment. Sternlieb et al. (1974) defined an abandoned building as a residential structure that the owner has removed from the housing stock by neglecting property ownership duties regarding functional, financial, and physical maintenance. Hillier et al. (2003) also identified three distinct aspects of abandonment: functional, financial, and physical. Functional abandonment concerns a vacant property that is not suitable for residency, such as one that lacks sealed doors and windows. Financial abandonment happens when an owner stops meeting his or her financial responsibilities, such as making property tax or mortgage payments. Physical abandonment occurs when a property is unfit for occupation because the owner neglected to maintain the inside or outside of the residence.

The definition of abandonment varies depending on state and municipal governments. The New Jersey Abandoned Properties Rehabilitation Act of 2004 defines abandoned property as any property that has not been legally occupied for 6 months and which also meets any one of the following criteria: (1) the property requires rehabilitation, and no rehabilitation has taken place during those 6 months; (2) construction began but was discontinued before the property was suitable for occupancy or use, and no construction has taken place for those 6 months; (3) at least one installment of property tax is delinquent; or (4) the property is determined to be a nuisance by the public officer.[1] The City of Kansas City, Missouri, classifies a vacant property as "vacant" or "dangerous." A property is defined as "vacant" if it lacks the habitual presence of human beings who have a legal right to be on the property or if any substantial lawful residential occupancy or business operation has ceased (City of Kansas City, Missouri, n.d.b). The following factors are

---

[1] Housing and Community Development Network of New Jersey. 2004. Abandoned Properties Rehabilitation Act, P.L. 2003, c.210. https://hcdnnj.memberclicks.net/assets/documents/npt_abandonedproprehabact.pdf

considered to determine whether a property is vacant: the proportion of vacant to occupied space, the condition and value of any items on the property, the presence of rental or for-sale signs on the property (implying the property is currently marketed by a licensed real estate professional), and water service not being shut off (City of Kansas City, Missouri, n.d.b). The City of Kansas City labels some vacant properties as dangerous buildings. Dangerous buildings exhibit the most severe type of residential code violations that pose the highest risk to surrounding areas (City of Kansas City, Missouri, n.d.a). Common conditions of dangerous buildings include exterior walls that are leaning; the building or any portion thereof is in danger of collapse; the building has been damaged by fire or earthquake; electrical, plumbing or other mechanical systems are dangerous or inoperable; the roof or walls have holes exposing the entry of weather; or a foundation that has settled or is damaged (City of Kansas City, Missouri, n.d.a).

Problems with Using Proxy Measures to Estimate Housing Abandonment

The lack of a universal definition, and thus the difficulty of obtaining an objective indicator of housing abandonment, makes it difficult to measure the scope of housing abandonment accurately. Municipal governments and scholars use indirect measures to estimate the scope of abandonment. The U.S. Census Bureau, for example, collects data on housing vacancy, classifying a housing unit as "Other vacant" if it is vacant year-round for reasons other than being vacant for sale or rent or occasional use. The "Other vacant" category also includes vacant and abandoned units, as well as units to be demolished or condemned. It also includes foreclosed housing units, units that are vacant because of legal issues, units that are currently being prepared to rent or sell, units that need repairs or are vacant for 6 months or longer, and units that are vacant for unknown reasons (U.S. Census Bureau, 2022). Thus, the U.S. Census data on housing vacancy fail to reflect the full depth of housing abandonment. The USPS provides data on addresses it identifies as "Vacant" or "No-Stat" (HUD PD&R, n.d.). Vacant addresses refer to those not collecting their mail for 90 days or longer. No-Stat addresses include those not collecting their mail for 90 days or longer and businesses or homes under construction and not yet occupied. USPS data categorizes these vacant addresses by the length of vacancy. Some scholars use the addresses that have been vacant for more than 24 months as a proxy for abandonment and blight. While USPS data provide information on occupancy status and the length of vacancy, USPS does not capture information on the nature of vacancy or the physical condition of the property. Therefore, it is difficult to use these data to identify abandoned houses because they overlook functionally abandoned houses that are unfit for human habitation or other authorized uses (Schilling, 2002; U.S. Government Accountability Office [GAO], 2011).

On a local level, municipal governments and scholars often use land bank data to understand the scope of abandonment. Land banks are governmental entities or nonprofit organizations that acquire vacant, abandoned, and tax-delinquent properties to return these properties to productive use (Heins and Abdelazim, 2014). However, land banks operate in different structures with varying abilities to acquire and redevelop their problem properties. One of the common challenges land banks face is the lack of precise data on these problem properties (Alexander, 2005). One reason for the lack of precise data on problem properties is fragmented functions among local government agencies and departments. For example, one department records housing and building code

violations while another agency records tax delinquency. Alexander (2005) argues that community development corporations or neighborhood agencies know their communities and can more accurately identify vacant, abandoned, and tax-delinquent properties.

Multiple warning signs often indicate the process of abandonment. For example, code violations and unpaid utility accounts are signs that a property is in the abandonment process (Hillier et al., 2003). Some scholars stress the importance of including property information, such as delinquent water and sewer bills and property-based nuisance complaints, as additional indicators of housing abandonment. A 311 call system can provide useful information on residential code violations. The data on 311 calls for service have been recently used in research to measure potential indicators of neighborhood physical disorder, capturing information such as neglected properties and public nuisance. However, some argue that 311 data can introduce biases, such as some neighborhoods or residents being unwilling or unable to call and log a 311 report (Theall et al., 2021).

Literature Review

Related Work on Detecting Abandoned Houses: Street Audits

Street audits are one of the traditional techniques for assessing neighborhood conditions and, by extension, the environmental factors that could lead to abandonment. Street audits and the use of streetscapes have received traction among researchers concerned with neighborhood conditions (Harvey and Aultman-Hall, 2016). The approach has been of particular interest to healthcare researchers interested in the health-related impacts of neighborhood conditions (Badland et al., 2015; Cain et al., 2014). In a novel approach, researchers have used "citizen scientists" to compile information about neighborhood conditions in disadvantaged communities as well as improve community engagement (Winter et al., 2016)

The Microscale Audit of Pedestrian Streetscapes (MAPS) was developed to assess details of streetscapes considered relevant for physical activity (Fox et al., 2021; Millstein et al., 2013; Vanwolleghem et al., 2016). Due to the time and resource commitments needed to support MAPS using observation, researchers have explored the validity of on-street versus online assessments and reported generally positive results (Cleland et al., 2021; Fox et al., 2021; Phillips et al., 2017; Queralt et al., 2021; Zhu et al., 2017). Virtual street audits have received significant attention as a means of assessing neighborhood conditions as compared to physical audits (Badland et al., 2010; Pliakas et al., 2017) Health-related researchers have also applied virtual street audits with an emphasis on neighborhood disorder (Mooney et al., 2017; Nesoff et al., 2020).

Related Work on Detecting Abandoned Houses: Deep Learning

Cities have increasingly sought to leverage the precipitous drop in the cost of sensors and cloud computing to better allocate finite municipal resources to the myriad of resident needs. Likewise, social scientists and engineers have leveraged artificial intelligence, generally, and visual analytics through deep learning, more specifically, to assist governments in better understanding the complex problems of their communities and to drive down costs. For example, Bloch (2020) assessed neighborhood characteristics by analyzing visual data from Google Street View and

municipal 311 calls. Recently, Zuo et al. (2020) evaluated mobility and sociability trends during COVID-19 using an interactive data visualization and analytics tool.

One promising technique has been the use of neural networks, and more specifically deep learning, as a means to classify aspects of images accurately. For example, Zou and Wang (2021) proposed a novel technique for detecting abandoned houses based on hierarchical deep learning. They designed two deep Convolutional Neural Network (ConvNet/CNN) models and extracted global visual features from the abandoned house using scene-based classification and local visual features from buildings, vegetation, and other objects using patch-based classification. This method allows researchers to accurately identify sub-elements of an image, such as doors or weeds; Zou and Wang successfully classified Google Street View images in five cities with an F1 score, which assesses both accuracy and precision of identification, at 0.84.

Stevenson and Bravo (2021) proposed to combine deep learning and K-means to cluster a series of convenient task-agnostic tile elevation embeddings with socioeconomic outcomes in predicting seven English deprivation indices for small geographies in the Greater London area. Briefly, task-agnostic tile elevation embedding is a deep learning model for learning structural node embeddings. It captures higher-level features of the urban environment numerically and can be imported into any analytics tool. The generated embeddings, either alone or combined with standard structured data, are used to assess their potential as auxiliary sources of data. Task-agnostic tile embeddings can also be derived from elevation data using unsupervised deep learning. After evaluating various model/embedding configurations, coherent tile segments enable the interpretation of latent embedding features, resulting in an increase in Root Mean Squared Error (RMSE) of up to 21 percent compared to standard demographic features alone.

Related Work on Detecting Abandoned Houses: Google Street View

Over the past decade, researchers have relied on Google Street View to observe and track neighborhood conditions (Anguelov et al., 2010; Clarke et al., 2010; Rundle et al., 2011). The validity and reliability of using Google Street View to supplant more traditional time- and resource-intensive observations of neighborhood conditions have been vetted in the research with generally positive results (Griew et al., 2013; Marco et al., 2017; Odgers et al., 2012; Wu et al. 2014).

Nesoff et al. (2018) used Google Street View to develop a measure of pedestrian safety as an alternative to maintaining a database of city infrastructure. Google Street View has been an integral component of several applications developed to assess streetscapes, including the Virtual Systematic Tool for Evaluating Pedestrian Streetscapes (Virtual-STEPS), a Google Street View-based auditing tool specifically designed to remotely assess microscale characteristics of the built environment (Steinmetz-Wood et al., 2019).

Related Work on Detecting Abandoned Houses: Images

The use of images to monitor street-level conditions has received attention from researchers applying a variety of techniques to compile images. Cannuscio et al. (2009) collaborated with residents to compile photographs from "outsiders" (i.e., staff photographers) along with "insiders" (or residents) and a combination of the two teams. This approach provided insights into health-

related conditions that benefited from a form of coproduction with residents. Wilson et al. (2012) collected health-related data on residents and built-environment conditions using a hybrid approach blending traditional field audits with omnidirectional imagery. The researchers concluded that the use of image-based audits could reliably supplant manual data collection processes. Researchers have also sought out low-cost alternatives for compiling images in conducting research in markets where Google Street View is not a viable option and have reported favorable results (de Souza-Daw et al., 2015).

Huang et al. (2014) provide an overview of image classification techniques that are relied on in computer vision and pattern recognition. The favorable results of image-based research into spatial conditions have fostered an increase in techniques for image classification and object detection (Liu et al., 2020; Sharma and Mir, 2020; Zhang et al., 2018). Researchers have combined the use of machine learning algorithms and Google Street View to measure walkability (Hu et al., 2020; Yin and Wang, 2016).

Empirical Studies on Predicting Abandoned Houses

Despite the challenges researchers have faced in their attempts to unambiguously identify and measure housing abandonment, some efforts to predict housing abandonment emerged in the early 1990s. For example, Anas and Arnott (1993) presented a stationary-state model which isolated land and improvements and treated both as investment assets in which improvements could be converted or rehabilitated, land could be improved, or improvements could be demolished. They also explored the relationships among construction costs, demolition costs, land costs, and maintenance costs, which all can affect investor decisions and help predict the outcomes of individual decisionmakers who rely on economics in approaching demolition and development decisions.

Hillier et al. (2003) used the Philadelphia Neighborhood Information System data to predict which residential properties are most likely to become abandoned. They found that vacancy, outstanding housing code violations, tax arrearages, and characteristics of nearby properties were significant predictors. That study stated that defining abandonment, integrating administrative data from multiple sources, and modeling temporal and spatial data are major challenges common to developing data to predict abandonment. Morckel (2012) predicted residential abandonment in two cities in Ohio, Columbus and Youngstown, using neighborhood-level factors including market conditions, physical neglect, socioeconomic conditions, and financial neglect. Morckel found that among these factors, market conditions, gentrification, and physical neglect significantly predicted the probability of a house being abandoned.

Method

Houses that have been abandoned or are about to be abandoned are a significant problem. This challenge has the potential to directly impact significant community issues, such as property violations and social-economic instability. Open data from a range of sources can be leveraged to robustly diagnose abandoned housing. Understanding the socioeconomic status of our neighbors and their concerns about property violations can aid in comprehending the critical issues

underlying house abandonment, raise awareness of this issue, and institute the necessary prevention and intervention strategies to address the impact of abandonment on communities.

This paper introduces an innovative framework for detecting abandoned houses. The framework presents a system of detecting abandoned houses, using deep learning by building an ensemble model that takes global and local contexts into account. Global contexts relate to all potential data that could be used in a model, whereas local contexts apply to a particular application; images of houses from any American city would speak to the global context, whereas images from a particular neighborhood in a single city would speak to the local context. The challenge of building models is to have the greatest potential scope of application (global) without significantly diminishing the applicability of the model to particular local contexts. To accomplish this goal, two image-based classification models for recognizing abandoned homes in both global and local contexts were developed using a variety of sources, including the house images from Web and Google Street View. A grid-search strategy was used to aggregate the models, therefore increasing the prediction accuracy for any input (either local or global context). While many deep-learning models can be highly accurate, they can be black-box analytical solutions where it is difficult to explain the result. A benefit of the model developed here is that it is highly explainable to stakeholders. It assesses the predictions of the ensemble approach model by explaining the best, worst, and conflicting cases with global and location variables. The training of the models consists of the following phases: (1) Data Collection and Preprocessing, (2) Deep Learning Modeling for Detecting Abandoned Houses, and (3) Model Validation and Explainability.

Data Collection and Preprocessing

Data Collection with Multiple Sources

Both image and structured data of houses in Kansas City, Missouri (KCMO) were gathered from various sources (exhibit 1).

Exhibit 1

Web and Kansas City Data Used for Abandoned House Detection (1 of 2)

Data	Size	Time	Modality	#Feature	Description
Neighborhood	246	Current	Geocode	246	KCMO's 246 Neighborhood in Geocode
311 call property violation (OpenData, 2021)	800K	2007-2021	Structured in a tabular format	23	Property Violation ID, Case ID, Status, Case Opened Date, Case Closed Date, Days Open, Violation Code, Violation Description, Ordinance Number, Ordinance Chapter, Violation Entry Date, Address, County, State, Zip Code, Latitude, Longitude, KIVA PIN, Council District, Police Patrol Area, Inspection Area, Neighborhood, Code Violation Location
Land Bank (2021)	6,256	2012-2020	Structured in a tabular format	16	Parcel Number, Property Class, Property Status, Address, City, State, Postal Code, County, Neighborhood, Council District, Sold Date, School District, Market Value Year, Market Value, Date Evaluated

Exhibit 1

Web and Kansas City Data Used for Abandoned House Detection (2 of 2)

Data	Size	Time	Modality	#Feature	Description
Google Street Views (Google, 2021)	398	2021	Image	2	Two categories (Abandoned houses vs. Occupied houses): A total of 398 house images were retrieved through Google Street View API based on (LandBank, 2021) and Property Violation Data from (OpenData KC, 2021)
Web Scraper (ours)	1,181	2022	Image	2	Two categories (Abandoned houses vs. Occupied houses): A total of 1181 house images were retrieved using Web Scraper.

API = application programming interface. KCMO = Kansas City, Missouri.
Source: Web Scraper images from *https://images.google.com/*

Neighborhoods in Kansas City

The neighborhoods used in this study were the officially recognized neighborhoods as published through the OpenData KC portal as a GeoJSON file. This type of data allows the most precise polygonal representation of the neighborhoods on a digital map. The dataset contained 246 neighborhoods with a corresponding id, name, and multi-polygonal annotations in terms of latitude and longitude coordinates. Five out of 246 neighborhoods had no name.

Web Scraping Image Data

Images of occupied homes and abandoned houses were scraped from Google Images, disregarding contextual information. These images were preprocessed manually and resized automatically. Exhibit 2 illustrates an example of abandoned and occupied houses collected by web scraping. These images were used to create a base model.

Google Street View Image Data

This dataset consists of abandoned and occupied house images in Kansas City, Missouri; 311 Call property violation data (City of Kansas City, MO, 2021b) and Land Bank data (City of Kansas City, MO, 2021a) from OpenData KC were used to identify addresses of abandoned and occupied houses. First, the Land Bank data contain data on abandoned homes in Kansas City, MO, including their current property status. A list of abandoned house addresses was compiled based upon this dataset. Second, the list of the occupied houses was compiled to establish a balance between abandoned and occupied properties in terms of socioeconomic situations and the degree of property violations. As a result, the majority of abandoned and occupied house images were acquired from comparable neighborhoods in Kansas City. Third, images of abandoned and occupied houses on these two lists were scraped from Google Street View. Diligent preprocessing is necessary for the accuracy of any subsequent analysis.

Exhibit 2

Google Images and Google Street View Images of Abandoned and Occupied Houses

Preprocessing

In a previous work, the authors designed an algorithm to merge the data from various sources such as 311 calls, property violations, and Land Bank data.[2] The essential factor in integrating data was to assign each data point to its own neighborhood using a geocode (coordinate pairs of latitude and longitude) given by the address of each data point (e.g., 311 request). Between 2007 and 2020, about 800,000 311 calls related to property violations were recorded in addition to data gathered from the Kansas City, Missouri Land Bank. Python was used to build the data retrieval system (version 3) using the Google Street View API. Exhibit 2 illustrates an example of abandoned and occupied houses. A total of 398 house images of 190 abandoned houses (more than one house could be associated with the 140 addresses from the Land Bank), and 208 occupied houses (from 311 property violation data) were retrieved from Google Street View.

Image-Based Abandoned House Detection

For abandoned house detection, transfer learning is applied in a three-stage ensemble approach used in a study by Anguelov et al. (2010) to identify abandonment. Transfer learning can construct domain-specific models with a comparatively small set of image data, which, when linked through the incremental learning and the ensemble approach, can optimize the prediction based upon multiple knowledge sources (Tan et al., 2018). The methodology overcomes the limitations of relying on non-geo-specific Google images and older geo-specific images from Google Street View (exhibit 2). A three-phase incremental learning algorithm was used, in which three models were created by transfer learning by adding iteratively different image datasets: **Model 1:** a global perspective based upon web images; **Model 2:** a local perspective based upon Google Street View of Kansas City, Missouri houses recorded in Land Bank records and 311 calls for suspected property infractions; and **Model 3:** an ensemble model of Model 1 and Model 2 with grid search.

The following three types of datasets were created to develop deep learning models. First, a total of 1,181 house images were acquired for this study (561 for abandoned houses and 620 for occupied houses) with Web scraping techniques using the keywords "abandoned houses," "empty houses," and "occupied." The images have been preserved in preparing for the deep learning network to be

---

[2] Reference is withheld to protect blind review process.

used to generate the initial model (Model 1). Second, a total of 398 house images were collected from Kansas City neighborhoods. This dataset was used to build out a specific local context compared with the global context using Web data. As discussed previously, 190 abandoned house images and 208 occupied house images were retrieved via Google Street View's API. Similar to the first step, these images were also preprocessed for training the local context model (Model 2). Third, a total of 160 images from both global and local domains (80 images each) were utilized as the ensemble model's testing dataset.

Model 1: Instead of training an image model from scratch, Model 1 was created with the ResNet-50 model. The ResNet-50 model is a pre-trained powerful image model, and the Web images of houses were used to tune the ResNet-50 model for the classification task of abandoned and occupied houses. The ResNet-50 model consists of five stages of 50-layers applying Convolutional Neural Networks (ConvNet/CNN). Each convolution block has three convolution layers, and each identity block has three convolution layers, trained on more than one million images of ImageNet (as shown in exhibit 3). CNN can learn without having separate filtering or featuring engineering steps constructing deep neural networks layer-by-layer. To tune the ResNet-50 model, a classification layer was added that specifically detects image features of abandoned versus occupied houses.

Exhibit 3

Architecture of CNN Model for Abandoned House Detection

Source: Image generated by the authors based on the ResNet-50 architecture diagram from He et al., 2016

The web image dataset was passed to the training process. As more training data would improve the accuracy of the classification model, data augmentation techniques were employed in this work to improve training generalization and reduce overfitting. Each image was rotated, zoomed, and flipped to create more training data points. Specially, the rotation range was set to 20 degrees, and the zoom range to 50 percent. The image was flipped horizontally, and the height and width shift ranges were set to 30 degrees. As a result, data augmentation increased the number of image variations available during each of the training epochs. One image became 16 images, equal to the batch size passed to the model. Thus, Model 1 was a general model that could classify whether a Web image of a house was abandoned or not.

Model 2: Model 2 aimed at gaining additional knowledge based upon new data and the general knowledge of abandoned houses from Model 1. The incremental knowledge was obtained by fine-tuning Model 1 using the 398 house images from the local community via Google Street View. The second model distinguishes itself from the first in that it incorporates real-world data in order to account for contextual viewpoints. Through leveraging Model 1, subsequent models are able to train much more efficiently. According to the learning curve that is a plot of model learning

performance over experience or time, the effectiveness of training is defined by the number of epochs required to complete the training. While the images from Google Street View are of lower quality than the online data, this rapid learning performance with acceptable accuracy demonstrates the effectiveness of incorporating Model 1's analysis.

Model 3: This model is a combination of Models 1 and 2. Model 1 is a global perspective on abandoned and occupied houses, and Model 2 is based on local cases of abandoned and occupied properties in Kansas City, Missouri, using local data sources such as Kansas City Land Bank data and 311 calls associated with supposed property violation data from the OpenData KC. The ensemble model was created primarily by integrating two image-based classification models utilizing the grid search method developed in a study by Qiu et al. (2014) on a weighted average basis. Grid search is a simple yet exhaustive approach for obtaining the ensemble members' weights. Grid search enables models with varying degrees of confidence or predicted performance to contribute to a prediction. Depending upon the weight, models with better performance could contribute more to the overall effort, and models with lower performance could contribute less. As a result, the ensemble model often outperforms a single model. The contribution of each model is weighted by a coefficient indicating its predicted performance. Weight values are between 0 to 100 percent, with the total weights of the ensemble models equal to one. This ensemble approach is crucial because it enables the most accurate prediction of abandoned homes by taking both global and local perspectives on vacant properties into account.

Model 4: This model was constructed as a check when compared to the proposed model (Model 3). Model 4 is a unified network model that was built with data from the Web (global) and Google Street View domains (local). In contrast to Model 3, this model was developed using the combined data. Comparing model-based fusion with data-based fusion may be an interesting point of comparison for AI and deep learning researchers.

Model Validation and Explainability

To overcome the explainability challenges associated with backbox machine learning and deep learning, one must first understand why and how prediction models work effectively or badly. Without an explanation or justification for the prediction models' decisions, prediction is not compelling. It is critical to explicate misclassified or disputed conditions in particular. Local Interpretable Model-Agnostic Explanations (LIME) developed by Ribeiro, Singh, and Guestrin (2016), was used, a well-known technique of incorporating interpretable data representations into prediction models to address black-box machine learning. Image classification evidence can be represented as a collection of super-pixels and converted to a binary on/off vector indicating which super-pixels remain visible while others are obscured. However, adoption of LIME is limited due to its presentation as a one-size-fits-all explainability tool with limited customization options. Creating a customized surrogate explainer for a specific task can significantly improve the quality of explanations produced and allow users to make informed choices, justifying the use of the instance segmentation model in conjunction with the classification model.

Results

All the datasets were randomly divided into three types of data (Training/Validation/Testing) in the ratio of 80 percent, 10 percent, 10 percent, and 65 percent, 15 percent, 20 percent for Model 1 and Model 2, respectively. A more aggressive strategy was used, which means allocating more training data (10 percent) for the global model compared to testing data (20 percent) for the local model in terms of data split since the local context could be subjective compared to the global context. Exhibit 4 summarizes the dataset's properties.

Exhibit 4

Breakdown of Image and Structured Dataset

Task/Data	Training			Validation			Testing			Total
	AB	OC	Total	AB	OC	Total	AB	OC	Total	
Web Image Data	449	496	945	56	62	118	56	62	118	1181
GSV Image Data	121	133	254	31	33	64	38	42	80	398
Ensemble Inferencing	N/A			N/A			80	80	160	160

AB = abandoned houses. GSV = Google Street View. OC = occupied houses.

Sources: Web Image Data from https://images.google.com/; GSV Image Data from https://developers.google.com/maps/documentation/streetview/overview

As mentioned previously, Model 1 is the global model, Model 2 is the local model (KCMO), Model 3 is an ensemble model with Model 1 and Model 2 (where Model 1 contributes 0.55 and model 2 contributes 0.45), and Model 4 is the single model with mixed data of global and local house images). In this paper, the authors propose Model 3 based on Model 1 and Model 2 while comparing it with Model 4. The results of testing the four unique models (Model 1, Model 2, Model 3, Model 4) are shown in exhibit 5 and exhibit 6. The confusion matrix categorizes the output into four distinct groups based on a comparison of the actual output to the ground truth: the upper left corner is True Positives (TP); the upper right corner is False Negative (FN); the lower left is False Positive (FP); and the lower right is True Negative (TN).

Accuracy—Accuracy is the most intuitive performance metric, consisting of the ratio of properly predicted observations to total observations.

$$Accuracy = \frac{TP + TN}{TP + FP + FN + TN}$$

Precision—The precision of a prediction is defined as the ratio of accurately predicted positive observations to all anticipated positive observations.

$$Precision = \frac{TP}{TP + FP}$$

Recall (Sensitivity)—The recall ratio is the number of accurately predicted positive observations divided by the total number of observations in the actual class: abandoned house.

$$Recall = \frac{TP}{TP + FP}$$

F1 score—The F1 Score is calculated by averaging Precision and Recall. As a result, this score accounts for both false positives and negatives. While the F1 score is not as intuitive as accuracy, it is frequently more helpful than accuracy, especially when the class distribution is unequal.

$$F1\ score = \frac{2 \times Recall \times Precision}{Recall + Precision}$$

Exhibit 5

Abandoned House Prediction Performance

|	Model 1 (%)	Model 2 (%)	Model 3 (%)	Model 4 (%)
Accuracy	68	88	**91**	86
Precision	91	91	**96**	85
Recall	80	84	**85**	88
F1 Score	85	87	**90**	86

Source: Data from experiments and evaluation by the authors

As shown in exhibit 5, the proposed model, Model 3, is superior to the other three models. The performance of Model 1 is demonstrably inferior to that of Model 3 or Model 4. Model 2 was created using more realistic house images as opposed to the house images of Model 1 found online that were devoid of context. Understandably, Model 2 outperforms Model 1, with an accuracy of 88 percent and an F1-score of 87 percent, compared to 68 percent and an F1-score of 85 percent, respectively. Additionally, Model 2 is superior to Model 4, which was constructed using combined images from both global and local settings. The most intriguing discovery is the performance difference between Model 3 and Model 4: Model 3 achieved 91 percent accuracy and a 90 percent F-score, whereas Model 4 achieved 86 percent and 86 percent, respectively. This demonstrates an increase of more than 5 percent, suggesting an ensemble model that integrates knowledge from separate models trained on two distinct homogeneous datasets (global and local datasets) is more successful than a single model trained on the mixed data.

Exhibit 6

Confusion Matrix

Model 1: Global Abandoned House Classification; Model 2: Local Abandoned House Classification;
Model 3: Ensemble Abandoned House Classification; Model 4: Fusion Abandoned House Classification

Model 1: Testing Results (160 images)

	Abandoned	Occupied
Abandoned	74 / 46.25%	6 / 3.75%
Occupied	16 / 10.00%	64 / 40.00%

Predicted Label
Accuracy = 0.863
Precision = 0.914
Recall = 0.800
F1 Score = 0.853

Model 2: Testing Results (160 images)

	Abandoned	Occupied
Abandoned	73 / 45.62%	7 / 4.38%
Occupied	13 / 8.12%	67 / 41.88%

Predicted Label
Accuracy = 0.875
Precision = 0.905
Recall = 0.838
F1 Score = 0.870

**Weighted Ensemble with Grid Search
(Model 1: 0.55, Model 2: 0.45)**

	Abandoned	Occupied
Abandoned	77 / 48.12%	3 / 1.88%
Occupied	12 / 7.50%	68 / 42.50%

Predicted Label
Accuracy = 0.906
Precision = 0.958
Recall = 0.850
F1 Score = 0.901

Model 4: Testing Results (160 images)

	Abandoned	Occupied
Abandoned	68 / 42.50%	12 / 7.50%
Occupied	10 / 6.25%	70 / 43.75%

Predicted Label
Accuracy = 0.863
Precision = 0.854
Recall = 0.875
F1 Score = 0.864

Source: Generated by the authors

The best cases in exhibit 7 feature three images in a series of an abandoned house, an occupied house, and an abandoned house with a 100 percent confidence score, which indicates that the images are accurate to the ground truth. These are the cases in which the three models (Models 1, 2, and 3) produce the correct prediction with the greatest confidence score (all are 100 [percent] since the house's look and location fit a logical pattern (cues such as boarded windows and peeling paint). The LIME visualization highlights several essential positive aspects of these properties (highlighted pixels) suggest they have been abandoned, including their entrances, roofs, pillars, front yards, and surrounding regions.

Exhibit 7

Best Cases: Model 1, 2, and 3 all Return the Correct Predictions

Filename: 2805_SPRUCE_AVE_facing=90.jpg.Abandoned.jpg
Model 1: Abandoned, Score = 1.0
Model 2: Abandoned, Score = 1.0
Model 3: Abandoned, Score = 1.0

Filename: Occupied.348.jpg
Model 1: Occupied, Score = 1.0
Model 2: Occupied, Score = 1.0
Model 3: Occupied, Score = 1.0

Filename: Abandoned.58.jpg
Model 1: Abandoned, Score = 1.0
Model 2: Abandoned, Score = 1.0
Model 3: Abandoned, Score = 1.0

Sources: Input images from Google Street Views; explainability images generated by the authors

The worst situations depicted in exhibit 8 are forecasts that are inaccurate for all three models. Three occupied houses are used as ground truth. For all three houses, Model 1 (global context) has a very high confidence score range (0.96–0.99), but Model 2 (local context) has a broad confidence score range of 0.96, 0.65, and 0.81. The inaccurate predictions often have a lower confidence score range (0.65–0.96) since their appearances are harder than those of the other cases, which have a more recent and clean condition. Additionally, as with the top row, the image is incomplete, which means that models are unable to collect valuable visual signals, resulting in erroneous scores and predictions. Also, the house is partially obscured by trees (as shown in the second row of the same

image), creating confusion in the three models. Finally, the explainability demonstrated by LIME visualization revealed numerous critical positive characteristics of these properties (highlighted pixels) that indicate they have been occupied, such as their trees, bushes, and neighboring areas, resulting in a misleading categorization of these properties in comparison to occupied houses.

Exhibit 8

Worst cases: Model 1, 2, and 3 all Return the Incorrect Predictions

Filename: 2611_E_27th_St_facing=180.jpg.Occupied.jpg
Model 1: Abandoned, Score = 0.993
Model 2: Abandoned, Score = 0.964
Model 3: Abandoned, Score = 0.98

Filename: 3801_E_12th_Ter_facing=60.jpg.Occupied.jpg
Model 1: Abandoned, Score = 0.999
Model 2: Abandoned, Score = 0.657
**Model 3: Abandoned, Score = 0.845**

Filename: 4404_Independence_Ave_facing=60.jpg.Occupied.jpg
Model 1: Abandoned, Score = 0.961
Model 2: Abandoned, Score = 0.814
Model 3: Abandoned, Score = 0.895

*Sources: Input images from Google Street Views; explainability images generated by the authors*

There are many cases in which Model 1 and Model 2 conflict with each other. Conflicting predictions cause the ensemble model to make incorrect predictions, as it is difficult to make a judgment based on disputed observations. The conflict situations depicted in exhibit 9 are

examples of Model 1 and Model 2 conflicting with each other. Model 2 (local context) correctly identified the ground truth for all three houses, while Model 1 (global context) incorrectly identified with relatively low confidence at around 0.7. Some of the images are incomplete, which means that models are unable to collect valuable visual signals, resulting in erroneous scores and predictions. With LIME, we can clearly tell that Model 2 has complete control over which portion of the image is crucial to its decision (highlighted pixels). If the owner of the house forgot to trim the lawn (as in the top row), the house is more likely to have been neglected for a while.

Exhibit 9

Conflicting Cases: Model 1 Returned an Incorrect Prediction, but Model 3 Corrected It


Filename: 1416_DENVER_AVE_facing=90.jpg.Abandoned.jpg
Model 1: Occupied, Score = 0.796
Model 2: Abandoned, Score = 0.976
Model 3: Abandoned, Score = 0.552

Filename: 3435_HOLMES_ST_facing=90.jpg.Occupied.jpg
Model 1: Abandoned, Score = 0.702
Model 2: Occupied, Score = 0.879
Model 3: Occupied, Score = 0.559

Filename: 3213_E_59th_St_facing=180.jpg.Occupied.jpg
Model 1: Abandoned, Score = 0.709
Model 2: Occupied, Score = 0.896
Model 3: Occupied, Score = 0.563

Sources: Input images from Google Street Views; explainability images generated by the authors

The conflicts can happen on the other side too, however, where Model 2 fails to classify correctly, as in exhibit 10. The top row indicates that the model was able to capture the trash in front of the house (highlighted pixels for abandoned in this case), but the positive area is dominant instead; it has given too much sensitivity to the roof and door instead of the surroundings—hence, a low confidence score (0.547) and a misclassification. In the remaining two cases, Model 2 tends to be confused by both the layout, scenery, and coloration of the images. These oddly lit scenes do not exist in a typical Google Street View image. Therefore, despite having a medium to high confidence score, Model 2 was not fully prepared for these.

Exhibit 10

Conflicting Cases: Model 2 Returned an Incorrect Prediction, but Model 3 Corrected It

Filename: 3807_E_68th_St_facing=180.jpg.Abandoned.jpg
Model 1: Abandoned, Score = 0.988
Model 2: Occupied, Score = 0.547
Model 3: Abandoned, Score = 0.747

Filename: Abandoned.544.jpg
Model 1: Abandoned, Score = 1.0
Model 2: Occupied, Score = 0.991
Model 3: Abandoned, Score = 0.554

Filename: Abandoned.13.jpg
Model 1: Abandoned, Score = 1.0
Model 2: Occupied, Score = 0.781
Model 3: Abandoned, Score = 0.648

Sources: Input images from Google Street Views; explainability images generated by the authors

Conclusion

The housing abandonment crisis that has contributed to neighborhood decline and created a financial drain on metropolitan areas must be addressed. One key obstacle has been the inability to identify individual abandoned houses and track them over the cycle of abandonment. While some progress has been made by using secondary data, fundamental challenges exist for the time- and cost-effective collection of primary data. Traditionally, housing surveys conducted by city officials physically walking the streets and cataloging housing conditions have faced high labor costs that render them unaffordable for many cities, particularly when considering the rapid change in housing conditions and the need to maintain such data over time. Municipal officials also face the challenge of convincing elected leaders and citizens to allocate resources toward data collection, something that does not yield immediate tangible results, as opposed to other tangible policies such as public safety or solid waste collection. This paper presents a methodology that uses technological innovations of image analysis to fill the void by substituting technological innovations to complement and extend manual efforts in a more cost-effective manner, especially when conducted at scale and on a continuous basis.

The image analysis method presented in this study has two paths forward for adoption by municipalities that can help reduce incremental costs. As noted previously, the ability to detect abandoned houses earlier in the process, enabled by the use of real-time images collected by low-cost cameras mounted on private vehicles and driven through the sample neighborhood, could be leveraged by blending the data with imagery from Google Street View. Furthermore, the recent drop in the price of cameras and image sensors opens the field for the widespread use of the technology. In a real-world application, the images could be collected at a relatively low cost by mounting cameras on municipal fleet vehicles. Trash vehicles, for example, that drive fixed routes every week, could provide images where subtle changes in housing conditions could lead to preventative maintenance even before the structure is abandoned. This level of image collection could help develop valid and reliable predictive models of abandonment, allowing cities to develop more effective intervention programs. Furthermore, the methodology presented herein is scalable and could be applied in other neighborhoods and communities.

In addition to collecting images through primary efforts, municipalities could also leverage a coproduction process wherein residents could be empowered to take images of houses that may need immediate attention. This could reduce costs and increase community engagement by encouraging resident participation to develop low-cost crowdsourcing of images that could complement municipal efforts. Having "eyes on the streets" would provide a real-time set of images that could be used to help train AI models to detect abandoned housing earlier in the process of abandonment. However, it should be recognized that overreliance on this method by municipalities has the potential for inequitable allocation of resources due to differences in the capacity and willingness of residents to engage in capturing images and submitting them for review, especially in blighted neighborhoods where the abandonment problem may be more ingrained. At the same time, the process could also serve to empower residents who have historically felt disengaged with their governments. Care needs to be placed in creating a process that empowers residents to collect data yet does not allow the process to be weaponized by disaffected neighbors or other actors.

In addition to data collection, the approach presented in this paper also takes into consideration the deployment of the models that can contribute to the success of city and land bank staff members in their missions to serve the public interest in an equitable, timely, and resource-efficient manner. Transfer learning could be used to mitigate the cost of development and implementation of the proposed method. That is, the approaches, methodology, and models presented in this study created a set of publicly available pretrained models. By building on these more general image models that detect abandonment in this case study, municipalities could build locally valid models using a smaller amount of locally collected data. This approach would reduce one of the biggest costs encountered in this study: the collection of up-to-date local data. To encourage the adoption of the methodology presented in this study, an open-source project of the proposed work will be published in the near future that will contain source-code, data, models, tools, and apps. Municipalities looking to leverage this method will be able to upload housing images and tune the model for their particular local contexts; this open-source tool will allow for continual updating and tuning of the deep learning model in order to generate more accurate predictions for future users. The goal is to guide potential users through the end-to-end process, which includes (1) gathering data relevant to their own context, (2) sharing their own data, (3) building deep learning models, and (4) distributing as a web application to a wider audience. This will reduce the cost and time associated with managing and scaling up the deployment of these models to the level of serving a whole community or city.

The proposed approach is not without limitations. This study focused on the front views of houses utilizing land bank data and 311 calls to develop methods for detecting abandoned houses based on house images. According to the literature review, markers of the pedestrian environment, such as the design of street crossings and the quality of sidewalks, can also help detect abandoned houses. In future research, the success of detection methods focused on street-front house images presented in this study can be enhanced by incorporating more images on the houses (e.g., backyards), as well as information on conditions in the surrounding environment that may be crowdsourced from local community members. It should also be noted that another limitation of this study was the reliance on housing images collected during the summer and fall. Seasonal differences will be explored in future research, incorporating thermal cameras and observing other variables (e.g., snow removal) that provide additional insights into abandonment. Since abandonment is a process that spans seasons, adapting the model to seasonal differences will help improve the reliability of the methodology. In making this extension, the methodology will correct for the influence of seasonal factors, such as vegetation or snow covering camera views, that impair the effectiveness of detection for the proposed models. In addition to seasonal recalibration, future research will explore the contributions and challenges imposed in the development of time-aware models (i.e., day and night), as well as weather-aware models (i.e., sunny versus rainy days). Leveraging the open-source tools available to municipal officials will allow the team to efficiently introduce images that contain seasonal differences.

Finally, it is important to recognize some caveats with respect to the use of multiple sources of data presented in this study. For example, efforts such as those used in this study must recognize the importance of privacy. Since 2020, the Census Bureau has used a technology known as differential privacy (DP) to protect individuals' privacy. A previous study by the authors of this study (Ho et

al., 2022) incorporated the DP technique for 311 call data and census data entirely in order to determine the optimal level of privacy for protecting neighbors' information. The intention is to include privacy-preserving abandoned home identification into future work while retaining a high detection rate for abandoned houses.

Additionally, public domain data, such as criminal or census bureau data, present a variety of challenges, including missing values, invalidated values, inconsistent values, and categories generated by a number of inconsistent or duplicate entries. Ho et al. (2022) addressed these data quality challenges when combining data from many sources for Kansas City neighborhood data. A future study will also address issues in image domains with inconsistencies or low image quality.

Authors

Jim DeLisle, Ph.D., is an associate professor at the Bloch School of Management, University of Missouri—Kansas City, delislej@umkc.edu. Hye-Sung Han, Ph.D., is an assistant professor at the School of Public Affairs and Administration, University of Kansas, hshan@ku.edu. Duy H. Ho is a doctoral candidate at the School of Computing and Engineering, University of Missouri—Kansas City, dhh3hb@mail.umkc.edu. Yugyung Lee, Ph.D., is a professor at the School of Computing and Engineering, University of Missouri—Kansas City, leeyu@umkc.edu. Brent Never, Ph.D., is an associate professor at the Bloch School of Management, University of Missouri—Kansas City, neverb@umkc.edu. Ye Wang, Ph.D., is an associate professor at the College of Arts and Sciences, University of Missouri—Kansas City, WanYe@umkc.edu.

References

Alexander, Frank. S. 2005. Land Bank Authorities: A Guide for the Creation and Operation of Local Land Banks. New York: Local Initiatives Support Corporation.

Anas, Alex, and Richard J. Arnott. 1993. "Technological Progress in a Model of the Housing-Land Cycle," *Journal of Urban Economics* 34 (2): 186.

Anguelov, Dragomir, Carole Dulong, Daniel Filip, Christian Frueh, Stéphane Lafon, Richard Lyon, Abhijit Ogale, Luc Vincent, and Josh Weaver. 2010. "Google Street View: Capturing the World at Street Level," *Computer* 43 (6): 32–38.

Badland, H.M., S. Opit, K. Witten, R.A. Kearns, and S. Mavoa. 2010. "Can Virtual Streetscape Audits Reliably Replace Physical Streetscape Audits?" *Journal of Urban Health: Bulletin of the New York Academy of Medicine* 87 (6): 1007–1016. https://doi.org/10.1007/s11524-010-9505-x.

Bak, Xian F., and Geoffrey J.D. Hewings. 2017. "Measuring Foreclosure Impact Mitigation: Evidence from the Neighborhood Stabilization Program in Chicago," *Regional Science and Urban Economics* 63: 38–56.

Bell, E.J., and D. Kelso. 1986. "The Demolition of Downtown Low-Income Residential Buildings: A Discriminant Analysis," *Socio-Economic Planning Sciences* 17.

Bieretz, B., and Joseph Schilling. 2019. *Pay for Success and Blighted Properties: Insights and Opportunities for Funding Vacant Property Reclamation and Neighborhood Stabilization*. Washington, DC: Urban Institute.

Bloch, S. 2020. "An On-the Ground Challenge to Uses of Spatial Big Data in Assessing Neighborhood Character," *Geographical Review* 110 (1–2): 210–223.

Cain, K.L., J.F. Sallis, T.L. Conway, K.A. Gavand, C.M. Geremia, R.A. Millstein, L.D. Frank, B.E. Saelens, J. Chapman, M.A. Adams, K. Glanz, and A.C. King. 2014. "Contribution of Streetscape Audits to Explanation of Physical Activity in Four Age Groups Based on the Microscale Audit of Pedestrian Streetscapes (MAPS)," *Social Science and Medicine* 116: 82–92.

Cannuscio, C.C., E.E. Weiss, H. Fruchtman, J. Schroeder, J. Weiner, and D.A. Asch. 2009. "Visual Epidemiology: Photographs as Tools for Probing Street-Level Etiologies," *Social Science & Medicine* 69 (4): 553–564.

Cannon, D.S., M.L. Lachman and A.S. Bernhard. 1977. "Identifying Neighborhoods for Preservation and Renewal," *Growth and Change* 8 (1): 35.

City of Kansas City, Missouri. n.d.a. *Dangerous Buildings and Demolitions*. https://www.kcmo.gov/city-hall/departments/neighborhoods-housing-services/neighborhood-preservation/dangerous-buildings-and-demolitions.

———. n.d.b. Vacant Property Registration. Retrieved 2021, from https://www.kcmo.gov/city-hall/departments/neighborhoods-housing-services/neighborhood-preservation/vacant-property-registration.

———. 2021a. Land Bank Data. Kansas City, Missouri. Retrieved from https://data.kcmo.org/Neighborhoods/Land-Bank-Data/j7gh-5trq.

———. 2021b. Property Violations. Kansas City, Missouri. Retrieved from https://data.kcmo.org/Neighborhoods/Property-Violations-Historical-/nhtf-e75a.

Clarke, P., J. Ailshire, R. Melendez, M. Bader, and J. Morenoff. 2010. "Using Google Earth to Conduct a Neighborhood Audit: Reliability of a Virtual Audit Instrument," *Health Place* 16 (6): 1224–1229.

Cleland, C.L., S. Ferguson, F. Kee, P. Kelly, A.J. Williams, G. Nightingale, A. Cope, C. Foster, K. Milton, M.P. Kelly, R. Jepson, and R.F. Hunter. 2021. "Adaptation and Testing of a Microscale Audit Tool to Assess Liveability Using Google Street Vew: MAPS-Liveability," *Journal of Transport & Health* 22.

Decker, A. 2018. "Community Land Trusts and State Legislation: A Model Act to Enable This Affordable Housing Tool," *Journal of Affordable Housing and Community Development Law* 26 (3): 489–530.

de Souza-Daw, T., R. Ross, T.D. Nhan, L.A. Hung, N.D.Q. Trung, L.H. Chau, H.M. Phuong, L.H. Ngoc, and M. Nkhoma. 2015. "Design and Evaluation of a Low-Cost Street-Level Image Capturing Vehicle for South-East Asia," *Journal of Engineering, Design and Technology* 13 (4): 579–595.

Foster, Kirk A., and J. Aaron Hipp. 2011. "Defining Neighborhood Boundaries for Social Measurement: Advancing Social Work Research," *Social Work Research* 35 (1): 25–35.

Fox, E.H., J.E. Chapman, A.M. Moland, N.E. Alfonsin, L.D. Frank, J.F. Sallis, T.L. Conway, K.L. Cain, C. Geremia, E. Cerin, G. Vanwolleghem, D. Van Dyck, A. Queralt, J. Molina-García, A.A.F. Hino, A.A.D.S. Lopes, J. Salmon, A. Timperio, and S.E. Kershaw. 2021. "International Evaluation of the Microscale Audit of Pedestrian Streetscapes (MAPS) Global Instrument: Comparative Assessment Between Local and Remote Online Observers," *International Journal of Behavioral Nutrition and Physical Activity* 18 (1): 1–15.

Fraser, James C., and Deirdre Oakley. 2015. "The Neighborhood Stabilization Program: Stable For Whom?" *Journal of Urban Affairs* 37 (1): 38–41.

Fujii, Y. 2016. "Putting the Pieces Together: How Collaboration Between Land Banks and Community Land Trusts Can Promote Affordable Housing in Distressed Neighborhoods," *Cities* 56: 1–8.

Goldstein, J., M. Jensen, and E. Reiskin. 2001. Urban Vacant Land Redevelopment: Challenges and Progress. Working paper 37. Lincoln Institute of Land Policy.

Griew, P., M. Hillsdon, C. Foster, E. Coombes, A. Jones, and P. Wilkinson. 2013. "Developing and Testing a Street Audit Tool Using Google Street View to Measure Environmental Supportiveness for Physical Activity," *International Journal of Behavioral Nutrition and Physical Activity* 10: 103.

Grinstein-Weiss, Michal, Yeong Hun Yeo, Kim R. Manturuk, Mathieu R. Despard, Krista A. Holub, Johanna K. P. Greeson, and Roberto G. Quercia. 2013. "Social Capital and Homeownership in Low- To Moderate-Income Neighborhoods," *Social Work Research* 37 (1): 37–53.

Harvey, C., and L. Aultman-Hall. 2016. "Measuring Urban Streetscapes for Livability: A Review of Approaches," *Professional Geographer* 68 (1): 149–158.

He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. "Deep Residual Learning for Image Recognition," In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770-778. https://doi.org/10.1109/CVPR.2016.90.

Heins, P., and T. Abdelazim. 2014. *Take it to the Bank: How Land Banks Are Strengthening America's Neighborhoods*. Flint, MI: Center for Community Progress.

Hillier, A.E., D.P. Culhane, T.E. Smith, and C.D. Tomlin. 2003. "Predicting Housing Abandonment with the Philadelphia Neighborhood Information System," *Journal of Urban Affairs*, 25 (1): 91–105.

Ho, Duy, Srichakradhar Nagireddy, Charan Thota, Yugyung Lee, Ye Wang, and Brent Never. 2022. "Open Community Platform for Data Integration and Privacy Preserving for 311 Calls," *Sustainable Cities and Society* (forthcoming).

Hu, C.-B., F. Zhang, F.-Y. Gong, C. Ratti, and X. Li 2020. "Classification and Mapping of Urban Canyon Geometry Using Google Street View Images and Deep Multitask Learning," *Building and Environment* 167.

Huang, Y., Z. Wu, L. Wang, and T. Tan. 2014. "Feature Coding in Image Classification: A Comprehensive Study," *IEEE Transactions on Pattern Analysis & Machine Intelligence* 36 (3): 493–506.

Immergluck, Dan, and Geoff Smith. 2006. "The Impact of Single-family Mortgage Foreclosures on Neighborhood Crime," *Housing Studies* 21 (6): 851–866.

Keating, W.D. 2010. "Redevelopment of Vacant Land in the Blighted Neighbourhoods of Cleveland, Ohio, Resulting from the Housing Foreclosure Crisis," *Journal of Urban Regeneration and Renewal* 4 (1): 39–52.

Leonard, T., N. Jha, and L. Zhang. 2017. "Neighborhood Price Externalities of Foreclosure Rehabilitation: An Examination of the Neighborhood Stabilization Program," *Empirical Economics* 52 (3): 955–975.

Liu, L., W. Ouyang, X. Wang, P. Fieguth, J. Chen, X. Liu, and M. Pietikäinen. 2020. "Deep Learning for Generic Object Detection: A Survey," *International Journal of Computer Vision* 128 (2): 261–318.

Mallach, A. 2006. *Bringing Buildings Back: From Abandoned Properties to Community Assets*. Montclair, NJ: National Housing Institute.

Marco, M., E. Gracia, M. Martín-Fernández, and A. López-Quílez. 2017. "Validation of a Google Street View-Based Neighborhood Disorder Observational Scale," *Journal of Urban Health* 94 (2): 190–198.

Margulis, H.L., and C. Sheets. 1985. "Housing Rehabilitation Impacts on Neighborhood Stability in a Declining Industrial City," *Journal of Urban Affairs*: 19.

Martin, D.G., A. Hadizadeh Esfahani, O.R. Williams, R. Kruger, J. Pierce, and J. DeFilippis. 2020. "Meanings of Limited Equity Homeownership in Community Land Trusts," *Housing Studies* 35 (3): 395–414.

Mennis, J., S.L. Dayanim, and H. Grunwald. 2013. "Neighborhood Collective Efficacy and Dimensions of Diversity: A Multilevel Analysis," *Environment and Planning A* 45 (9): 2176–2193.

Millstein, R.A., K.L. Cain, J.F. Sallis, T.L. Conway, C. Geremia, L.D. Frank, J. Chapman, D. Van Dyck, L.R. Dipzinski, J. Kerr, K. Glanz, and B.E. Saelens. 2013. "Development, Scoring, and Reliability of the Microscale Audit of Pedestrian Streetscapes (MAPS)," *BMC Public Health* 13 (1): 1–15.

Mooney, S.J., M.D.M. Bader, G.S. Lovasi, J.O. Teitler, K.C. Koenen, A.E. Aiello, S. Galea, E. Goldmann, D.M. Sheehan, and A.G. Rundle. 2017. "Street Audits to Measure Neighborhood Disorder: Virtual or In-Person?" *American Journal of Epidemiology* 186 (3): 265–273.

Morckel, V. 2014. "Predicting Abandoned Housing: Does the Operational Definition of Abandonment Matter?" *Community Development* 45 (2): 121.

Morckel, V.C. 2012. *Predicting the Probability of Housing Abandonment Using Hierarchical and Spatial Models*. Columbus, OH: Ohio State University.

Nesoff, E.D., A.J. Milam, C.B. Barajas, and C.D.M. Furr-Holden. 2020. "Expanding Tools for Investigating Neighborhood Indicators of Drug Use and Violence: Validation of the NIfETy for Virtual Street Observation," *Prevention Science* 21 (2): 203–210.

Nesoff, Elizabeth D., Adam J. Milam, Keshia M. Pollack, Frank C. Curriero, Janice V. Bowie, Andrea C. Gielen, and Debra M. Furr-Holden. 2018. "Novel Methods for Environmental Assessment of Pedestrian Injury: Creation and Validation of the Inventory for Pedestrian Safety Infrastructure." *Journal of Urban Health* 95 (2): 208–221.

Odgers, C.L., A. Caspi, C.J. Bates, R.J. Sampson, and T.E. Moffitt. 2012. "Systematic Social Observation of Children's Neighborhoods Using Google Street View: A Reliable and Cost-Effective Method," *Journal of Child Psychology and Psychiatry* 53 (10): 1009–1017.

Phillips, C.B., J.K. Engelberg, C.M. Geremia, W. Zhu, J.M. Kurka, K.L. Cain, J.F. Sallis, T.L. Conway, and M.A. Adams. 2017. "Online Versus In-Person Comparison of Microscale Audit of Pedestrian Streetscapes (MAPS) Assessments: Reliability of Alternate Methods," *International Journal of Health Geographics* 16 (1): 27.

Pagano, M.A., and A.O. Bowman. 2000. *Vacant Land in Cities: An Urban Resource*. Washington, DC: Brookings Institution.

Pliakas, T., S. Hawkesworth, R.J. Silverwood, K. Nanchahal, C. Grundy, B. Armstrong, J.P. Casas, R.W. Morris, P. Wilkinson, and K. Lock. 2017. "Optimising Measurement of Health-Related Characteristics of the Built Environment: Comparing Data Collected by Foot-Based Street Audits, Virtual Street Audits and Routine Secondary Data Sources," *Health Place* 43: 75–84.

Queralt, A., J. Molina-García, M. Terrón-Pérez, E. Cerin, A. Barnett, A. Timperio, J. Veitch, R. Reis, A. Silva, A. Ghekiere, D. Van Dyck, T.L. Conway, K.L. Cain, C.M. Geremia, and J.F. Sallis. 2021. "Reliability of Streetscape Audits Comparing On-Street and Online Observations: MAPS-Global in 5 Countries." *International Journal of Health Geographics* 20 (1): 6. https://doi.org/10.1186/s12942-021-00261-5.

Qiu, X., L. Zhang, Y. Ren, P.N. Suganthan, and G. Amaratunga. 2014. "Ensemble Deep Learning for Regression and Time Series Forecasting." In *2014 IEEE Symposium on Computational Intelligence in Ensemble Learning (CIEL)* 1–6.

Ribeiro, M.T., S. Singh, and C. Guestrin. 2016. "Why Should I Trust You? Explaining the Predictions of any Classifier," In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144.

Robbins, I.S. 1935. "Program for the Demolition of Unsafe and Insanitary Buildings," *American City* 50: 45–45.

Rosenthal, E. 1953. "Housing Market Behavior in a Declining Area: Long-Term Changes in Inventory and Utilization of Housing on New York's Lower East Side," *American Sociological Review* 18 (2): 210–211.

Rundle, A.G., M.D.M. Bader, C.A. Richards, K.M. Neckerman, and J.O. Teitler. 2011. "Using Google Street View to Audit Neighborhood Environments," *American Journal of Preventive Medicine* 40 (1): 94–100.

Schilling, J.M. 2002. *The Revitalization of Vacant Properties: Where Broken Windows Meet Smart Growth*. International City/County Management Association.

Schneider, W.W. 1941. "Residential Construction and Demolition, 1936 to 1938," *Monthly Labor Review* 52: 721–729.

Schuetz, J., J. Spader, and A. Cortes. 2016. "Have Distressed Neighborhoods Recovered? Evidence from the Neighborhood Stabilization Program," *Journal of Housing Economics* 34: 30–48.

Sharma, V., and R.N. Mir. 2020. "A Comprehensive and Systematic Look Up Into Deep Learning Based Object Detection Techniques: A Review," *Computer Science Review* 38.

Sjoberg, G. 1955. "Urban Community Theory and Research: A Partial Evaluation," *American Journal of Economics and Sociology* 14 (2): 199–206.

Steinmetz-Wood, M., K. Velauthapillai, G. O'Brien, and N.A. Ross. 2019. "Assessing the Micro-Scale Environment Using Google Street View: The Virtual Systematic Tool for Evaluating Pedestrian Streetscapes (Virtual-STEPS)," *BMC Public Health* 19 (1): 1246.

Sternlieb, G., R.W. Burchell, J.W. Hughes, and F.J. James. 1974. "Housing Abandonment in the Urban Core," *Journal of the American Institute of Planners* 40 (5), 321–332.

Stevenson, M., C. Mues, and C. Bravo. 2021. "Deep Residential Representations: Using Unsupervised Learning to Unlock Elevation Data for Geo-Demographic Prediction," *arXiv* preprint arXIv:2112.01421.

Tan, C., F. Sun, T. Kong, W. Zhang, C. Yang, and C. Li. 2018. "A Survey on Deep Transfer Learning," in *International Conference on Artificial Neural Networks and Machine Learning* - ICANN 2018, Springer. 270–279.

Theall, K.P., C.N. Morrison, S.F. Jacoby, A. Tucker, M.E. Wallace, M.C. Kondo, C.C. Branas, and J. Gustat. 2021. "Neighborhood Blighted Property Removal and 311 Calls for Non-Emergency Services: A Test of a Marker of Social Control," *Geography Analysis*. doi:https://doi.org/10.1111/gean.12286.

U.S. Census Bureau. 2022. Housing Vacancies and Homeownership: Definitions and Explanations. Washington, D.C.: U.S. Census Bureau.

U.S. Department of Housing and Urban Development (HUD) Office of Policy Development and Research (PD&R). 2014. "Vacant and Abandoned Properties: Turning Liabilities Into Assets," *Evidence Matters*, Winter.

———. n.d. *HUD Aggregated USPS Administrative Data on Address Vacancies*. https://www.huduser.gov/portal/datasets/usps.html.

U.S. Government Accountability Office (GAO). 2011. Vacant Properties: Growing Number Increases Communities' Costs and Challenges. Washington, D.C.

Vanwolleghem, G., A. Ghekiere, G. Cardon, I. De Bourdeaudhuij, S. D'Haese, C.M. Geremia, M. Lenoir, J.F. Sallis, H. Verhoeven, and D. Van Dyck. 2016. "Using an Audit Tool (MAPS Global) to Assess the Characteristics of the Physical Environment Related to Walking for Transport in Youth: Reliability of Belgian Data," *International Journal of Health Geographics* 15: 1–11.

Varady, D.P. 1984. "Residential Mobility in the Urban Homesteading Demonstration Neighborhoods," *Journal of the American Planning Association* 50 (3): 346.

Wachsmuth, D. 2008. From Abandonment to Affordable Housing: Policy Options for Addressing Toronto's Abandonment Problem. Toronto: Cities Centre, University of Toronto.

Williams, S., G. Galster, and N. Verma. 2013. "Home Foreclosures as Early Warning Indicator of Neighborhood Decline," *Journal of the American Planning Association* 79 (3): 201–210.

Wilson, J.S., C.M. Kelly, M. Schootman, E.A. Baker, A. Banerjee, M. Clennin, and D.K. Miller 2012. "Assessing the Built Environment Using Omnidirectional Imagery," *American Journal of Preventive Medicine* 42 (2): 193–199.

Winter, S., L. Goldman Rosas, P. Padilla Romero, J. Sheats, M. Buman, C. Baker, and A. King. 2016. "Using Citizen Scientists to Gather, Analyze, and Disseminate Information About Neighborhood Features That Affect Active Living," *Journal of Immigrant and Minority Health* 18 (5): 1126–1138.

Whitaker, S., and T.J. Fitzpatrick. 2016. "Land Bank 2.0: An Empirical Evaluation," *Journal of Regional Science* 56 (1): 156–175.

Woodbury, C. 1937. "Procedure for Securing Loans and Grants Under U.S. Housing Act of 1937," *American City* 52: 105.

Wu, Y.-T., P. Nash, L.E. Barnes, T. Minett, F.E. Matthews, A. Jones, and C. Brayne. 2014. "Assessing Environmental Features Related to Mental Health: A Reliability Study of Visual Streetscape Images," *BMC Public Health* 14: 1094.

Yin, L. and Z. Wang. 2016. "Measuring Visual Enclosure for Street Walkability: Using Machine Learning Algorithms and Google Street View Imagery," *Applied Geography* 76: 147–153.

Zhang, W., C. Witharana, W. Li, C. Zhang, X. Li, and J. Parent. 2018. "Using Deep Learning to Identify Utility Poles with Crossarms and Estimate Their Locations from Google Street View Images," *Sensors* (14248220) 18 (8): 2484.

Zhu, W., Y. Sun, J. Kurka, C. Geremia, J.K. Engelberg, K. Cain, T. Conway, J.F. Sallis, S.P. Hooker, and M.A. Adams. 2017. "Reliability Between Online Raters with Varying Familiarities of a Region: Microscale Audit of Pedestrian Streetscapes (MAPS)," *Landscape and Urban Planning* 167: 240–248.

Zou, S., and L. Wang. 2021. "Detecting Individual Abandoned Houses from Google Street View: A Hierarchical Deep Learning Approach," *ISPRS Journal of Photogrammetry and Remote Sensing* 175: 298–310.

Zuo, Fan, Jingxing Wang, Jingqin Gao, Kaan Ozbay, Xuegang Jeff Ban, Yubin Shen, Hong Yang, and Shri Iyer. 2020. "An Interactive Data Visualization and Analytics Tool to Evaluate Mobility and Sociability Trends During COVID-19." In *Proceedings of UrbComp 2020: The 9th SIGKDD International Workshop on Urban Computing*, San Diego, CA, August 24.