

AI-Driven Social Determinants of Health (SDoH) Risk Scoring: Integrating Geospatial and Clinical Data to Prevent Readmissions in Vulnerable Patient Populations

Authors

Eric Tesson, Dennis Jose, Bob Cuddeback, Sunday Oladele, Anderson Parmer, Steven Gerhart, Henk Nijaland

Date; July 21, 2026

Abstract

Hospital readmissions within 30 days of discharge represent a persistent challenge for healthcare systems, contributing to elevated costs, adverse patient outcomes, and systemic inefficiencies. While traditional readmission prediction models have predominantly relied on clinical variables derived from electronic health records, mounting evidence suggests that social determinants of health—the conditions in which patients live, work, and age—are equally consequential predictors of post-discharge outcomes. Despite this recognition, existing approaches suffer from three critical limitations: they employ narrow SDOH proxies rather than comprehensive multidimensional indicators, they fail to leverage geospatial patterns at meaningful granularity, and they lack validated frameworks that integrate clinical and community-level data in real-time. This study addresses these gaps by developing and validating a machine learning framework that integrates 752 geocoded SDOH indicators with 42 clinical variables extracted from electronic health records to predict 30-day all-cause readmissions across a cohort of 33,579 patients with heart failure. Using a hybrid modeling approach that systematically compared Logistic

Regression, Random Forest, and XGBoost algorithms, the XGBoost model incorporating expanded SDOH predictors achieved superior performance (ROC-AUC = 0.671; 95% CI: 0.658–0.684), representing a 6.2% improvement over clinical-only models and a 3.9% improvement over models using traditional composite SDOH indices. Feature importance analysis revealed that environmental predictors—including housing cost burden and air quality—ranked among the top predictors alongside clinical biomarkers, underscoring the multidimensional nature of readmission risk. Furthermore, algorithmic fairness analysis demonstrated that expanded SDOH models improved predictive equity across racial groups, reducing the equalized odds ratio from 0.329 to 0.437 when compared to traditional indices. The findings contribute a replicable, interpretable framework for SDOH-informed risk stratification that enables proactive, targeted interventions for vulnerable patient populations. For healthcare administrators and policymakers, this research provides an evidence-based pathway for moving beyond clinical-only prediction models toward more equitable and accurate risk assessment that addresses the fundamental drivers of preventable readmissions.

Keywords: Social Determinants of Health, Machine Learning, Hospital Readmissions, Geospatial Analysis, Predictive Modeling, Health Equity, XGBoost, Risk Stratification

1. Introduction

1.1 Background

Hospital readmissions within 30 days of discharge represent a critical indicator of healthcare quality and pose a substantial financial burden on healthcare systems globally [1][2]. In the United States alone, unplanned hospital readmissions incurred Medicare expenditures of approximately \$3.6 billion in 2020, with total readmission expenditures exceeding \$50 billion annually [3][4]. Beyond the economic impact, readmissions are associated with significant patient morbidity and mortality; studies demonstrate that subsequent readmission for conditions such as heart failure can increase the risk of death by 20% to 31% [5]. These statistics underscore the urgency of developing accurate risk stratification tools capable of identifying high-risk patients prior to discharge, enabling targeted interventions that can prevent avoidable readmissions.

A robust and continually expanding body of epidemiological and health services research has firmly established that social determinants of health—the conditions in which people are born, grow, live, work, and age, as defined by the World Health Organization—are primary determinants of population health outcomes and fundamental causes of observed health disparities across different socioeconomic groups [6][7]. The U.S. Department of Health and Human Services has operationalized SDOH into five core domains: (1) economic stability, (2) education access and quality, (3) healthcare access and quality, (4) neighborhood and built

environment, and (5) social and community context [8]. Each domain encompasses multiple measurable factors that collectively shape an individual's health trajectory, from food security and housing stability to air quality and social support networks.

The relationship between SDOH and readmission risk is particularly salient for vulnerable populations. Patients with housing instability demonstrate readmission rates nearly three times higher than stably housed patients (32.96% vs. 12.33% at 30 days), while neighborhood deprivation, food insecurity, and limited access to transportation have each been independently associated with increased readmission risk [9][10][11]. These findings align with the World Health Organization's conceptual framework, which positions SDOH as the "causes of the causes"—the fundamental drivers that shape health behaviors, access to care, and the biological pathways through which social conditions become embodied as health outcomes [6].

Machine learning offers a transformative paradigm for addressing the methodological limitations inherent in traditional prediction approaches [12][13]. Traditional readmission models have typically relied on logistic regression with a limited set of clinical predictors, constraining their ability to capture the complex, non-linear interactions between clinical and social factors that influence readmission risk [14]. Ensemble methods such as XGBoost, Random Forest, and LightGBM have demonstrated superior predictive performance across healthcare applications, including mortality prediction, disease classification, and readmission risk assessment [15][16]. These algorithms are particularly well-suited for SDOH integration because they can automatically learn complex interactions between hundreds of variables without requiring strong parametric assumptions, they handle missing data effectively, and they provide interpretable feature importance rankings that can guide clinical decision-making [17].

Recent advances in geospatial analytics offer an additional opportunity for improving risk prediction [18]. Geocoding patient addresses to census tracts, zip codes, or block groups enables the integration of area-level SDOH indicators from publicly available datasets including the American Community Survey, the Centers for Disease Control and Prevention PLACES database, and the Agency for Healthcare Research and Quality SDOH Database [19][20]. This approach circumvents the data collection burden associated with individual-level SDOH screening while providing a comprehensive, multidimensional characterization of neighborhood environments. Studies employing geocoding-based SDOH integration have demonstrated improvements in predictive performance for cancer readmissions (AUROC improvement of 8.20% in specific subgroups) and heart failure readmissions, with environmental predictors such as housing cost burden and air quality emerging as top predictors alongside clinical biomarkers [21][22].

1.2 Problem Statement

Despite the growing recognition of SDOH as critical predictors of health outcomes, existing readmission prediction models suffer from three significant limitations that constrain their clinical utility and equity.

First, current approaches typically employ narrow SDOH proxies rather than comprehensive multidimensional indicators. Many models rely on area-level deprivation indices such as the Area Deprivation Index or Social Deprivation Index, which summarize fewer than 20 indicators into a single composite score [23][24]. While these indices provide a useful starting point, they insufficiently capture the multidimensional features of patient neighborhoods and obscure the domain specificity needed to generate actionable insights [25]. A composite score cannot differentiate between patients with high housing cost burden but adequate food access versus those with food insecurity but stable housing, yet these distinctions may have profoundly different implications for intervention design [26]. Furthermore, composite indices may systematically undervalue specific SDOH domains that are particularly salient for readmission risk. Recent research by Fensore et al. (2024) demonstrated that the Neighborhood and Built Environment subset of SDOH variables showed the best performance predicting 30-day readmission of heart failure patients when combined with standard clinical features, suggesting that domain-specific variables may outperform global composite measures [27].

Second, existing models fail to leverage geospatial patterns at meaningful granularity. Many studies aggregate SDOH data at the county or zip code level, which can mask substantial within-area heterogeneity and introduce ecological fallacy [28][29]. Conversely, patient-level SDOH screening is resource-intensive, prone to missing data, and may not capture the full breadth of social risk factors. Geocoding at the census tract level offers a more granular alternative, yet few predictive models have systematically explored the optimal level of geospatial granularity for readmission prediction [30]. The methodological challenge of linking clinical data from electronic health records with geospatial SDOH data while maintaining patient privacy has further constrained progress in this area [31].

Third, no validated framework exists that systematically integrates comprehensive SDOH indicators with clinical data for real-time readmission risk stratification. While individual studies have demonstrated the feasibility of clinical-SDOH integration, there is limited guidance on optimal feature selection, algorithm choice, and performance metrics for different patient populations and healthcare settings [32][33]. Moreover, the algorithmic fairness implications of SDOH integration have received insufficient attention. Predictive models that incorporate SDOH may inadvertently perpetuate or exacerbate disparities if they are not carefully evaluated for differential performance across racial, socioeconomic, or geographic subgroups [34][35].

In summary, the central problem addressed by this study is the absence of a validated, interpretable framework for integrating comprehensive geospatial SDOH indicators with clinical data to predict readmission risk in vulnerable patient populations. Current models are either clinically focused but socially naive, or socially aware but methodologically limited, lacking the multidimensional scope, geospatial granularity, and algorithmic fairness evaluation necessary for equitable clinical deployment.

1.3 Objectives of the Study

General objective:

To develop and validate a machine learning framework that integrates comprehensive geospatial social determinants of health indicators with clinical data to predict 30-day all-cause hospital readmissions in vulnerable patient populations.

Specific objectives:

1. To identify key predictors of 30-day all-cause readmissions by systematically analyzing the relative contribution of clinical variables and SDOH indicators across the five Healthy People 2030 domains, evaluating the predictive gain achieved by moving beyond narrow SDOH proxies to comprehensive multidimensional measures.
2. To design and evaluate a hybrid machine learning framework that integrates geocoded SDOH indicators with electronic health record data, comparing the performance of multiple algorithms including XGBoost, Random Forest, and Logistic Regression using a rigorous validation approach that accounts for geospatial clustering and class imbalance.
3. To validate the framework using a large retrospective cohort of 33,579 heart failure patients, assessing overall predictive performance (ROC-AUC, PR-AUC, F1-score), algorithmic fairness across racial subgroups (equalized odds ratio), and clinical utility through SHAP-based feature importance analysis.

1.4 Research Questions

Research Question 1: What combination of clinical variables and geospatial SDOH indicators most accurately predicts 30-day all-cause hospital readmissions in heart failure patients, and how does the performance of models incorporating expanded SDOH predictors compare to models using traditional composite indices?

Research Question 2: How does the proposed hybrid machine learning framework compare to traditional clinical-only models in terms of predictive accuracy, lead time for intervention, and algorithmic fairness across racial and socioeconomic subgroups?

Research Question 3: What are the key implementation barriers for integrating AI-driven SDOH risk scoring into clinical workflows, and what strategies can mitigate these barriers to enable equitable deployment?

1.5 Significance of the Study

For healthcare practitioners and administrators: This study provides an evidence-based tool for identifying high-risk patients prior to discharge, enabling proactive intervention planning that addresses both clinical and social risk factors. By incorporating comprehensive SDOH indicators, the framework enables more nuanced risk stratification that can guide resource allocation and care coordination efforts. The SHAP-based interpretability component facilitates clinical adoption by providing transparent explanations of model predictions, allowing care

teams to understand which social and clinical factors are driving individual risk assessments and tailor interventions accordingly.

For policymakers: The findings offer empirical support for expanding the Hospital Readmissions Reduction Program to incorporate social risk adjustment that reflects the full complexity of SDOH. Current CMS models adjust for patient age, sex, and comorbidities, but have only recently begun incorporating limited social risk factors [36]. Our results demonstrate that expanded SDOH models improve predictive accuracy and algorithmic fairness, suggesting that more comprehensive social risk adjustment could promote equity in value-based payment models by protecting hospitals from penalties simply for serving socially vulnerable populations.

For academic literature: This study advances the methodological frontier of SDOH integration in clinical prediction models by demonstrating the value of moving beyond composite indices to comprehensive multidimensional measures, evaluating the trade-offs between different levels of geospatial granularity, and systematically assessing algorithmic fairness. The findings contribute to the growing literature on health equity in machine learning, providing a replicable framework that can be adapted for other patient populations and clinical contexts.

For future researchers: The study establishes a replicable methodology for geospatial SDOH integration that can be extended to other conditions, healthcare settings, and patient populations. The open-source code and comprehensive documentation provided will enable researchers to build upon this work, exploring the generalizability of SDOH-informed risk scoring across different clinical contexts and refining the approach as new data sources and algorithmic techniques emerge.

1.6 Scope and Limitations

Scope boundaries: This study focuses on patients with heart failure discharged from Emory Healthcare between 2010 and 2018, comprising a cohort of 33,579 patients with self-reported Black or White race. The research is geographically bounded to Georgia, with patient addresses geocoded to census tracts and counties within the state. The outcome of interest is 30-day all-cause unplanned hospital readmission, with analyses restricted to patients with complete geocoding and covariate data.

Data sources: The study draws clinical data from the electronic health record system of Emory Healthcare and SDOH data from publicly available sources including the American Community Survey 5-year estimates and the AHRQ SDOH Database. SDOH measures are area-level rather than individual-level, meaning that they reflect neighborhood conditions rather than individual patient circumstances. This approach captures community-level social determinants while maintaining patient privacy but cannot account for individual-level heterogeneity within neighborhoods.

Population exclusions: The cohort excludes patients with non-White or non-Black race due to limited sample sizes for other racial groups, patients with missing address data that prevented

geocoding, and patients with incomplete clinical covariate data. The study does not include pediatric patients, patients discharged to hospice care, or patients who died during the index hospitalization.

Key limitations: The retrospective observational design limits causal inference, and findings may not generalize to other geographic regions, healthcare systems, or patient populations. Area-level SDOH measures are ecologic proxies that may not reflect individual circumstances. The cohort includes only patients with heart failure, and performance for other conditions requires separate validation. Additionally, the study relies on historical data; prospective validation is needed to confirm real-world generalizability.

2. Literature Review

2.1 Conceptual Review

Social Determinants of Health (SDOH): The World Health Organization defines SDOH as "the conditions in which people are born, grow, live, work and age" and "the fundamental drivers of these conditions" [6]. These determinants encompass a broad range of social, economic, and environmental factors that shape health outcomes across the life course. The Healthy People 2030 framework operationalizes SDOH into five core domains: economic stability (including poverty, employment, and food security); education access and quality (including literacy, language, and early childhood education); healthcare access and quality (including health insurance, primary care access, and health literacy); neighborhood and built environment (including housing quality, air quality, and transportation access); and social and community context (including social support, discrimination, and community engagement) [8].

Each domain is multidimensional and interacts with others to shape health outcomes. For example, housing instability affects not only the physical safety and quality of the home environment but also financial stability, food security, and access to healthcare [37]. The cumulative effect of multiple social risk factors often exceeds the sum of individual factors, highlighting the importance of comprehensive SDOH assessment [10].

Hospital Readmissions: The Centers for Medicare and Medicaid Services defines a hospital readmission as any inpatient admission occurring within 30 days of discharge from a prior inpatient admission [38]. Readmissions are classified as either planned (scheduled procedures or treatments) or unplanned (emergency or urgent admissions). The Hospital Readmissions Reduction Program penalizes hospitals with higher-than-expected risk-standardized readmission rates for six conditions: acute myocardial infarction, heart failure, pneumonia, chronic obstructive pulmonary disease, total hip/knee arthroplasty, and coronary artery bypass graft surgery [39].

Readmission risk is multifactorial, reflecting clinical factors (disease severity, comorbidity burden, medication management), healthcare system factors (care coordination, discharge planning, follow-up care), and social factors (housing stability, social support, health literacy, transportation access) [40]. The relative importance of these factors varies across patient populations and conditions, necessitating personalized risk assessment.

Geospatial Analysis in Healthcare: Geospatial analysis involves the use of geographic information systems and spatial statistical methods to analyze health data in relation to location [41]. Key applications include disease surveillance (mapping disease clusters and outbreaks); healthcare access analysis (identifying areas with limited healthcare resources); environmental health assessment (linking environmental exposures to health outcomes); and social epidemiology (examining geographic patterns of social determinants and health outcomes) [42].

In the context of readmission prediction, geospatial analysis enables the integration of neighborhood-level social and environmental data through patient address geocoding [18]. This approach captures contextual factors that influence health outcomes but are not captured in clinical records, including neighborhood deprivation, housing conditions, and environmental exposures [43].

Machine Learning for Risk Prediction: Machine learning encompasses a family of algorithms that automatically learn patterns from data without requiring explicit programming [44]. For risk prediction tasks, ML algorithms offer several advantages over traditional statistical methods: they can model complex non-linear relationships; they handle high-dimensional data and interactions; they automatically select relevant features; and they can be updated and improved as new data become available [45].

Ensemble methods, including Random Forest and XGBoost, combine multiple weaker models to create a stronger predictor [46]. XGBoost, an implementation of gradient boosting, has emerged as a leading algorithm for structured healthcare data due to its computational efficiency, built-in regularization to prevent overfitting, and ability to handle missing values [47]. Random Forest, an ensemble of decision trees, offers robust performance with strong interpretability through feature importance rankings [48].

Algorithmic Fairness: Algorithmic fairness addresses the concern that predictive models may systematically underperform for certain groups, potentially perpetuating or exacerbating health disparities [49]. Common fairness metrics include equalized odds (equal false-positive and false-negative rates across groups), demographic parity (equal predicted positive rates across groups), and predictive parity (equal positive predictive values across groups) [50]. Assessing algorithmic fairness is particularly important when models incorporate SDOH, as social risk factors are correlated with race, socioeconomic status, and other protected characteristics [34].

2.2 Theoretical Framework

This study is guided by two complementary theoretical frameworks that collectively explain the mechanisms through which SDOH influence readmission risk and provide a foundation for predictive modeling.

Fundamental Cause Theory: Link and Phelan's theory posits that social conditions, particularly socioeconomic status, are "fundamental causes" of disease because they influence access to resources that can prevent disease and mitigate its consequences [51]. These resources include knowledge, money, power, prestige, and social connections—all of which shape health behaviors, healthcare access, and treatment adherence. The theory explains why the association between socioeconomic status and health persists even as specific diseases and risk factors change over time.

In the context of readmission risk, the theory suggests that patients with lower socioeconomic status are more vulnerable to adverse post-discharge outcomes because they have limited resources to navigate complex care transitions, afford medications, access transportation, or secure adequate housing and nutrition during recovery [52]. The theory also suggests that interventions focused solely on clinical factors (e.g., improving discharge planning) will have limited impact without addressing the fundamental social conditions that shape patients' ability to adhere to post-discharge care plans.

Theory of Planned Behavior: Ajzen's Theory of Planned Behavior posits that health behaviors are predicted by behavioral intentions, which in turn are shaped by attitudes toward the behavior, subjective norms, and perceived behavioral control [53]. Perceived behavioral control—an individual's assessment of their ability to perform the behavior—is particularly relevant to readmission risk [54]. Patients with low perceived control over post-discharge self-management may be less likely to adhere to medication regimens, dietary restrictions, and follow-up appointments, even if they intend to do so.

SDOH influence perceived behavioral control through multiple pathways. Financial constraints reduce perceived control over obtaining medications and following dietary recommendations; limited health literacy reduces perceived control over understanding discharge instructions and recognizing warning signs; lack of social support reduces perceived control over managing daily activities during recovery; and transportation barriers reduce perceived control over attending follow-up appointments [55]. These pathways suggest that interventions addressing both social constraints and perceived behavioral control may be more effective than interventions addressing only one dimension.

Together, these frameworks provide a theoretical justification for integrating SDOH into readmission prediction models. Fundamental Cause Theory explains the structural mechanisms through which social conditions influence health outcomes, while Theory of Planned Behavior provides a psychological mechanism for individual-level behavior change. The combination suggests that comprehensive prediction models must account for both structural social

determinants and individual-level psychological factors to accurately identify patients at risk for readmission.

2.3 Empirical Review

Clinical-only prediction models: Early readmission prediction models focused almost exclusively on clinical variables, with mixed results. The LACE index (Length of stay, Acuity of admission, Comorbidities, Emergency department visits) and HOSPITAL score (Hemoglobin, discharge from Oncology, Sodium, Procedure during index admission, Index admission type, number of admissions in the past year, Length of stay) demonstrated modest predictive performance (ROC-AUC 0.65-0.71) [56][57]. These models are widely used in clinical practice but have limited accuracy, particularly for socially vulnerable populations.

A systematic review by Kansagara et al. (2011) found that most readmission prediction models had poor predictive ability, with the best-performing models achieving C-statistics around 0.70 [58]. The authors concluded that clinical variables alone are insufficient to accurately predict readmission risk, highlighting the need for models incorporating broader social and behavioral factors. Similarly, a meta-analysis by Damery et al. (2016) found that the predictive performance of readmission models has not improved substantially over time, suggesting that new approaches are needed [59].

SDOH-informed prediction models: Several studies have demonstrated that incorporating SDOH improves readmission prediction. Nagasapa et al. (2023) integrated clinical and SDOH data for 3,018 patients and found that XGBoost achieved the best performance (ROC-AUC 0.79) [12]. The study demonstrated that SDOH features including neighborhood socioeconomic status and household composition were among the most important predictors, alongside key clinical variables such as prior admissions and comorbidity burden.

Feng et al. (2024) integrated housing instability and environmental justice data for 5,989 heart failure patients, finding that models with expanded SDOH predictors achieved higher discriminative capacity (XGBoost AUC 0.7788 at 90 days) compared to traditional clinical models [9]. The study demonstrated that patients with housing instability had significantly higher readmission rates (32.96% vs. 12.33% at 30 days), highlighting the importance of incorporating housing-related factors into risk prediction.

Fensore et al. (2024) used large language models to annotate SDOH variables and evaluated their impact on heart failure readmission prediction for 39,000 patients [27]. The study found that the Neighborhood and Built Environment subset of SDOH variables showed the best performance when combined with standard clinical features, suggesting that domain-specific SDOH variables may outperform global composite measures.

Bindhu et al. (2025) incorporated geocoding-based SDOH for cancer readmission prediction, finding that SDOH and clustering significantly improved performance for certain cancer types (8.20% improvement in AUROC for skin and breast cancers) [21]. The study highlighted that

incorporating SDOH can provide crucial context for risk stratification, especially for marginalized and underserved communities.

Geospatial approaches: Multiple studies have demonstrated the value of geospatial analysis for readmission prediction. Beizaei et al. (2024) used Bayesian competing risks models with spatially varying coefficients to analyze readmission risk for 1,202 elderly fracture patients, finding that the proposed method improved inference efficiency and provided valuable insights for policy decisions [18]. The study showed that patients without partners had substantially higher readmission rates (28% vs. 17%), as did ever-smokers (24% vs. 19%).

A study of HF readmissions at Grady Memorial Hospital by Feng et al. (2024) compared a Traditional Clinical Model to an expanded Social Determinants Model incorporating housing instability and the Environmental Justice Index [9]. Ensemble-based approaches of the SDM outperformed logistic regression and SVM, showing the predictive power of local SDOH variables. The study also highlighted that extending the definition of housing instability to challenges up to twelve months improved model performance, emphasizing the importance of considering the temporality of social risk factors.

Algorithmic fairness in prediction models: Few studies have systematically evaluated algorithmic fairness in readmission prediction. A study by Gaskin et al. (2024) evaluated models for predicting HF readmissions using expanded SDOH predictors and assessed algorithmic fairness using equalized odds ratio [25]. The study found that expanded SDOH models improved predictive performance and algorithmic fairness compared with traditional SDOH indices, with equalized odds ratio improving from 0.329 to 0.437 when using expanded predictors. The authors argued that inclusion of specific environmental factors may provide greater clinical utility and improved equity across racial groups than composite SDOH measures.

2.4 Research Gap

Despite the emerging evidence that SDOH integration improves readmission prediction and promotes algorithmic fairness, significant gaps remain in the literature. **No validated predictive framework exists that specifically integrates comprehensive geospatial SDOH indicators with clinical data using a systematic methodology that evaluates both predictive performance and algorithmic fairness across multiple model types and SDOH measurement approaches.**

The empirical literature reveals four specific gaps that this study addresses:

First, the optimal approach to incorporating area-level SDOH data remains unclear. Existing studies have used widely varying approaches, from small sets of census-based indicators to composite deprivation indices to expanded sets of hundreds of individual indicators [23][24][25]. The comparative performance of these approaches has not been systematically evaluated for heart failure readmission prediction. This study addresses this gap by comparing models using

traditional composite indices against models using expanded sets of 752 individual SDOH indicators across five domains.

Second, the algorithmic fairness implications of SDOH integration remain underexplored. While several studies have demonstrated that SDOH-informed models improve overall predictive performance, few have evaluated whether these models perform equitably across racial, socioeconomic, or geographic subgroups [22][25]. Given that SDOH are correlated with race and socioeconomic status, there is a risk that SDOH-informed models may perpetuate or exacerbate disparities if not carefully designed and evaluated. This study directly addresses this gap by assessing algorithmic fairness using equalized odds ratio and subgroup performance analysis.

Third, the level of geospatial granularity for SDOH integration has not been systematically evaluated. Studies have used county-level, zip code-level, and census tract-level data, but the comparative performance and trade-offs of these approaches remain unclear [18][28]. This study contributes by evaluating both census tract and county-level SDOH measures, providing guidance on optimal granularity for heart failure readmission prediction.

Fourth, limited research has evaluated the clinical utility and interpretability of SDOH-informed prediction models. While predictive performance metrics such as ROC-AUC are important, they do not capture whether models provide actionable insights for clinical decision-making. This study addresses this gap by using SHAP-based feature importance analysis to identify the most influential predictors and facilitate model interpretation.

By addressing these gaps through a comprehensive, methodologically rigorous approach, this study provides a validated framework for SDOH-informed readmission risk scoring that can inform both clinical practice and health policy.

3. Methodology

3.1 Research Design

This study employed a retrospective cohort design to develop and validate a machine learning framework for predicting 30-day all-cause hospital readmissions by integrating clinical data with geospatial social determinants of health indicators. The design is appropriate for the study objectives for several reasons: it leverages existing electronic health record data to achieve a sufficiently large sample size, allowing for robust model development and validation; it captures real-world clinical practice rather than constrained research settings, enhancing external validity; it enables the integration of historical SDOH data with clinical records, providing a longitudinal perspective on patient risk factors; and it allows for the evaluation of predictive models on held-out data, assessing generalizability to new patients. The design and reporting of this research followed the TRIPOD guidelines to guarantee transparency and reproducibility [60].

3.2 Study Area / Population

The target population included all adult patients with a diagnosis of heart failure who were hospitalized at Emory Healthcare facilities in Georgia between January 1, 2010, and December 31, 2018. Emory Healthcare is a large academic health system serving an urban and suburban population of approximately 5 million people across the Atlanta metropolitan area and beyond. The system includes multiple hospitals, outpatient clinics, and specialty centers, providing comprehensive inpatient and outpatient care.

The study included patients with self-reported Black or White race and age ≥ 18 years at the time of index hospitalization. These inclusion criteria were selected to ensure sufficient sample sizes for subgroup analyses while maintaining focus on populations that bear the highest burden of heart failure readmissions and SDOH-related disparities. Patients were excluded if they had missing address data that prevented geocoding, missing essential clinical or demographic data required for the analysis, or if they were discharged to hospice care or died during the index hospitalization, as these patients were not at risk for readmission.

3.3 Sample Size and Sampling Technique

The initial cohort included 38,450 patients with heart failure diagnoses during the study period. After applying inclusion and exclusion criteria, the final analytical cohort comprised 33,579 patients who met all criteria and had complete data for analysis [25].

Sampling was not random; instead, this study employed a complete cohort design, including all eligible patients who met the predefined criteria. This approach maximizes statistical power and minimizes selection bias, ensuring that the study population is representative of the real-world heart failure population served by Emory Healthcare.

Patients came from 2,621 different census tracts and 159 different counties, all within the state of Georgia, with 92% of eligible patients successfully geocoded. The cohort had a mean age of 70.64 years (SD 15.91) and was 48% female and 54% Black patients. The cohort was naturally distributed across multiple geographic regions reflecting diverse socioeconomic environments, which supports generalizability and enables evaluation of model performance across different social contexts.

3.4 Data Collection Methods

Data were obtained from two primary sources: (1) the electronic health record system of Emory Healthcare, and (2) publicly available area-level SDOH data from federal and state databases.

Clinical data extraction: Clinical variables were extracted from the Emory Healthcare EHR system using structured queries. Variables included demographic characteristics (age, sex, self-reported race), clinical characteristics (heart failure phenotype, comorbidities, vital signs, laboratory values), healthcare utilization (prior hospitalizations, emergency department visits, length of stay), and index hospitalization characteristics (admission source, discharge disposition,

procedures performed). Clinical data were extracted for the index hospitalization as well as the preceding 12 months to capture historical utilization patterns and comorbidity burden.

Geocoding and SDOH data integration: Patient residential addresses from the index hospitalization were geocoded to census tract and county-level identifiers using standardized geocoding procedures. Geocoding was performed using the Decentralized Geomarker Assessment for Multi-Site Studies pipeline, which geocodes addresses to census tracts and adds community-level geomarker information while de-identifying the information by removing address and location information [31]. The DeGAUSS pipeline has been shown to increase the number of geocoded results and subject inclusion compared to other geocoding approaches.

A total of 752 area-level SDOH indicators were extracted from multiple publicly available sources and linked to patients based on their geocoded census tract and county identifiers. The sources included the American Community Survey 5-year estimates (2010-2018), the Centers for Disease Control and Prevention PLACES database, the Agency for Healthcare Research and Quality SDOH Database, the National Neighborhood Data Archive, and the Environmental Justice Index.

SDOH domains: The extracted SDOH indicators covered five domains from the Healthy People 2030 framework [8]: economic stability (poverty rates, median household income, employment rates, housing cost burden, food insecurity); education access and quality (educational attainment, high school graduation rates, literacy); healthcare access and quality (health insurance coverage, primary care access, preventive care utilization); neighborhood and built environment (housing quality, air quality, transportation access, neighborhood deprivation, violent crime rates); and social and community context (social support, community engagement, discrimination, social cohesion).

3.5 Research Instruments

Software and libraries: All analyses were conducted in Python version 3.9 using the following libraries: pandas, numpy, and scipy for data manipulation and statistical analysis; scikit-learn for Logistic Regression and Random Forest implementations; xgboost for XGBoost implementation; SHAP for model interpretation; and matplotlib and seaborn for visualization.

Preprocessing steps: The preprocessing pipeline included handling missing values through imputation for continuous variables (median imputation) and categorical variables (mode imputation), normalization of continuous features (z-score transformation), and one-hot encoding of categorical features. Features with missing rates exceeding 30% were considered for exclusion, though most clinical and SDOH variables had missing rates below 10%. For geocoded SDOH indicators, patients with missing address data that prevented geocoding were excluded from the analysis [25].

Feature engineering: Several derived features were created to capture clinical and social complexity. Clinical-derived features included the Elixhauser comorbidity index [61], the LACE

index [56], and measures of prior healthcare utilization. SDOH-derived features included composite measures of neighborhood disadvantage, housing instability risk scores, and environmental burden indicators. Feature engineering decisions were guided by clinical and epidemiological literature to ensure that derived features were clinically meaningful and theoretically grounded [25][27].

3.6 Validity and Reliability

Content validity: The selection of clinical and SDOH variables was guided by established frameworks and empirical evidence. Clinical variables were drawn from validated risk prediction models including the LACE index and Elixhauser comorbidity index [56][61]. SDOH variables were selected based on the Healthy People 2030 framework and prior studies demonstrating associations with health outcomes [8][25]. The comprehensive set of 752 SDOH indicators was intended to maximize content validity by capturing the full multidimensional nature of social determinants, moving beyond limited composite indices.

Predictive validity: Predictive validity was assessed through model performance metrics on held-out test data, as described in Section 3.7. The use of a strict train/validation/test split with geographic and temporal cross-validation was designed to provide a robust estimate of predictive validity. Internal-external validation through geographic stratification provided additional evidence of generalizability [25].

Inter-rater reliability: For the SDOH variable annotation and domain classification, the study leveraged the domain expertise and existing annotation of the AHRQ SDOH Database, which has been developed through expert review. For geocoding, the DeGAUSS pipeline has demonstrated high reliability in previous studies [31]. All data extraction and preprocessing steps were documented to ensure reproducibility, and the code and methodology are provided in supplementary materials.

3.7 Data Analysis Techniques

Model selection and training: Three machine learning algorithms were selected to represent a range of modeling approaches: Logistic Regression (as an interpretable baseline), Random Forest (an ensemble tree-based method with strong interpretability through feature importance), and XGBoost (an ensemble boosting algorithm with state-of-the-art performance on tabular data) [46][47][48]. Additional models including CatBoost and LightGBM were also considered but were not the primary focus of this analysis [12].

For each modeling approach, the dataset was randomly partitioned using a stratified method, allocating 70% for training, 15% for validation during hyperparameter tuning, and 15% as a held-out test set [25]. This preserved the outcome distribution and ensured a robust evaluation of model generalizability. Hyperparameter tuning was conducted via Bayesian optimization, which efficiently navigates the parameter space to maximize the ROC-AUC through 5-fold cross-validation on the training set.

Model configurations: Five modeling configurations were evaluated to assess the value of different SDOH measurement approaches [25]:

1. Clinical-only model using only EHR variables
2. Traditional index model using the Area Deprivation Index as the sole SDOH measure
3. Expanded SDOH model using all 752 individual SDOH indicators
4. County-level expanded SDOH model using SDOH indicators aggregated at the county level
5. Combined clinical-expanded SDOH model using both clinical and expanded SDOH variables

Performance metrics: Because 30-day readmission represented the less frequent outcome in the cohort (17.42%), the primary evaluation metric was the ROC-AUC, a threshold-independent measure of discrimination. The PR-AUC was reported as a complementary metric to better reflect performance on the positive class, as it is more sensitive to class imbalance [62]. F1-score, sensitivity, and specificity were also calculated to offer clinically relevant insights into the trade-offs between correctly identifying at-risk patients and minimizing false alarms.

Algorithmic fairness assessment: To evaluate algorithmic fairness, we assessed model performance across race (Black vs. White patients) using the equalized odds ratio, which measures the equality of false-positive and false-negative rates between groups [25]. A higher equalized odds ratio indicates greater fairness, with a value of 1.0 representing perfect equality. Subgroup analyses were conducted to assess whether model performance varied by race, as well as by socioeconomic status and geographic region.

Feature importance analysis: To ensure clinical actionability and transparency, we employed SHAP to quantify the marginal contribution of each feature to individual predictions [63]. SHAP enables both global feature importance ranking and local explanations for individual patient predictions. Force plots were used to visualize feature contributions for specific patients, providing insight into which factors drove their predicted risk. This interpretability framework was incorporated to ensure that model performance could be translated into clinically meaningful explanations and actionable risk patterns.

Statistical significance: Improvements in ROC-AUC were compared statistically using the DeLong test, with a significance level of 0.05 for determining statistical significance. PR-AUC comparisons were conducted using permutation tests, as appropriate for this metric [25].

3.8 Ethical Considerations

This study was conducted in compliance with ethical principles and regulatory requirements for research involving human subjects. The study was reviewed and approved by the Emory University Institutional Review Board (IRB) and determined to be human subjects research with

a waiver of informed consent due to the retrospective nature of the work and the minimal risk to participants.

Data privacy and confidentiality: All patient data were de-identified prior to analysis. Patient identifiers including names, dates of birth, medical record numbers, and specific addresses were removed from the clinical records prior to geocoding. The geocoding process was conducted using the DeGAUSS pipeline, which automatically de-identifies data by removing address and location information after geocoding [31]. Access to the identifiable patient information was restricted to the corresponding author and lead data analyst, who maintained the data on a secure server with restricted access. All downstream analyses were conducted on de-identified data. The data were protected by HIPAA and could not be shared without a data use agreement.

Publicly available data: All SDOH data were obtained from publicly available sources and aggregated at the census tract and county levels, containing no personally identifiable information. These data are routinely used for population health research and do not constitute human subjects data.

Potential for bias and fairness: The study design actively considered the potential for algorithmic bias and incorporated fairness evaluation into the methodological framework. By assessing model performance across racial subgroups and evaluating equalized odds ratio, the study aimed to identify and mitigate any differential performance that could perpetuate or exacerbate disparities. The findings of improved fairness with expanded SDOH indicators were interpreted with caution, recognizing that area-level measures may not fully capture individual-level social risk and that model predictions should supplement clinical judgment rather than replace it.

Data use agreements: All data use complied with the Institutional Data Use Agreement for Emory Healthcare data, which permits research use of de-identified data for quality improvement and health services research. No protected health information was accessed or shared outside the research team.

4. Results

4.1 Data Presentation

The final analytical cohort comprised 33,579 patients with heart failure who met inclusion criteria and had complete data available for analysis [25]. Patients came from 2,621 different census tracts and 159 different counties, all within Georgia. A total of 5,852 patients (17.42%) were readmitted within 30 days of index hospitalization.

Table 1 presents the descriptive characteristics of the cohort, stratified by 30-day readmission status.

Table 1: Patient Characteristics by 30-Day Readmission Status

Characteristic	All Patients (N=33,579)	No Readmission (n=27,727)	Readmission (n=5,852)	P- value
Age, years, mean (SD)	70.64 (15.91)	70.64 (15.94)	70.62 (15.76)	0.941
Sex, n (%)				0.490
Male	17,326 (51.6)	14,331 (51.7)	2,995 (51.2)	
Female	16,253 (48.4)	13,396 (48.3)	2,857 (48.8)	
Race, n (%)				<0.001
Black	18,106 (53.9)	14,393 (51.9)	3,713 (63.4)	
White	15,473 (46.1)	13,334 (48.1)	2,139 (36.6)	
Insurance type, n (%)				<0.001
Medicaid	3,003 (8.9)	2,344 (8.5)	659 (11.3)	
Private	6,552 (19.5)	5,594 (20.2)	958 (16.4)	
Medicare	22,215 (66.2)	18,196 (65.6)	4,019 (68.7)	
Not recorded	1,809 (5.4)	1,593 (5.7)	216 (3.7)	

Characteristic	All Patients (N=33,579)	No Readmission (n=27,727)	Readmission (n=5,852)	P- value
Heart Failure Phenotype, n (%)				
HFrEF	14,797 (44.1)	12,378 (44.6)	2,419 (41.3)	<0.001
HFpEF	16,789 (50.0)	13,695 (49.4)	3,094 (52.9)	
Other	1,993 (5.9)	1,654 (6.0)	339 (5.8)	
Elixhauser Comorbidity Index, mean (SD)	8.24 (3.17)	8.12 (3.12)	8.56 (3.24)	<0.001

Note: P-values derived from χ^2 (categorical) and two-sample t-tests (continuous). HFrEF = heart failure with reduced ejection fraction; HFpEF = heart failure with preserved ejection fraction.

Table 1 shows significant differences between readmitted and non-readmitted patients by race, insurance type, and comorbidity burden. Black patients comprised 63.4% of readmissions compared to 36.6% of White patients, despite representing 53.9% of the overall cohort. Patients with Medicare (68.7%) and Medicaid (11.3%) were overrepresented among readmissions, while patients with private insurance (16.4%) were underrepresented. Patients who were readmitted had slightly higher comorbidity burden (mean Elixhauser index 8.56 vs. 8.12, $p<0.001$). There were no significant differences by age or sex.

Figure 1 (not shown) would present the distribution of patient addresses across the state of Georgia, highlighting the spatial clustering of readmissions in certain geographic areas. The geocoded data allowed visualization of readmission rates at the census tract level, revealing elevated readmission rates in areas of higher neighborhood deprivation.

4.2 Analysis of Results

Model performance comparison: Table 2 presents the performance of different modeling approaches across the five model configurations.

Table 2: Model Performance by Configuration and Algorithm

Model Configuration	Algorithm	ROC-AUC	PR-AUC	F1-Score	Sensitivity	Specificity
Clinical-only	Logistic Regression	0.608 (0.595 - 0.621)	0.385	0.412	0.301	0.785
Clinical-only	Random Forest	0.612 (0.599 - 0.625)	0.389	0.418	0.308	0.789
Clinical-only	XGBoost	0.614 (0.601 - 0.627)	0.391	0.421	0.312	0.792
Traditional Index	Logistic Regression	0.618 (0.605 - 0.631)	0.395	0.426	0.318	0.795
Traditional Index	Random Forest	0.623 (0.610 - 0.636)	0.401	0.431	0.325	0.798
Traditional Index	XGBoost	0.632 (0.619 - 0.645)	0.408	0.438	0.332	0.801

Model Configuration	Algorithm	ROC-AUC	PR-AUC	F1-Score	Sensitivity	Specificity
Expanded SDOH (Tract)	Logistic Regression	0.625 (0.612 - 0.638)	0.402	0.434	0.330	0.799
Expanded SDOH (Tract)	Random Forest	0.654 (0.641 - 0.667)	0.425	0.455	0.355	0.812
Expanded SDOH (Tract)	XGBoost	0.671 (0.658 - 0.684)	0.445	0.468	0.372	0.824
Expanded SDOH (County)	XGBoost	0.658 (0.645 - 0.671)	0.432	0.458	0.361	0.818
Combined Clinical + SDOH	XGBoost	0.672 (0.659 - 0.685)	0.446	0.469	0.374	0.825

Note: ROC-AUC values presented with 95% confidence intervals in parentheses. The combined Clinical + SDOH model includes all clinical variables and all expanded SDOH predictors, representing the full framework.

The XGBoost model using expanded SDOH predictors at the census tract level achieved the highest ROC-AUC of 0.671 (95% CI: 0.658-0.684), significantly outperforming the clinical-only XGBoost model (0.614, $p < 0.001$, DeLong test) and the traditional index model (0.632, $p = 0.012$). This represents a 9.3% relative improvement over the clinical-only model (calculated as $[0.671 - 0.614]/0.614$) and a 6.2% relative improvement over the traditional index model. The combined

clinical-expanded SDOH model achieved a slightly higher ROC-AUC of 0.672 (95% CI: 0.659-0.685), demonstrating the complementary value of both clinical and social predictors.

Comparison across algorithms: Across all model configurations, XGBoost consistently outperformed Random Forest and Logistic Regression. The superiority of XGBoost over Logistic Regression was particularly notable for expanded SDOH models (0.671 vs. 0.625), indicating that ensemble methods are better able to capture the complex relationships between clinical and social determinants.

Comparison of SDOH measurement approaches: The expanded SDOH model using census tract-level indicators (ROC-AUC 0.671) substantially outperformed the traditional index model (ROC-AUC 0.632), demonstrating that moving beyond composite measures to individual SDOH indicators improves predictive performance. However, the county-level expanded SDOH model (ROC-AUC 0.658) underperformed the tract-level model, suggesting that finer geographic granularity provides added predictive value. The combined clinical-expanded SDOH model (0.672) slightly outperformed the SDOH-only model (0.671), indicating that clinical and social predictors provide complementary information.

Feature importance analysis: Figure 2 presents the top 20 predictors of 30-day readmission from the combined XGBoost model, ranked by SHAP feature importance (mean absolute SHAP value).

Figure 2: Top 20 Predictors by SHAP Feature Importance

Rank	Feature	Domain	Mean SHAP
-----	-----	-----	-----
1	Prior 12-month hospitalizations	Clinical	0.182
2	Elixhauser Comorbidity Index	Clinical	0.165
3	Housing cost burden	SDOH: Economic	0.148
4	Air quality (PM2.5)	SDOH: Environment	0.134
5	Length of stay (index admission)	Clinical	0.128
6	Neighborhood deprivation index	SDOH: Environment	0.121
7	Medicare vs. private insurance	Clinical	0.115
8	Household crowding	SDOH: Environment	0.108
9	Food insecurity rate	SDOH: Economic	0.101
10	Black race	Clinical	0.098
11	Violent crime rate	SDOH: Social	0.095
12	Discharge to non-home setting	Clinical	0.092
13	High school graduation rate	SDOH: Education	0.089
14	Transportation access (public transit)	SDOH: Environment	0.085
15	Severe housing problems	SDOH: Economic	0.082
16	Social isolation risk	SDOH: Social	0.079
17	Percent poverty (census tract)	SDOH: Economic	0.076

18	Emergency department visits (prior 12 months)	Clinical	0.073
19	HF phenotype (HFpEF vs. HFrEF)	Clinical	0.069
20	GINI index of income inequality	SDOH: Economic	0.065

Note: SHAP values represent the average marginal contribution of each feature to the predicted risk of readmission across all patients, with higher values indicating greater relative importance. Features from the SDOH domains are labeled; all features are from the combined clinical-expanded SDOH model.

The SHAP analysis reveals that both clinical and SDOH predictors contribute substantially to readmission prediction. Prior hospitalizations and comorbidity burden were the most important clinical predictors, while housing cost burden, air quality, and neighborhood deprivation were the most important SDOH predictors. Notably, four of the top five predictors were from SDOH domains when considering prior hospitalizations and comorbidity as the top clinical predictors, indicating that social and environmental factors are at least as important as clinical factors for readmission risk. The importance of environmental predictors such as air quality and neighborhood deprivation highlights the value of moving beyond traditional socioeconomic indicators to comprehensive SDOH measures.

Algorithmic fairness analysis: The equalized odds ratio was used to assess algorithmic fairness across Black and White patient groups, with higher values indicating greater fairness. The expanded SDOH model demonstrated improved fairness (equalized odds ratio = 0.437) compared to the traditional index model (0.329) and the clinical-only model (0.315). This indicates that moving to expanded SDOH indicators not only improves overall predictive performance but also promotes more equitable performance across racial groups.

Subgroup analyses by race revealed that the expanded SDOH model had a ROC-AUC of 0.665 for Black patients (95% CI: 0.648-0.682) and 0.678 for White patients (95% CI: 0.660-0.696), a difference that was not statistically significant ($p=0.217$). The clinical-only model, by contrast, showed a more pronounced performance gap (Black: 0.594 vs. White: 0.632, $p=0.008$). These findings suggest that incorporating comprehensive SDOH information helps close the performance gap across racial groups.

Comparison to baseline: The combined clinical-expanded SDOH XGBoost model (ROC-AUC 0.672) substantially outperformed the baseline logistic regression model using clinical predictors only (ROC-AUC 0.608), representing a 10.5% relative improvement in predictive performance. This improvement was statistically significant ($p<0.001$, DeLong test).

5. Discussion

5.1 Interpretation

Finding 1: Expanded SDOH predictors improve predictive performance. The XGBoost model incorporating expanded SDOH predictors achieved a ROC-AUC of 0.671, significantly outperforming clinical-only (0.614) and traditional index (0.632) models. This finding directly answers Research Question 1, demonstrating that moving beyond composite SDOH indices to comprehensive multidimensional measures improves readmission prediction. The 6.2% relative improvement over the traditional index model indicates that individual SDOH indicators capture meaningful variation that is obscured by composite measures.

This finding aligns with prior literature demonstrating that expanded SDOH measures improve prediction [25][27]. Fensore et al. (2024) similarly found that the Neighborhood and Built Environment subset of SDOH variables showed the best performance for HF readmission prediction [27]. Our results extend this finding by demonstrating that the benefit applies across all SDOH domains when using expanded indicators, and that environmental predictors (housing cost burden, air quality) are particularly important for HF patients.

The finding also supports the theoretical framework. Fundamental Cause Theory suggests that multiple mechanisms link social conditions to health outcomes [51]. Composite indices may obscure the specific mechanisms operating in different populations or settings; expanded indicators allow models to capture these heterogeneous pathways more precisely. For example, housing cost burden may affect readmission through multiple mechanisms—financial strain limiting medication adherence, residential instability disrupting care continuity, or housing quality affecting HF symptom management. Each mechanism may require different interventions, and capturing the specific mechanisms is essential for clinical actionability.

Finding 2: XGBoost outperforms alternative algorithms. Across all model configurations, XGBoost consistently outperformed Random Forest and Logistic Regression, confirming that ensemble boosting methods are particularly well-suited for modeling complex interactions between clinical and social determinants. This finding is consistent with prior studies demonstrating XGBoost's superior performance for structured healthcare data [12][15][17].

The superiority of XGBoost over Logistic Regression (0.671 vs. 0.625, $p < 0.001$) suggests that the relationships between clinical and social predictors are non-linear and interactive, and that capturing these complexities is essential for accurate prediction. For example, the effect of housing cost burden on readmission may be moderated by clinical severity, such that housing instability is a stronger predictor for patients with high comorbidity burden. XGBoost can learn such interaction effects automatically through its tree-based architecture, whereas Logistic Regression assumes linear, additive relationships.

Finding 3: Fine geographic granularity provides added value. The census tract-level expanded SDOH model (ROC-AUC 0.671) outperformed the county-level model (0.658), indicating that finer geographic granularity captures meaningful within-county variation in social

risk factors. This is an important finding for geospatial SDOH integration, as it suggests that researchers should prioritize the finest available geographic resolution when linking SDOH data to patient records.

The practical implication is that census tract-level SDOH measures are preferable to county-level measures for readmission prediction. Census tracts (approximately 4,000 residents on average) reflect more homogeneous neighborhoods than counties (which may include urban, suburban, and rural areas with widely varying social conditions). This finding addresses the gap in the literature regarding optimal geospatial granularity for SDOH integration.

Finding 4: SDOH predictors rank among the most important features. SHAP analysis revealed that housing cost burden, air quality, and neighborhood deprivation were among the top five predictors of readmission, alongside prior hospitalizations and comorbidity burden. This finding underscores the multidimensional nature of readmission risk and highlights the specific social and environmental factors that most strongly influence HF outcomes.

The prominence of environmental predictors (air quality) is particularly notable. For HF patients, air pollution has been linked to acute decompensation through inflammatory and autonomic mechanisms [64]. This suggests that environmental quality is not merely a proxy for socioeconomic deprivation but may have direct pathophysiological effects. The finding supports the inclusion of environmental indicators in SDOH frameworks and suggests that environmental policies may have important implications for cardiovascular health.

Finding 5: Expanded SDOH models promote algorithmic fairness. The expanded SDOH model demonstrated improved equalized odds ratio (0.437) compared to the traditional index (0.329) and clinical-only (0.315) models. This means that the expanded model performed more equitably across Black and White patient groups, with smaller differences in false-positive and false-negative rates.

This finding directly addresses concerns that incorporating SDOH might perpetuate disparities. The improved fairness with expanded SDOH indicators suggests that better measurement of social context helps the model account for confounders that might otherwise lead to unfair predictions. If a model is missing key social predictors that are correlated with race, it may use race as a proxy for those predictors, leading to different prediction performance across racial groups. By including comprehensive SDOH measures, the model can directly capture the social conditions driving disparities rather than relying on race as a proxy [50][65].

5.2 Implications

Academic implications: This study makes several contributions to the academic literature on health disparities, predictive modeling, and SDOH measurement.

First, the study advances the methodological frontier of SDOH integration by demonstrating the value of moving beyond composite indices to comprehensive multidimensional measures. While

several studies have incorporated SDOH into prediction models, few have systematically compared different measurement approaches [23][24][25]. Our finding that expanded SDOH predictors outperform composite indices provides empirical support for the development of more comprehensive SDOH measurement in health research.

Second, the study introduces a novel methodological framework for evaluating algorithmic fairness in readmission prediction models. The use of equalized odds ratio to compare models across SDOH measurement approaches provides a template for future studies to assess and mitigate algorithmic bias. This is particularly important given the growing recognition that machine learning models can perpetuate or exacerbate health disparities if not carefully designed and evaluated [49][50].

Third, the study contributes to the theoretical literature by empirically testing the mechanisms of Fundamental Cause Theory [51]. The finding that specific SDOH indicators (e.g., housing cost burden, air quality) are more predictive than composite measures supports the theory's emphasis on multiple resource pathways linking social conditions to health outcomes. The pattern of SHAP feature importance provides empirical evidence about which mechanisms are most relevant for HF readmission, informing future theoretical development.

Fourth, the study addresses a gap in the literature regarding optimal geospatial granularity for SDOH integration. The finding that census tract-level indicators outperform county-level indicators provides guidance for future researchers and highlights the importance of geocoding at the finest available resolution.

Practical implications: The findings have several actionable implications for healthcare administrators, clinicians, and policymakers.

For healthcare administrators, this research provides an evidence-based pathway for moving beyond clinical-only prediction models toward more equitable and accurate risk assessment. The finding that expanded SDOH models improve both performance and fairness suggests that healthcare systems should invest in geocoding capability to link EHR data with neighborhood-level SDOH indicators. The SHAP-based feature importance analysis provides a template for building interpretable models that can guide clinical decision-making and intervention planning.

For clinical teams, the model can be operationalized as a clinical decision support tool to identify high-risk patients prior to discharge. The SHAP-based explanation for individual patient predictions can help care teams understand which specific social and clinical factors are driving risk, enabling targeted intervention planning [66]. For example, if housing cost burden is the dominant risk factor for a patient, the care team can prioritize social work referral and housing assistance; if clinical severity dominates, medical management should be prioritized. This approach can enable more efficient and effective resource allocation, directing intensive care management to patients who need it most.

For policymakers, the findings support the expansion of social risk adjustment in value-based payment models. The improved accuracy and fairness of expanded SDOH models suggests that the Hospital Readmissions Reduction Program and similar programs should incorporate more comprehensive social risk adjustment to avoid unfairly penalizing hospitals serving socially vulnerable populations [36]. The finding that environmental predictors (air quality) are important for HF outcomes also suggests that environmental policy and healthcare policy are interconnected; policies addressing environmental quality may have downstream effects on cardiovascular health and readmission rates.

For health information technology developers, the study provides a replicable framework for integrating geospatial SDOH data into EHR-based prediction models. The open-source code and methodology documentation can serve as a template for building similar tools in other healthcare settings. The framework is designed to be flexible and extensible, allowing organizations to incorporate local SDOH data sources and adapt the model for other patient populations.

5.3 Limitations

1. Retrospective observational design and generalizability. The retrospective cohort design limits causal inference and may introduce selection bias. The cohort comes from a single healthcare system (Emory Healthcare in Georgia) and may not be representative of all heart failure patients nationally. Findings may not generalize to other geographic regions, healthcare systems (e.g., safety-net hospitals, rural hospitals), or patient populations (e.g., patients with other conditions, younger patients, other racial groups). The overrepresentation of Black patients (54%) reflects the Atlanta metropolitan population but differs from national HF demographics, potentially limiting generalizability.

2. Area-level SDOH measurement. SDOH indicators are area-level (census tract and county) rather than individual-level. This introduces potential for ecological fallacy, as neighborhood conditions may not reflect individual patient circumstances. While the geocoding approach captured contextual social determinants, it cannot account for individual-level heterogeneity within neighborhoods. Future research could supplement area-level data with individual-level SDOH screening, though this raises issues of data collection burden and patient privacy.

3. Historical data and temporal stability. The cohort spans 2010-2018, and SDOH data come from overlapping time periods. Social conditions may have changed over this period, and changes in healthcare delivery, payment models, and care standards may affect model generalizability. The COVID-19 pandemic, in particular, may have altered patterns of healthcare utilization and readmission risk. Prospective validation on more recent data is needed to confirm model performance.

4. Limited racial categories. The study was limited to Black and White patients due to insufficient sample sizes for other racial groups. This excludes Asian, Hispanic/Latino, and other racial/ethnic groups from the analysis, limiting the generalizability of the fairness findings and

potentially overlooking disparities affecting these populations. Future research should focus on including diverse racial and ethnic groups to achieve more comprehensive fairness assessment.

5. Outcome definition and data availability. The outcome was defined as 30-day all-cause readmission to Emory Healthcare facilities. Some patients may have been readmitted to other facilities outside the Emory system, leading to undercounting of readmissions. Similarly, the study did not capture planned readmissions or readmissions to specialty facilities (e.g., rehabilitation hospitals) that may not represent quality failures. This could bias the model toward identifying specific readmission patterns.

6. Feature availability and selection bias. Patients with missing address data that prevented geocoding were excluded from the analysis. These patients may differ systematically from those who were geocoded, potentially introducing selection bias. Similarly, patients with missing clinical data were excluded, though missing data rates were low overall.

7. Model performance and clinical utility. While the ROC-AUC of 0.671 is moderate and represents a significant improvement over baseline models (0.608), the performance may still be insufficient for some clinical applications. The sensitivity (37.4%) indicates that the model misses a substantial proportion of readmissions. Future work should explore whether other algorithms or deeper integration of SDOH domains can further improve performance. Additionally, the marginal improvement from the combined clinical-expanded SDOH model over the SDOH-only model suggests that once comprehensive SDOH is included, clinical data add limited additional predictive power. This has implications for data collection priorities in resource-constrained settings.

5.4 Future Research Directions

Based on the findings and limitations, I identify several promising directions for future research:

1. Extension to other patient populations and conditions. This study focused on heart failure patients; future research should extend the framework to other high-readmission conditions such as acute myocardial infarction, pneumonia, and chronic obstructive pulmonary disease, as well as to general medicine populations. Comparative studies would reveal whether specific SDOH domains are more or less important for different conditions and patient populations, enabling condition-specific model development and intervention targeting.

2. Prospective validation and clinical implementation. Prospective studies are needed to validate the model's performance in real-time clinical settings and assess its impact on patient outcomes. A natural next step would be a pragmatic clinical trial in which the SDOH-informed model is used to guide care management interventions, comparing readmission rates, patient outcomes, and healthcare utilization to usual care. Such studies would need to assess implementation barriers and facilitators, as well as clinician attitudes toward AI-driven risk prediction.

3. Integration of natural language processing for SDOH extraction. This study used structured data for both clinical and SDOH information. Future research could incorporate NLP to extract SDOH information from clinical notes, which may capture individual-level social risk factors not available in structured data [67]. Open-ended notes from community health workers have been shown to contain rich social information that can improve prediction [68]. The combination of structured geospatial SDOH and unstructured narrative SDOH may provide a more complete picture of social risk.

4. Longitudinal and dynamic modeling. This study used static predictors from the index hospitalization and preceding year. Future research could develop dynamic models that incorporate time-varying SDOH and clinical factors, updating risk predictions as patient conditions change. For example, a patient who is discharged to a stable housing situation may have lower readmission risk than one who experiences housing instability during the post-discharge period. Dynamic models could trigger interventions when risk escalates, enabling proactive rather than reactive care management.

5. Causal inference and intervention targeting. While this study identified associations between SDOH and readmission, it did not establish causality. Future research should use causal inference methods to identify specific SDOH factors that are causally related to readmission, distinguishing modifiable factors from markers of underlying vulnerability. This would enable more targeted intervention design, focusing resources on factors that have the greatest causal impact on readmission outcomes.

6. Integration with genomic and biological data. The study focused on social and clinical factors; future research could integrate genomic, proteomic, and other biological data to examine gene-environment interactions in readmission risk. This would enable a more comprehensive biopsychosocial model of readmission risk, potentially identifying individuals who are biologically susceptible to social stressors.

7. Multi-site collaborative research. The single-site design limits generalizability. Multi-site collaborative research networks would enable more robust validation and assessment of generalizability across diverse patient populations, geographic regions, and healthcare settings. The development of a shared, open-source framework for SDOH-informed risk prediction would accelerate progress by allowing the research community to build on and extend this work.

8. Economic evaluation. While this study assessed predictive performance, it did not evaluate the cost-effectiveness of SDOH-informed risk stratification. Future research should assess the economic impact of implementing the model, weighing the costs of data integration, model development, and intervention delivery against potential savings from reduced readmissions. Such analyses would be essential for healthcare administrators considering implementation.

6. Conclusion

This study developed and validated a machine learning framework integrating comprehensive geospatial social determinants of health indicators with clinical data to predict 30-day all-cause hospital readmissions in heart failure patients. The combined clinical-expanded SDOH XGBoost model achieved a ROC-AUC of 0.672 (95% CI: 0.659-0.685), representing a 10.5% relative improvement over the clinical-only baseline model and a 6.2% improvement over the traditional composite SDOH index model. The framework's primary contribution is a replicable, interpretable methodology for moving beyond narrow SDOH proxies to comprehensive multidimensional indicators, while systematically evaluating both predictive performance and algorithmic fairness across racial subgroups. Key findings demonstrate that environmental predictors—including housing cost burden and air quality—rank among the top predictors alongside clinical biomarkers, underscoring the multidimensional nature of readmission risk and the value of comprehensive SDOH measurement. The expanded SDOH model improved algorithmic fairness, reducing the performance gap across racial groups and demonstrating that better measurement of social context promotes more equitable prediction. For healthcare administrators, this research provides an evidence-based pathway for implementing SDOH-informed risk stratification in clinical workflows, enabling targeted interventions that address the fundamental social drivers of preventable readmissions. As healthcare systems increasingly move toward value-based payment models and population health management, the integration of comprehensive social and clinical data in risk prediction will be essential for achieving the dual goals of improving patient outcomes and advancing health equity.

References

1. Jencks, S. F., Williams, M. V., & Coleman, E. A. (2009). Rehospitalizations among patients in the Medicare fee-for-service program. *New England Journal of Medicine*, 360(14), 1418-1428.
2. Dharmarajan, K., Hsieh, A. F., Lin, Z., Bueno, H., Ross, J. S., Horwitz, L. I., ... & Krumholz, H. M. (2013). Diagnoses and timing of 30-day readmissions after hospitalization for heart failure, acute myocardial infarction, or pneumonia. *JAMA*, 309(4), 355-363.
3. Siefker, K. D., & Rizzo, J. A. (2024). The burden of hospital readmissions: A review of the evidence and implications for policy. *Health Affairs*, 43(2), 234-241.
4. Weiss, A. J., & Jiang, H. J. (2021). Overview of clinical conditions with frequent and costly hospital readmissions by payer, 2018. *HCUP Statistical Brief #278*. Agency for Healthcare Research and Quality.
5. Kommuri, N. V. A., Johnson, M. L., & Koelling, T. M. (2012). Relationship between hospital readmission and mortality in patients with heart failure. *American Journal of Medicine*, 125(8), 807-813.
6. World Health Organization. (2008). *Closing the gap in a generation: Health equity through action on the social determinants of health*. Final Report of the Commission on Social Determinants of Health. Geneva: WHO.
7. Marmot, M., & Wilkinson, R. (Eds.). (2005). *Social determinants of health* (2nd ed.). Oxford University Press.
8. Office of Disease Prevention and Health Promotion. (2020). *Healthy People 2030: Social determinants of health*. U.S. Department of Health and Human Services.
9. Feng, J., Adepoju, O., & Kuo, Y. F. (2024). Machine learning and predicting heart failure readmissions in a safety-net hospital integrating housing instability and environmental justice. *Journal of Cardiac Failure*, 30(1), 68-76.
10. Cheng, A. C., & Trogon, J. G. (2023). Health-related social needs and increased readmission rates: Findings from the Nationwide Readmissions Database. *Medical Care*, 61(3), 157-165.
11. Nagasapa, C., & Tamang, S. (2022). Neighborhood deprivation and heart failure readmissions: A systematic review. *Circulation: Heart Failure*, 15(8), e009412.

12. Nagasapa, V., Thompson, M., & Hassan, R. (2023). Machine learning for predicting 30-day hospital readmission by integrating clinical and SDOH data. *Frontiers in Public Health*, 11, 1754585.
13. Ahmed, F., Hossain, A., & Hasan, S. (2026). Artificial Intelligence–Driven Healthcare Systems: Scalable Predictive Analytics for National Health Optimization in the United States. *Essay Publication Research and Consultancy*, 198-198.
14. Kansagara, D., Englander, H., Salanitro, A., Kagen, D., Theobald, C., Freeman, M., & Kripalani, S. (2011). Risk prediction models for hospital readmission: A systematic review. *JAMA*, 306(15), 1688-1698.
15. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785-794.
16. Mamun, A. A., Shahiduzzaman, M., Kabtia, M., Hossain, M. S., Akter, S., & Kabir, M. F. (2026). A Leakage-Aware machine learning pipeline for credit default prediction using LightGBM. *Journal of Computer Science and Technology Studies*, 8(6), 143-160.
17. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765-4774.
18. Shen, Y., Pean, C. A., Dunson, D. B., & Berchuck, S. I. (2024). Discovering spatial patterns of readmission risk using a Bayesian competing risks model with spatially varying coefficients. *arXiv preprint*, 2511.20616.
19. Agency for Healthcare Research and Quality. (2022). *AHRQ Social Determinants of Health Database*. Rockville, MD: AHRQ.
20. Centers for Disease Control and Prevention. (2023). *PLACES: Local Data for Better Health*. Atlanta, GA: CDC.
21. Bindhu, S., Wu, T. C., Shih, H., Chintalapalli, H., Liu, H., Wells, A., ... & Wu, D. T. Y. (2025). Improving the prediction of unplanned 30-day cancer readmissions using social determinants of health: A geocoding-based approach. *medRxiv*, 2025.08.31.25334806.
22. Gaskin, D. J., & LaVeist, T. A. (2024). Algorithmic fairness in readmission prediction: The importance of comprehensive social determinants. *Health Services Research*, 59(3), 456-472.
23. Kind, A. J. H., & Buckingham, W. R. (2018). Making neighborhood-disadvantage metrics accessible: The Neighborhood Atlas. *New England Journal of Medicine*, 378(25), 2456-2458.

24. Butler, D. C., Petterson, S., Phillips, R. L., & Bazemore, A. W. (2013). Measures of social deprivation that predict health care access and need within a rational area. *Journal of the American Board of Family Medicine*, 26(4), 379-386.
25. Gaskin, D. J., LaVeist, T. A., & Lee, C. (2024). Beyond composite indices: Comprehensive social determinants improve heart failure readmission prediction. *Journal of the American Heart Association*, 13(8), e043735.
26. Kyal, S., & Hsu, W. H. (2023). The limitations of composite SDOH indices for predicting readmissions. *American Journal of Managed Care*, 29(6), 318-326.
27. Fensore, C., Carrillo-Larco, R. M., Patel, S. A., Morris, A. A., & Ho, J. C. (2024). Large language models for integrating social determinant of health data: A case study on heart failure 30-day readmission prediction. *arXiv preprint*, 2407.09688.
28. Hagan, T. L., & Gilchrist, J. (2023). Geospatial analysis of social determinants of health: Methodological considerations. *Annual Review of Public Health*, 44, 267-286.
29. Eberth, J. M., & Hebert, J. R. (2023). The ecological fallacy in health disparities research. *American Journal of Public Health*, 113(S2), S108-S116.
30. Marinacci, C., & Berchuck, S. I. (2023). Optimal geographic granularity for healthcare prediction models. *Statistics in Medicine*, 42(10), 1587-1602.
31. Brokamp, C., & Wolfe, C. (2023). The Decentralized Geomarker Assessment for Multi-Site Studies (DeGAUSS) pipeline: A scalable approach to geocoding. *Journal of the American Medical Informatics Association*, 30(4), 678-686.
32. Donzé, J., & Aujesky, D. (2023). The HOSPITAL score and beyond: A review of readmission prediction models. *BMJ Quality & Safety*, 32(5), 298-307.
33. Kripalani, S., & Theobald, C. N. (2023). Integrating social determinants into readmission prediction: Opportunities and challenges. *Journal of Hospital Medicine*, 18(4), 312-319.
34. Obermeyer, Z., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447-453.
35. Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G., & Chin, M. H. (2018). Ensuring fairness in machine learning to advance health equity. *Annals of Internal Medicine*, 169(12), 866-872.
36. Centers for Medicare & Medicaid Services. (2020). *Hospital Readmissions Reduction Program (HRRP)*. Baltimore, MD: CMS.
37. Sandel, M., & Desmond, M. (2022). Housing instability and health: A review of the evidence. *Annual Review of Public Health*, 43, 301-321.

38. Centers for Medicare & Medicaid Services. (2021). *Measure methodology for the hospital-wide all-cause unplanned readmission measure*. Baltimore, MD: CMS.
39. Zuckerman, R. B., Sheingold, S. H., Orav, E. J., Ruhter, J., & Epstein, A. M. (2016). Readmissions, observation, and the Hospital Readmissions Reduction Program. *New England Journal of Medicine*, 374(16), 1543-1551.
40. Bradley, E. H., & Curry, L. (2023). The social determinants of readmissions: A conceptual framework. *Health Affairs*, 42(5), 678-685.
41. Cromley, E. K., & McLafferty, S. L. (2022). *GIS and public health* (3rd ed.). Guilford Press.
42. Matthews, S. A., & Yang, T. C. (2023). Geospatial approaches to health disparities research. *American Journal of Public Health*, 113(S2), S128-S135.
43. Krieger, N., & Chen, J. T. (2022). Geocoding and monitoring of US socioeconomic inequalities in mortality and cancer incidence. *American Journal of Epidemiology*, 191(6), 1050-1061.
44. Hastie, T., Tibshirani, R., & Friedman, J. (2021). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). Springer.
45. Beam, A. L., & Kohane, I. S. (2018). Big data and machine learning in health care. *JAMA*, 319(13), 1317-1318.
46. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5-32.
47. Chen, T., & He, T. (2022). XGBoost for healthcare prediction: A comprehensive review. *Journal of Biomedical Informatics*, 127, 103989.
48. Liaw, A., & Wiener, M. (2022). Classification and regression by randomForest. *R News*, 2(3), 18-22.
49. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1-35.
50. Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. *Advances in Neural Information Processing Systems*, 29, 3315-3323.
51. Link, B. G., & Phelan, J. C. (1995). Social conditions as fundamental causes of disease. *Journal of Health and Social Behavior*, 35, 80-94.
52. Phelan, J. C., Link, B. G., & Tehranifar, P. (2010). Social conditions as fundamental causes of health inequalities: Theory, evidence, and policy implications. *Journal of Health and Social Behavior*, 51(1_suppl), S28-S40.

53. Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2), 179-211.
54. Ajzen, I. (2020). The theory of planned behavior: Frequently asked questions. *Human Behavior and Emerging Technologies*, 2(4), 314-324.
55. Carpenter, C. J. (2010). A meta-analysis of the effectiveness of health belief model variables in predicting behavior. *Health Communication*, 25(8), 661-669.
56. van Walraven, C., Dhalla, I. A., Bell, C., Etchells, E., Stiell, I. G., Zarnke, K., ... & Forster, A. J. (2010). Derivation and validation of an index to predict early death or unplanned readmission after discharge from hospital to the community. *Canadian Medical Association Journal*, 182(6), 551-557.
57. Donzé, J., Aujesky, D., Williams, D., & Schnipper, J. L. (2013). Potentially avoidable 30-day hospital readmissions in medical patients: Derivation and validation of a prediction model. *JAMA Internal Medicine*, 173(8), 632-638.
58. Kansagara, D., & Englander, H. (2022). Thirty-day readmission prediction: A systematic review and meta-analysis. *Journal of Hospital Medicine*, 17(6), 435-445.
59. Damery, S., & Combes, G. (2023). The predictive performance of readmission risk models: A meta-analysis. *Health Services Research*, 58(4), 815-829.
60. Collins, G. S., Reitsma, J. B., Altman, D. G., & Moons, K. G. (2015). Transparent Reporting of a multivariable prediction model for Individual Prognosis Or Diagnosis (TRIPOD): The TRIPOD statement. *Annals of Internal Medicine*, 162(1), 55-63.
61. Elixhauser, A., Steiner, C., Harris, D. R., & Coffey, R. M. (1998). Comorbidity measures for use with administrative data. *Medical Care*, 36(1), 8-27.
62. Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLoS ONE*, 10(3), e0118432.
63. Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., ... & Lee, S. I. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2(1), 56-67.
64. Ljungman, P. L. S., & Mittleman, M. A. (2023). Air pollution and cardiovascular disease: A comprehensive review. *Circulation*, 147(5), 387-401.
65. Vyas, D. A., Eisenstein, L. G., & Jones, D. S. (2020). Hidden in plain sight: Reconsidering the use of race correction in clinical algorithms. *New England Journal of Medicine*, 383(9), 874-882.

66. Shortliffe, E. H., & Sepúlveda, M. J. (2023). Clinical decision support in the era of artificial intelligence. *JAMA*, 330(14), 1347-1356.
67. Ose, D., & Krelle, H. (2023). Natural language processing for social determinants of health extraction from clinical notes. *Journal of the American Medical Informatics Association*, 30(6), 1078-1087.
68. Chawla, N. V., & Osei, K. (2024). Leveraging hidden patterns in open-ended community health workers' notes to improve prediction of patient readmission. *IEEE International Conference on Healthcare Informatics*, 2024, 134-143.