



Lessons from a Human-in-the-Loop Machine Learning Approach for Identifying Vacant, Abandoned, and Deteriorated Properties in Savannah, Georgia

Xiaofan Liang¹ , Brian Brainerd², Tara Hicks²,
and Clio Andris³ 

Journal of Planning Education and Research
1–13

© The Author(s) 2024

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/0739456X241273945

journals.sagepub.com/home/jpe



Abstract

Addressing strategies for managing vacant, abandoned, and deteriorated (VAD) properties is important for maintaining healthy communities. Yet, the process of identifying these properties can be difficult. Here, we create a human-in-the-loop machine learning (HITLML) model called *VADecide* and apply it to a parcel-level case study in Savannah, Georgia. The results show a higher prediction accuracy than was achieved when using a machine learning model without human input in the training. The HITLML approach also reveals differences between machine versus human-generated results. Our findings contribute to knowledge about the advantages and challenges of HITLML in urban planning.

Keywords

housing, human-in-the-loop machine learning, spatial decision support system, civic tech, vacancy, blight

Introduction

Over time, neglected structures and land in communities degrade into unusable or unsafe infrastructure. This process adversely affects public health and welfare and can increase crime and lower property values. To protect and support growth in communities, local jurisdictions try to identify these vacant, abandoned, and deteriorated (VAD) properties. However, identifying VAD properties, that is, recording their location, is not a simple task: there is no national database nor a standardized definition with which to detect VAD properties, and as such, cities typically rely on block-by-block field surveys to count VAD properties (Mallach 2018). Often, housing planners drive around the neighborhoods to visually identify and record these properties, or they rely on calls from neighbors to the civic hotline to report problems with properties in their communities.¹ Some cities have used techniques such as spatial decision support systems (SDSSs) and machine learning (ML) models to find VAD properties (Appel et al., 2014; Hillier et al., 2003; Hillenbrand 2016; Reyes et al., 2016) but reports of these efforts emphasize the technology itself rather than the combination of human, technology, and data. Minimizing or ignoring the human aspect of these decision-making statistical models is dangerous because the models can provide suggestions that could

counter a practitioner's commonsense or what is known to be helpful for a community.

Accordingly, this research has two objectives. The first objective is to design a workflow that integrates human expertise into a ML model to efficiently identify VAD properties from a dataset of parcels with parcel attributes. The outcome of this objective is a human-in-the-loop machine learning (HITLML) model that we call *VADecide*. The second objective is to measure how *VADecide* improves or does not improve upon (1) a simple, classic ML model using existing datasets (i.e., the “ML model” in this paper) and (2) the city's current workflow based on their expert knowledge. An improved model will be more *efficient* (requiring fewer training samples and giving more predictions), more

Initial submission, July 2022; revised submissions, July 2023, March and June 2024; final acceptance, July 2024

¹University of Michigan, Ann Arbor, MI, USA

²City of Savannah Housing & Neighborhood Services Department, Savannah, GA, USA

³Georgia Institute of Technology, Atlanta, GA, USA

Corresponding Author:

Xiaofan Liang, University of Michigan, 2000 Bonisteel Blvd, Ann Arbor, MI 48109, USA.

Email: xfliang@umich.edu

accurate (yielding a higher prediction accuracy rate), more *reliable* (recalling a similar set of properties that were previously discovered using field surveys and U.S. Post Office [USPS] vacancy records), and more *robust* (capturing all features of VAD properties).

To meet these research objectives, we partnered with the City of Savannah, Georgia, in collaboration with four housing officials from Savannah's Housing and Neighborhood Services Department (HNSD) and the Chatham County/City of Savannah Land Bank Authority (LBA). The city government acquires a few dozen VAD properties per year for redevelopment as affordable housing through tax sales, declaration of public nuisance, or eminent domain. To help identify VAD properties, we built a model that inputs a set of parcels (i.e., properties) and outputs a label for each as "VAD" or "not VAD." In the model, each parcel is associated with attributes, such as code violations, tax, and crime, which were collected using census and municipal data records. To help integrate the "human factor," four housing officials labeled properties by hand in a spreadsheet. The model was trained on these three hundred training samples and returns 1,309 predicted VAD properties. We found that *VADecide* is more efficient and scalable than the classic ML model and the city's current workflow. Compared with the ML model, the improved model has a similar model prediction accuracy yet is more reliable (i.e., it has higher agreement rates) when validated with VADs identified through field surveys and USPS vacancy data. The ML model yields a high percentage of VADs with code violations, as it is trained on VADs identified by code violations alone, whereas the properties classified as VADs in *VADecide* have a wider range of "telling" variables (e.g., tax, code, crime). Compared with the city's workflow, the HITLML approach identifies more VAD candidates distributed across more neighborhoods, which helps reduce confirmation bias from city planners. The model also associated VAD properties with higher tax delinquency and crime, whereas the city's workflow tends to associate VADs with code violations, which are often manifested as visual features that housing workers see when they visit the neighborhoods.

As such, the contribution of this work is the use of a HITLML model to identify more VAD candidates in diverse neighborhoods, which can help experts create plans for identified parcels and preempt deterioration that adversely affects neighborhoods and residents. The method engages planning experts, and its development unearthed subjectivity in the labeling process (for instance, we found that in some cases, two experts labeled the same property differently). By examining the model's decision process a posteriori, the experts learned more about their process. While *VADecide*'s results may not be generalizable to all communities, other planners and researchers may benefit from the lessons learned in this case study, and from the academic-practitioner collaboration that integrates technology with professional knowledge.

Our article proceeds as follows: We review literature on the topics of vacant, abandoned, and deteriorated properties and spatial technologies. We next describe the Savannah, Georgia case study and details on training *VADecide*. We then show the results of *VADecide*, compared with the ML model and the city's current workflow, and discuss how HITLML aids the city's planning decision-making process.

Literature Review

Blighted and VAD Properties

Certain properties in our communities often exhibit problems that give them the distinction as "blighted" properties. Such issues include tax delinquency, prolonged vacancy, structural defects, features that are health hazards (e.g., asbestos, lead paint, or exposed wood shards), overgrowth in vegetation, missing utility infrastructure (e.g., pipes or ventilation systems), and may be subject to flooding (see Luce 2000). The term blight is often used as a legal basis for cities to acquire properties through condemnation (e.g., eminent domain) (Gold and Sagalyn 2010). However, this term has been criticized as a stigmatizing term for properties and their neighborhoods (Mallach 2018), and therefore, we use the distinction of VAD properties to refer to those that exhibit physical deterioration and neglect (Brachman 2005; Center for Community Progress 2023).

To our knowledge, there is no known universal definition of VAD (or blighted) properties, and definitions may be tailored to a specific need, such as raising awareness for substandard low-income housing, revitalizing downtown through eminent domain, or revealing the impacts of mortgage crises (Schilling and Pinzon 2016). VAD properties may have lower property values compared with those in the same block, as found in research comparing the sales price of foreclosure properties and the neighboring occupied units (Sumell 2009; Whitaker and Fitzpatrick 2013). Yet in practice, property value can be slow to update and vary significantly based on property attributes and thus not a stable indicator for VADs. Therefore, the definition of VAD must be examined in a local, operational context (see our definition in the Case Study section) and the interactions between VAD characteristics should be considered when distinguishing VAD properties from other properties.

Causes and Impacts of VAD Properties

A property can deteriorate due to macroeconomic and demographic shifts, a housing market failure, or neglect. In the United States, many legacy cities (i.e., post-industrial shrinking cities) with sustained job and population loss over the past decade have struggled to support neighborhoods with concentrated vacancy (Mallach & Brachman 2013). In other cities, the impact of redlining, a race-based exclusionary housing

policy from the mid-20th century, continues to divert investment from predominantly African American neighborhoods and reinforce residential segregation (Rothstein 2017). The housing bubble burst in 2006 and 2007 also led to large-scale foreclosure on properties as the cost to maintain them (e.g., mortgages) exceeded their values (Mallach 2018). At the local level, VAD properties with unclear ownership status or delinquent taxes exceeding the value of the property can be characterized as zombie properties (Alexander 2015, 32), because they have high transaction costs and low property values.²

VAD properties can disproportionately affect marginalized communities' local housing markets by decreasing neighboring property values and impacting neighbors' quality of life (Mallach 2018; Whitaker and Fitzpatrick 2013). Neighborhoods with many VAD properties are associated with poor school quality (Sun et al. 2019), high crime rates (Branas, Rubin, and Guo 2012), higher male unemployment rate (Appel et al. 2014), and slower growth in property sales price (Gilreath 2013). These neighborhoods are also more likely to be home to low-income and African American households (Silverman et al. 2013; Sun et al. 2019) and suffer from declining home ownership and pessimistic perceptions of neighborhood trajectories (Mallach 2021). Geographically, VAD properties also tend to cluster (Hillier et al. 2003; Reyes et al. 2016; Weaver and Bagchi-Sen 2013), reinforcing the concentration of income, race, and housing market inequality in marginalized communities and stunting economic mobility. VAD properties also cost millions of USD in lost tax revenue and unrecoverable costs of managing overgrown grass, litter, and illegal dumping; securing open structures; and demolition (Immergluck 2016; Mallach 2018; Sumell 2009).

Spatial Decision Support Systems for Identifying and Managing VAD Properties

Early examples of identifying and managing VAD properties with large, digital, spatial datasets include the use of SDSSs. SDSSs combine spatial and aspatial data with analytical models and geovisualization to facilitate decision-making in a spatial context (Armstrong, Densham, and Rushton 1986; Hopkins and Armstrong 1985). For example, the Philadelphia Neighborhood Information System integrated housing information, web mapping, and logistic regression to identify likely abandoned properties (Hillier et al. 2003). Other SDSSs have detected VAD properties from historical records such as code violations or field surveys. The City of New Orleans developed a decision support scorecard system using logistic regression to recommend that city officials sell (on the market or privately) or demolish a property, based on local experts' scoring of multiple criteria related to the property's condition (Hillenbrand 2016). A similar study in the City of Youngstown, Ohio,

overlaid property characteristics such as vacancy rates and crime rates to create a combined index where properties that "scored" highest in these collective factors were prioritized for demolition (Morckel 2016). Others have used ML models such as random forest, decision tree, and gradient boosting to detect VAD properties and discover how each factor contributes to the prediction (e.g., City of Syracuse in Appel et al. 2014; City of Cincinnati in Reyes et al. 2016). In these case studies, records were often limited to certain neighborhoods. Similarly, the City of Savannah only has "labels" for a few dozen properties per year. These small sample sizes can deprecate the predictive quality of ML models, which typically require larger training datasets, posing a challenge to adopting ML models and SDSSs at the municipal level.

Most case studies are designed to involve planners before or after building the ML models (Hill et al. 2003; Appel et al. 2014; Reyes et al. 2016; Hillenbrand 2016 is an exception), rather than consulting planners in tasks such as labeling training data and auditing the model's decision mechanisms. As such, planners may be reluctant to adopt the technology due to a lack of trust in and understanding of the model's innerworkings.

Human Engagement in Urban Artificial Intelligence

Traditional supervised ML processes involve training a model on big data that has been classified (e.g., whether a property is labeled as VAD) and iterating through multiple probability and statistical tests to predict classes for new data (e.g., properties with unknown VAD labels). In contrast, people use HITLML to create new training data from small samples, audit the decision mechanisms behind the algorithms, and calibrate human practices (Holzinger 2016; Monarch 2021; Zhou and Chen 2018). For instance, human experts can carefully examine the results of uncertainty sampling, which selects samples that are most uncertain for machine judgments, reflect on the causes of uncertainty, and then provide the appropriate labels (human judgments) to help the machine learn. Thus, HITLML is especially beneficial in scenarios where there is no large, reliable, labeled dataset and where the complexity of the subject matter necessitates a mutual learning dynamic between the humans and machines, rather than supervision or automation.

The human-in-the-loop approach is part of a rising urban artificial intelligence (AI) movement that outlines the power of technology to govern and serve as a voice for communities (Boehner and DiSalvo 2016; D'Ignazio and Klein 2020; Loukissas 2019). Technology can encourage civic participation in projects that range from authoritarian, with little public input, to citizen-controlled, in partnership with the public (Arnstein 1969; Cardullo and Kitchin 2019). Others argue that articulating issues, such as design goals, from different stakeholders' perspectives (Le Dantec 2016) and collaborating with

expert teams to build, quantify, ground-truth, share, and visualize data (Williams 2020), facilitate the use of technology for public good. While these prior studies call for human-AI collaboration in urban planning, critical discussions around mechanisms and activities for effectively engaging planners and denizens with AI are still developing (Wilson 2022). There are also arguments in urban planning that advocate for a reorientation of technocratic planning toward communicative and collaborative traditions (Mattern 2021). Specifically, a growing number of studies have operationalized HITLML in urban planning: Zheng and Sieber (2022) applied the technique to topic modeling on the text of a corpus of grant proposals on the topic of smart cities, and Anwar (2022) followed suit for land cover mapping applications. Yet, there do not appear to be examples of the benefits of HITLML as applied to issues of housing.

Our research attempts to fill these gaps by proposing an explorative HITLML approach to elevate human presence in ML and evaluate the advantages and challenges of a HITLML approach in a planning context, especially for identifying VAD properties.

Case Study

The City of Savannah is a historic city (city pop 147,780 [U.S. Census 2020]), a popular tourist destination, and home to the Savannah Port, one of the busiest seaports in the United States (International Trade Administration n.d.). Savannah is classified by the Lincoln Institute of Land Policy (n.d.) as a legacy city (i.e., a post-industrial shrinking city) whose population peaked in the 1960s.

The city of Savannah's municipal code defines blighted property as one with two or more of the following conditions, including uninhabitable structure, inadequate utilities, safety hazards, environmental contamination, repeated illegal activity, active code violations beyond one year, or crime or public health hazards to properties in immediate proximity (City of Savannah 2016). While this legal definition helps the city acquire such properties through eminent domain, few acquisitions used strategy and thus these cases do not adequately represent all aspects of VAD conditions that can be used for ML.

The city has a growing number of observed residential properties that show signs of deterioration and neglect, as measured through code violations and tax delinquency. In 2019, 1,319 properties had code violations that indicate severe physical deterioration and 1,404 properties had at least three years of tax delinquent history. These properties tend to be in minority-concentrated neighborhoods with a sizable black population.

In 2019, Savannah allocated \$10 million USD to repair and redevelop 1,000 properties into affordable housing over the subsequent ten years (called the "1K-in-10 initiative") in neighborhoods long neglected or exploited by

profit-driven investors (Housing and Neighborhood Service n.d.). To meet this goal, planning experts from Savannah Housing and Neighborhood Service Department (HNSD) experimented with a data-driven approach in 2018. They acquired spreadsheet data from civic departments, visually examined parcel information on an online platform, normalized and averaged the score of each parcel's records in crime, tax, and code violations, and conducted field surveys to validate vacancy on site. However, this process is time-consuming and labor-intensive. The data collected only capture a snapshot of conditions and the field surveys only cover a few neighborhoods. A transparent, scalable, and sustainable VAD identification system that also considers local context and interactions between variables could help mitigate these issues.

In Savannah, VAD properties have characteristics that are outlined in Savannah's blight standards (City of Savannah 2016) and that are likely to ease acquisition efforts (based on HNSD planners' experiential knowledge) in the city's 1K-in-10 initiative. In the past, Savannah acquired only dozens of properties annually and thus lacks a large VAD dataset for traditional ML prediction. Therefore, we use HITLML workflow to label and audit a small subset of VAD candidates with HNSD planners, which enables VAD prediction on a larger set of properties.

Data

Independent variables used in the city's current workflow are based on Savannah's legal requirements for government acquisition. Such properties must have active code violations, crime (Part I and Part II types), or tax delinquency to qualify for tax (foreclosure) sales, nuisance abatement, or eminent domain. The two ML models include additional independent variables suggested by HNSD staff and theories of VAD characteristics (see specific variable definitions in Table 1 and variable nuances in Supplementary Material Section G).

VAD labels for properties identified in the city's current workflow come from a field survey: HNSD staff drove to a few low-income neighborhoods with many highly ranked VAD properties and rated properties based on visual cues aligned with legal blight criteria. In contrast, the ML model's VAD labels are defined by active code violations, and *VADecide*'s VAD labels are defined by both data and human expertise (see more details in the next sections). To validate ML model and *VADecide*'s predictions, we use the city's field survey and vacancy data collected by U.S. Post Service (USPS) (see Table 1 for details).

We acquired parcel-level data for the City of Savannah through our partnership with HNSD staff members. Crime and tax data include records from 2010 to 2019. Code violations data include entries from 2012 to 2019. In the two

Table 1. Independent, Dependent, and Validation Variables per Parcel in the City's Current Workflow, the Simple ML Model, and VADecide.

Variable	Year	Description	Source	Used in
Independent variables				
Crime	2010–2019	The number of crime incidents (Part I and Part II crime). It is weighted by recency and type in ML model and VADecide.	Police Department	City's workflow, VADecide, ML model
Drug crime	2010–2019	The number of drug crime incidents weighted by recency.	Police Department	VADecide, ML model
Active code violation	2012–2019	The number of active code violations. It is weighted by recency and type in VADecide.	Code Compliance	City's workflow
Delinquent tax	2010–2019	Total amount of delinquent city and county tax and unpaid special assessment. The unit is USD. Due to collinearity, it is combined with delinquent year as one variable when calculating feature importance.	County Tax Office City Tax Office City Revenue Office	City's workflow, VADecide, ML model
Total delinquent years	2010–2019	The number of years that the property has tax delinquency or unpaid special assessment. Due to collinearity, it is combined with delinquent amount as one variable when calculating feature importance.	County Tax Office City Tax Office City Revenue Office	VADecide, ML model
Unpaid special assessment tax Pct	2010–2019	The percentage of unpaid special assessment in total delinquent tax.	Derived	VADecide, ML model
Property value	2019	Property values estimated by computer-assisted mass appraisal. The unit is \$1,000.	Savannah GIS Parcel Shapefile	VADecide, ML model
Dependent variables				
VADs labeled by field survey	2019	Properties identified as VAD candidates in a 2019 field survey by human experts in selected neighborhoods.	Field Survey	City's workflow
VADs labeled by active code violations	2019	The number of active code violations under designated type "Condemnation," "Vacant Property Clean/Mow," "Unsafe Secure," and "Unsafe Demolition" in code compliance database, which indicates VAD conditions.	Code Compliance	ML model
VADs labeled by data + human expertise	2019	Properties identified as VAD candidates in the human-in-the-loop workflow by human experts based on independent (input) variables in the model and their experience.	Human experts	VADecide
Validation/consensus variables				
VADs labeled by field survey	2019	(repeated from above) Properties identified as VAD candidates in a 2019 field survey by human experts in selected neighborhoods.	Field Survey	VADecide, ML model
USPS vacancy	2021	Properties identified as vacant (not collecting their mail) by USPS staff for ninety days or longer.	USPS	City's workflow, VADecide, ML model

ML models, code violations and crime are defined as the weighted count of incidents: an incident is weighted higher if it is more recent and if the type particularly contributes to

VAD (see Supplementary Material Section A for details). Property value was dated to 2019 and USPS vacancy data was acquired in 2021.

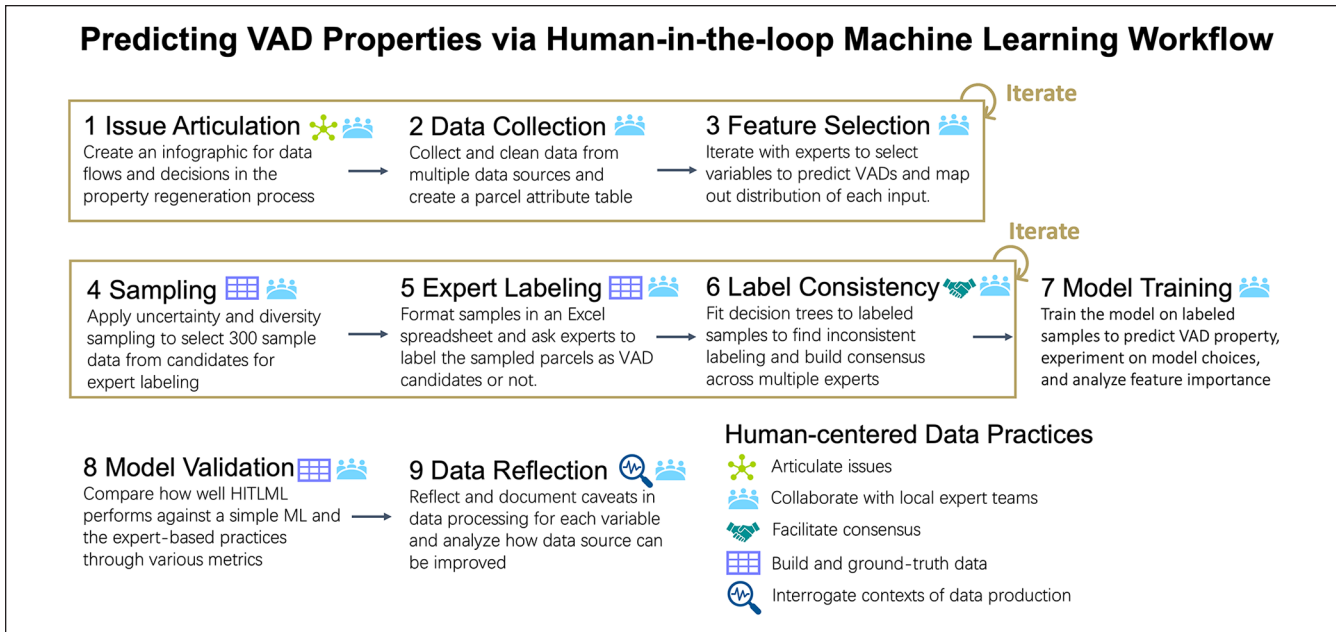


Figure 1. Human-in-the-loop machine learning (HILML) workflow used to create *VADecide*. The symbols represent various human-centered data practices we implemented in each step.

Methods

A Simple Machine Learning Model

This article defines a “simple” ML model (also referred to as the “ML model”) as a supervised model with few or no expert interactions. We labeled VAD properties as those with certain code violations, such as vacant property, need for cleaning/mowing, unsafe demolition, and condemnation. We removed code violation variables from the simple ML’s input features because we used code violation to construct the VAD labels. Unlike *VADecide*, the ML model learns VAD characteristics from labeled VADs in 2019 and then makes predictions without expert inputs, as is typical in similar studies.

Human-in-the-Loop Machine Learning (*VADecide*)

We created a workflow for *VADecide* (Figure 1) to identify/label example VAD properties. Steps 4 through 6 specifically distinguish a human-in-the-loop approach from a simple ML model and from the city’s current workflow (see detailed discussion of how the three approaches diverge in each step in Supplementary Material Section B).

Issue Articulation

To identify whether and where ML interventions help the decision-making process, we first sent a Q&A document (Supplementary Material Section C) to the HNSD experts

asking them to document the process of property regeneration from an administrative perspective. They reported that this process involves collecting data, identifying VAD candidates, determining acquisition strategies, acquiring properties, and regenerating the properties for various purposes. We next worked with the team at their offices in Savannah and online to codify how they decide to label a parcel. To help communicate these processes, we collaboratively created an infographic with the team (Supplementary Material Section D); this graphic shows that the biggest challenge is managing decentralized data that flows between institutions and leveraging the data for effective decision-making. By unveiling the decision process behind VAD property identification and acquisition, we realized that ML can help identify VAD candidates but is less successful at matching VAD candidates to acquisition strategies (e.g., tax sales, public nuisance, eminent domain), as this process requires contextual knowledge that is difficult to codify into data (e.g., some neighborhoods are more receptive to certain strategies). Thus, we narrowed the goal of the ML intervention as VAD identification.

Data Collection and Feature Selection

We created maps of Savannah symbolized by variables listed in Table 1 and examined them in person with the experts to decide on filtering criteria for model input data. This discussion found 5,372 residential structures (single-family, 2-4 family, and townhouse) and lands that have records in drug crime, tax issues, code violations, or fire incidents (required by Savannah city law for government property acquisition) and have little or no flood risk (which

was a priority of HNSD). While the expert workflow uses tax, code, and crime data as decision input, the ML models take in many variables and report an importance score for each. We experimented with fifteen variables, including building attributes (such as land size, number of bedrooms) and market indicators (such as number of unqualified sales) (see Supplementary Material Section A for all variables), but only kept a minimal list (see Table 1) that combines planners' recommendations and model suggestions for final prediction (Steps 2 and 3 in Figure 1).

VAD Properties Sampling and Labeling

Since Savannah does not have a database with properties labeled as VAD, we used a mixture of random sampling and active learning techniques (uncertainty and diversity sampling) to create training samples (i.e., which property can be a VAD candidate). Uncertainty sampling selects data points that are close to decision boundaries and thus have the highest uncertainty (Lewis and Catlett 1994), while diversity sampling ensures the geographic and representation diversity in samples (Monarch 2021) (see Supplementary Material Section E for technical details) (Step 4 in Figure 1).

After sampling, we asked a team of four housing experts to classify a subset of properties ($n = 300$) as VAD or not; this serves as a dependent variable for the model's training data (Step 5 in Figure 1). The experts were two males and two females ranging in age from their thirties to sixty-five. They have experience in the field of land and housing for twenty-five years, fifteen years, twelve years, and three years and have been working for the City of Savannah for eight years, three years (x2), and two years. They had no prior experience in ML or in labeling training data.

We sent each expert a spreadsheet of properties and their characteristics (see Supplementary Material Section E). In the spreadsheet, each column is a parcel, and each row is a parcel attribute. The experts could choose labels (i.e., VAD and Not VAD) from a dropdown menu in the spreadsheet and write comments to clarify how they made the decisions. One expert labeled 150 samples, and the others labeled fifty samples each.

After labeling the training data, we asked experts via e-mail to discuss labels that either conflicted with their professional knowledge or were deemed as highly uncertain by the ML algorithm (see Supplementary Material Section E) (Step 6 in Figure 1). We found that properties with very similar conditions were labeled differently, presenting opportunities to highlight implicit assumptions in the labeling process and to streamline the method to ensure equity across the process. To find these properties, we fitted decision trees to labeled data of parcels with and without structures, respectively, and visualized how the model made decisions for each type. For example, if a VAD property encountered multiple splits in the decision tree and was separated from a large group of non-VAD properties, then it was a good candidate

for discussion (see Supplementary Material Section D for visualization). When experts labeled a property differently, we presented the decision trees and the labeled samples, and the experts collaboratively agreed upon a resolution. This process helped expose nuances in the interactions of variables, differentiate whether false labels were due to low sample points or human error, and uncover tacit variables being used in the reasoning.

HITLML Model Training, Interpretation, and Prediction

To create *VADecide*, we used a *random forest algorithm* to classify properties into VAD and not VAD properties (Step 7 in Figure 1). The random forest model captures nonlinear relationships between variables by creating a "forest" of decision trees in which each tree decides on a sequence of variables and variable values to split the data to minimize label differences in each branch (Liaw and Wiener 2002). We chose a random forest model because it has higher model accuracy than logistic regression and because VAD status depends on a combination of variables (e.g., crime with tax delinquency); random forest can capture such feature interactions and is "smart" enough to pick between correlated variables (Basu et al. 2018) (see Supplementary Material Section E).

We used drop-column feature importance and partial dependence plots to experiment with different model settings, including different feature combinations and different numbers of training samples (see Supplementary Material Section E). In the end, our final HITLML model is trained on 300 samples (Land $n = 105$; Structure $n = 195$) with seven features (see Table 1).

Model Evaluation and Validation

We compared the *VADecide*-predicted VADs with ML-predicted VADs and VADs identified by the city's current workflow through five metrics (Step 8 in Figure 1).

Input refers to the number of VADs surveyed by the city or trained in the ML model and *VADecide*. A higher value means that more data collection is needed to be collected to make inference.

Output refers to the number of VADs found by the city or predicted by the ML model and *VADecide*. A higher value means that the process returned more VAD candidates.

Internal accuracy (ML and *VADecide* only) refers to the prediction accuracy trained with 80 percent of labeled VAD properties and validated with the remaining 20 percent (reserved from labeled VAD dataset, i.e., test sets). We used cross-validation and out-of-bag scores, two standard metrics in ML to measure the accuracy (see Supplementary Material Section F for Out-of-Bag (OOB) score results and technical explanations). A higher value indicates better prediction accuracy.

The external consensus rate refers to the percentage of properties that are also VAD candidates in other sources (i.e., the field survey and USPS vacancy data); a higher consensus rate indicates more reliable predictions. A higher value means more reliable prediction. Both field survey and USPS data are not perfect representations of the locations of VAD properties; the ideal “ground-truth” would be VAD properties that were identified and acquired in the past, but this population is too small and not diverse in our case study site.

Content sensitivity refers to the percentages of VAD properties dominated by different features. We focused on crime, code violation, tax delinquency, and low property values as VAD types because they are important features for predicting VADs in our random forest model (see Supplementary Material Section F for details on how we determine types). A higher value means that a method is less robust and is likely to retrieve VAD candidates based on a single variable.

Results

Comparison Between VADecide, Simple ML, and the City’s Current Workflow

Considering input-output, *VADecide* is the most efficient among the three, followed by the ML model and then the city’s current workflow. *VADecide* identified 1,309 VAD properties among 5,372 candidates, using only three hundred carefully selected expert-labeled samples. In contrast, the field survey from the city’s current workflow contains 890 effective samples and identifies only 192 as VAD candidates.

Next, both the *VADecide* and the simple ML model have high internal (prediction) accuracy. The minimum accuracy is *VADecide*’s prediction on structures (88%) and the maximum is ML’s prediction on structures (93%).

Regarding consensus rates, *VADecide* exceeds those produced by the ML and the city’s typical workflow. For example, *VADecide*’s predictions overlap with 60.49 percent VAD land and 63.96 percent VAD structures reported in the field survey and 66.67 percent VAD land and 34.61 percent VAD structures found in the USPS vacancy data. Since neither source represents the ground-truth status of VADs, the accuracy for *VADecide* may be even higher. The range of consensus rates for *VADecide* (34.61%–66.67%) is also much higher than the range in the ML model (14.81%–22.52%), indicating that the human-in-the-loop approach is more reliable, while the ML model performance may be biased by using a single variable (i.e., code violations) to label VAD samples. *VADecide* also outperforms the city’s current workflow (range, 26.15%–48%).

The three methods also exhibit distinct distributions of predominant VAD features, indicating that they have varying degrees of content sensitivity. This divergence implies that certain VAD characteristics exert greater influence (or bias)

than others. For instance, both the *VADecide* and the ML model predict more VAD properties with high crime and tax delinquency records, whereas the city’s current workflow identifies more properties with low property values. The *VADecide*-predicted VADs also have fewer code violations than those identified by the city. These discrepancies may be caused by differences in the human-versus-machine identification process. In the city’s *current* workflow, planning experts use low-income neighborhood location and low property value as heuristics for identifying easy-to-acquire VAD candidates. During field surveys, they are also more likely to rely on visual cues, such as dilapidated roofs and broken windows, which tend to be associated with code violations. The simple ML model predicted that VADs have higher rates of code violations than those detected from the city’s current workflow and from *VADecide*. This reveals that using one VAD feature (in this case, number of code violations) as training properties’ VAD labels may cause the ML model predictions to overrepresent properties with this feature. In contrast, *VADecide*’s results show that the percentage of identified VADs with various features (e.g., crime, code violations, tax delinquency, low properties values) falls between the percentages reported by the ML model and the city’s current workflow, indicating that they may be a “happy medium” between the techniques (Table 2). In comparison to VADs retrieved from the city’s workflow, *VADecide*-predicted VADs are spread more widely across neighborhoods and are distributed more evenly across land and structure parcel types (see Figure 2). This outcome may be the result of the city’s neighborhood visitation constraints, as field surveys cover a limited set of (low income) neighborhoods, while *VADecide* was trained with geographically diverse samples. The wider distribution of *VADecide*-predicted VADs mitigates human bias in associating VADs with marginalized communities. VADs predicted by the ML model are also distributed across more neighborhoods than those retrieved when using the city’s workflow (see Figure 2). Yet, the ML model’s predicted properties are concentrated around neighborhoods with available code violation data (see Supplementary Material Section G). For reference, USPS vacancy data have the widest geographic spread, but these data are biased toward vacant structures because structureless properties rarely receive mail.

Reflection

We met with HNSD staff and shared a table that outlined limitations and notes for each feature in the data collection and processing steps (see Supplementary Material Section G) (Step 9 in Figure 1). Together, we reflected on possible issues with interpretation, generalizability, and suggestions for improving data source quality. For example, in this post hoc examination, we noticed a lack of code cases from 311 calls (a civic hotline in Savannah and a key source of code violation data for the city) in certain neighborhoods. Using linear regression, we subsequently found that neighborhoods

Table 2. Comparison of City's Current Workflow, ML Model, and VADecide.

Metrics	Interpretation	City's workflow		ML model		VADecide	
		Land	Structure	Land	Structure	Land	Structure
Input	Total VADs surveyed in the city's current workflow or used to train the ML model or VADecide. Higher values indicate that more records are needed to make an inference.	483	407	183	353	105	195
Output	Total VADs identified in the city's current workflow or that were predicted by ML model or VADecide. Higher values mean that more VADs are identified.	81	111	96	281	472	837
Internal accuracy	Percent of ML-predicted VADs (trained with 80% labeled VADs) that are also labeled as VAD in the test set (20% of labeled VADs not used for ML training). Higher values mean better prediction accuracy.	NA	NA	90.72%	92.95%	93.33%	87.69%
External consensus	Percent of VADs identified/predicted that are also found in the city's field surveys (CV score)	NA	NA	14.81%	22.52%	60.49%	63.96%
	Percent of VAD candidates identified/predicted that are also found in the USPS vacancy dataset (CV score) Higher value in either indicates a more reliable prediction output.	48%	26.15%	16.67%	17.67%	66.67%	34.61%
Content sensitivity	Percent of identified VADs with <u>crime</u>	7.4%	22.3%	9.4%	30.6%	8.3%	46.5%
	Percent of identified VADs with <u>code violations</u>	34.6%	42.7%	82.3%	63%	29%	30.5%
	Percent of identified VADs with <u>tax delinquency</u>	75.3%	83.5%	100%	86.5%	99.8%	98.2%
	Percent of identified VADs with <u>low property value</u> Higher values mean that a higher percentage of VADs have these features and thus a particular method is less robust and is likely to depend on one data source versus others	96.3%	83.5%	86.5%	71.9%	90.3%	68.2%

Note: CV refers to cross-validation accuracy (see Supplementary Material Section F for further explanation). USPS refers to vacancy records compiled by USPS staff. NA values are present because the data collected using the city's current workflow (i.e., field surveys) could only be compared with USPS data, and do not run through the ML learning process.

with lower household income and higher percentage of African American population correlate with few 311 calls. We subsequently discussed possible ways to encourage more 311 calls from these neighborhoods (see Supplementary Material Section G).

Discussion and Conclusion

This article presented a HITLML approach (VADecide) that predicts VAD properties in Savannah, Georgia. We found that VADecide-predicted VADs not only have comparable accuracy to the classic ML model but are also more reliable than outcomes from both the ML model and the city's current workflow. We also discovered differences in what VADecide and HNSD (human experts) would classify as VAD: VADecide identifies more properties that are tax delinquent, prone to crime, and in neighborhoods with various income levels, while human experts often rely on visual

cues from code violations and low property value data. VADecide's predictions also covered a wider range of neighborhoods beyond the scope of human experts' field survey.

However, our study also has some limitations. First, our model focuses on challenges from the institutional perspective rather than the challenges of the property owners. As such, we did not incorporate the voices of community members, renters, and property owners, which reinforces a transactive, normative mode of governance (where denizens are consumers of public services) rather than a relational mode (where denizens are co-creators). Despite the benefits of integrating expert knowledge into the workflow, the human-in-the-loop approach has been criticized as a method that simply uses people as labor to generate and calibrate data (Janowicz, Sieber, and Crampton 2022). Second, involving housing experts to contextualize the selection of variables, build training data, adapt the model, and validate the model outcomes is costly. Model accuracy also requires constant



Figure 2. Geographic distribution of identified/predicted VADs in city's current workflow, ML model, and VADecide. USPS vacant properties are mapped as an external validation.

data updates, which is a common struggle in many smart city initiatives (Kitchin 2014). Finally, our model only identifies candidates based on VAD characteristics, but does not suggest how the government should proceed in managing the property. Acquisition in Savannah, and many other cities, is subject to political (e.g., minority communities' concerns with eminent domain) and financial factors (e.g., case-by-case acquisition costs and benefits).³ The extra properties identified by *VADecide* present opportunities for

early rehabilitation and intervention before these properties become uninhabitable. Future studies can investigate whether ML-identified VAD candidates are better for various acquisition strategies (e.g., tax sale, nuisance abatement, eminent domain, arms-length transaction) or regeneration options (e.g., as affordable housing, green space, rehabilitation), as compared with sites identified through fieldwork (see Figure S1 in Supplementary Appendix for the city's property regeneration workflow).

We recommend the human-in-the-loop model for municipal governments that lack an unbiased, large-scale VAD dataset for machine training, are ready to maintain and upgrade their digital infrastructures, and seek an alternative approach to evaluate their existing workflow. For cities without a consistent definition and acquisition of VAD properties, the HITLML model is beneficial as it requires only a small training sample to generate standardized, accurate, reliable, and robust predictions. The intentional process for human experts to define, deliberate, and dialogue with the HITLML model can also be enlightening and trust-building, even if it is a one-off practice. For instance, Savannah planners found the sample labeling process helpful in organizing their thoughts, making decisions, and reflecting on why they chose certain outcomes for the parcels. Two experts expressed increased confidence and knowledge about deploying the model in future operations. One expert also mentioned that interpreting the differences between human and machine-predicted outcomes made them rethink hidden assumptions in their current workflow.

The main barriers to consistently deploying the HITLML model in practice include technical and labor costs in maintenance. Critical digital infrastructures may include a central database of all VAD-relevant variables, often collected by separate civic departments, cloud resources to deploy the model, a visualization platform to overlay parcel details and predictions on the map, and an interface that allows humans to validate the predictions. Local governments may also need to consider the labor cost of hosting training workshops for staff to understand, recalibrate, and use the model and its outcomes, and develop guidelines for auditing and communicating results with community members. These challenges have also prevented the City of Savannah from adopting our HITLML model for the long term.

In conclusion, this research describes a HITLML approach for classifying land and residential structures as potentially in need of attention. By using a collaborative technique that engages experts while using large datasets and statistical analysis, we generated new insights into how to potentially automate or improve the local government's ability to identify VAD properties. The result is a more reliable method for managing assets in a municipal setting. We suggest that more researchers and practitioners collaborate and incorporate human input and expertise as they develop and test their ML applications.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was funded by the Georgia Partnership for Inclusive Innovation (PIN).

ORCID iDs

Xiaofan Liang  <https://orcid.org/0000-0002-0378-3608>

Clio Andris  <https://orcid.org/0000-0002-8559-5079>

Supplemental Material

Supplemental material for this article is available online.

Notes

1. Brian Brainerd, Discussion with authors, April, 2021.
2. Brian Brainerd, Discussion with authors, September, 2020.
3. Brian Brainerd, Discussion with authors, April, 2021.

References

- Alexander, Frank S. 2015. "Land Banks and Land Banking, 2nd Edition." Center for Community Progress. <https://community-progress.org/publications/land-banks-and-land-banking-2nd-edition/> (accessed July 13, 2024).
- Anwar, Sajjad. 2022. "Open Sourcing PEARL—Human-in-the-loop AI for Land Cover Mapping." *DevelopmentSEED*, <https://developmentseed.org/blog/2022-03-15-open-sourcing-pearl> (accessed March 15, 2022).
- Appel, Sheila U., Derek Botti, James Jamison, Leslie Plant, Jing Y. Shyr, and Lav R. Varshney. 2014. "Predictive Analytics Can Facilitate Proactive Property Vacancy Policies for Cities." *Technological Forecasting and Social Change* 89:161–73. doi:10.1016/j.techfore.2013.08.028.
- Armstrong, M. P., P. J. Densham, and G. Rushton. 1986. "Architecture for a Microcomputer Based Spatial Decision Support system." Proceedings of the Second International Symposium on Spatial Data Handling, Seattle, 120–131.
- Arnstein, Sherry R. 1969. "A Ladder of Citizen Participation." *Journal of the American Institute of Planners* 35 (4): 216–24. doi:10.1080/01944366908977225.
- Basu, Sumanta, Karl Kumbier, James B. Brown, and Bin Yu. 2018. "Iterative Random Forests to Discover Predictive and Stable High-order Interactions." *Proceedings of the National Academy of Sciences* 115 (8): 1943–48. doi:10.1073/pnas.1711236115.
- Boehner, Kirsten, and Carl DiSalvo. 2016. "Data, Design and Civics: An Exploratory Study of Civic Tech." Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, San Jose, 2970–81. doi:10.1145/2858036.2858326.
- Brachman, Lavea. 2005. "Vacant and Abandoned Property." <https://www.lincolnst.edu/publications/articles/vacant-abandoned-property> (accessed February 16, 2024).
- Branas, Charles C., David Rubin, and Wensheng Guo. 2012. "Vacant Properties and Violence in Neighborhoods." *ISRN Public Health* 2012:1–23. doi:10.5402/2012/246142.
- Cardullo, Paolo, and Rob Kitchin. 2019. "Being a 'Citizen' in the Smart City: Up and Down the Scaffold of Smart Citizen Participation in Dublin, Ireland." *GeoJournal* 84 (1): 1–13. doi:10.1007/s10708-018-9845-8.
- Center for Community Progress. 2023. "The Problem with Calling Neighborhoods with Vacant Properties 'Blighted.'" <https://communityprogress.org/blog/what-is-blight/> (accessed February 16, 2024).
- City of Savannah (2016) Georgia, Municipal Code § 2-18. <https://online.encodeplus.com/regs/savannah-ga/doc-viewer.aspx?secid=3757#secid-3757>

- D'ignazio, Catherine, and Lauren F. Klein. 2020. *Data Feminism*. Cambridge, MA: MIT Press.
- Gilreath, Morgan B. 2013. "A Model for Quantitatively Defining Urban Blight by Using Assessment Data." *Fair & Equitable*. https://www.iaao.org/media/Topics/FE_Aug13_Gilreath.pdf (accessed July 3, 2022).
- Gold, Martin E., and Lynne B. Sagalyn. 2010. "The Use and Abuse of Blight in Eminent Domain" *Fordham Urban Law Journal* 38:1119.
- Hillenbrand, Katherine. 2016. "New Orleans Brings Data-driven Tools to Blight Remediation." Harvard Data Smart City Solutions. Harvard University. <https://datasmart.ash.harvard.edu/news/article/new-orleans-brings-data-driven-tools-to-blight-remediation-915> (accessed February 16, 2022).
- Hillier, Amy E., Dennis P. Culhane, Tony E. Smith, and C. Dana Tomlin. 2003. "Predicting Housing Abandonment with the Philadelphia Neighborhood Information System." *Journal of Urban Affairs* 25 (1): 91–106. doi:10.1111/1467-9906.00007.
- Holzinger, Andreas. 2016. "Interactive Machine Learning for Health Informatics: When Do We Need the Human-in-the-loop?" *Brain Informatics* 3 (2): 119–31. doi:10.1007/s40708-016-0042-6.
- Hopkins, Lewis D., and Marc P. Armstrong. 1985. "Analytic and Cartographic Data Storage: A Two-tiered Approach to Spatial Decision Support Systems." Proceedings, Seventh International Symposium on Computer-Assisted Cartography, Washington, DC: American Congress on Surveying and Mapping.
- Housing and Neighborhood Service. n.d. "City of Savannah." <https://www.savannahga.gov/484/Housing-and-Neighborhood-Services> (accessed February 16, 2022).
- Immergluck, Dan. 2016. "Examining Changes in Long-term Neighborhood Housing Vacancy During the 2011 to 2014 US National Recovery." *Journal of Urban Affairs* 38 (5): 607–22. doi:10.1111/juaf.12267.
- International Trade Administration. n.d. "Maritime Services Trade Data." U.S. Department of Commerce. <https://www.trade.gov/maritime-services-trade-data> (accessed February 16, 2022).
- Janowicz, Krzysztof, Renee Sieber, and Jeremy Crampton. 2022. "GeoAI, Counter-AI, and Human Geography: A Conversation." *Dialogues in Human Geography* 12 (3): 446–58. doi:10.1177/20438206221132510.
- Kitchin, Rob. 2014. *The Data Revolution: Big Data, Open Data, Data Infrastructures and Their Consequences*. New York: Sage.
- Le Dantec, Christopher A. 2016. *Designing Publics*. Cambridge, MA: MIT Press.
- Lewis, David D., and Jason Catlett. 1994. "Heterogeneous Uncertainty Sampling for Supervised Learning." *Machine Learning Proceedings 1994*, New Brunswick, 148–56. doi:10.1016/B978-1-55860-335-6.50026-X.
- Liaw, Andy, and Matthew Wiener. 2002. "Classification and Regression by randomForest." *R News* 2 (3): 18–22.
- Lincoln Institute of Land Policy. n.d. "Legacy Cities." <https://www.lincolnst.edu/research-data/data-toolkits/legacy-cities> (accessed February 16, 2020).
- Loukissas, Yanni Alexander. 2019. *All Data Are Local: Thinking Critically in a Data-driven Society*. Cambridge, MA: MIT Press.
- Luce, Hudson Hayes. 2000. "The Meaning of Blight: A Survey of Statutory and Case Law." *Real Property, Probate and Trust Journal* 35:389–478.
- Mallach, Alan. 2018. *The Empty House Next Door: Understanding and Reducing Vacancy and Hypervacancy in the United States*. Cambridge, MA: Lincoln Institute of Land Policy.
- Mallach, Alan. 2021. *Making the Comeback: Reversing the Downward Trajectories of African American Middle-Income Neighborhoods in Legacy Cities*. Cambridge, MA: Lincoln Institute of Land Policy.
- Mallach, Alan, and Lavea Brachman. 2013. *Regenerating America's Legacy Cities*. Cambridge, MA: Lincoln Institute of Land Policy.
- Mattern, S., 2021. *A City Is Not a Computer: Other Urban Intelligences*. Princeton, NJ: Princeton Press.
- Monarch, Robert Munro. 2021. *Human-in-the-Loop Machine Learning: Active Learning and Annotation for Human-Centered AI*. New York: Simon & Schuster.
- Morekel, Victoria C. 2016. "Using Suitability Analysis to Prioritize Demolitions in a Legacy City." *Urban Geography* 38 (1): 90–111. doi:10.1080/02723638.2016.1147756.
- Reyes, Eduardo Blancas, Jennifer Helsby, Katharina Rasch, Paul van der Boor, Rayid Ghani, Lauren Haynes, and Edward P. Cunningham. 2016. "Early Detection of Properties at Risk of Blight Using Spatiotemporal Data." *Data for Policy 2016: Frontiers of Data Science for Government: Ideas, Practices and Projections*. Cambridge, United Kingdom. doi:10.5281/zenodo.556510.
- Rothstein, Richard. 2017. *The Color of Law: A Forgotten History of How Our Government Segregated America*. New York: Liveright Publishing.
- Schilling, Joseph, and Jimena Pinzon. 2016. *The Basics of Blight: Recent Research on Its Drivers, Impacts, and Interventions*. Alexandria, VA: Vacant Property Research Network.
- Silverman, Robert Mark, Li Yin, and Kelly L. Patterson. 2013. "Dawn of the Dead City: An Exploratory Analysis of Vacant Addresses in Buffalo, NY 2008–2010." *Journal of Urban Affairs* 35 (2): 131–52. doi:10.1111/j.1467-9906.2012.00627.x.
- Sumell, Albert. 2009. "The Determinants of Foreclosed Property Values: Evidence from Inner-city Cleveland." *Journal of Housing Research* 18 (1): 45–61. doi:10.1080/10835547.2009.12091996.
- Sun, Wei, Ying Huang, Ronald W. Spahr, Mark A. Sunderman, and Minxing Sun. 2019. "Neighborhood Blight Indices, Impacts on Property Values and Blight Resolution Alternatives." *Journal of Real Estate Research* 41 (4): 555–604. doi:10.22300/0896-5803.41.4.555.
- U.S. Census Bureau. 2020. Total Population, ACS 5-Years Estimates Detailed Tables. <https://data.census.gov/>
- Weaver, Russell C., and Sharmistha Bagchi-Sen. 2013. "Spatial Analysis of Urban Decline: The Geography of Blight." *Applied Geography* 40:61–70. doi:10.1016/j.apgeog.2013.01.011.
- Whitaker, Stephan, and Thomas J. Fitzpatrick IV. 2013. "Deconstructing Distressed-property Spillovers: The Effects of Vacant, Tax-delinquent, and Foreclosed Properties in Housing Submarkets." *Journal of Housing Economics* 22 (2): 79–91. doi:10.1016/j.jhe.2013.04.001.
- Wilson, Alexander, and Mark Tewdwr-Jones, M. 2022. *Digital participatory planning: Citizen engagement, democracy, and design*. New York: Routledge.
- Williams, Sarah. 2020. *Data Action: Using Data for Public Good*. Cambridge, MA: MIT Press.
- Zheng, Zhibin, and Renee Sieber. 2022. "Putting Humans Back in the Loop of Machine Learning in Canadian Smart Cities." *Transactions in GIS* 26 (1): 8–24.

Zhou, Jianlong, and Fang Chen. 2018. *Human and Machine Learning: Visible, Explainable, Trustworthy, and Transparent*. Cham: Springer.

Author Biographies

Xiaofan Liang is an assistant professor at Taubman College of Architecture and Urban Planning, University of Michigan. Her research is in spatial social networks, urban analytics, and participatory planning methods.

Brian Brainerd is a senior planner at the City of Savannah Housing & Neighborhood Services Department.

Tara Hicks is assistant director at the City of Savannah Housing & Neighborhood Services Department.

Clio Andris is an associate professor with a joint appointment in the School of City & Regional Planning and the School of Interactive Computing at Georgia Institute of Technology. Her research is in geographic information science, social networks, urban planning, and information visualization.