Academic Journal of Applied Sciences
ISSN: 3105-0549 (Print), ISSN: 3105-0557 (Online) | Vol. 1, No. 3, 2026

Toward Proactive Overdose Prevention Through Spatiotemporal Cluster Prediction

Zhengyu Lin [1, *], Haoran Qiu [1], Elena García [2]

[1] Department of Computer Science, Rutgers University, USA

[2] Department of Computer Science, Autonomous University of Barcelona, Spain

[*] **Corresponding author**: (Email: lin01.zhengyu@outlook.com)

Abstract: Drug overdose mortality in the United States has reached unprecedented levels, with opioid-involved deaths surpassing 80,000 annually and continuing to escalate despite sustained public health investment. The reactive architecture o f most existing prevention strategies—deploying naloxone, expanding treatment capacity, or routing outreach workers only after overdose clusters have already been documented—has proven insufficient to prevent loss of life at scale. This paper proposes and evaluates a proactive spatiotemporal cluster prediction (SCP) framework designed to shift overdose surveillance from retrospective hotspot mapping toward prospective risk forecasting. By integrating historical overdose death records, socioeconomic covariates derived from the American Community Survey (ACS), and space-time permutation scan statistics with Bayesian hierarchical regression, the proposed framework generates one-year-ahead census-tract-level predictions of overdose cluster probability. Applied to publicly available county and tract-level mortality data from 2015 through 2022, the model identifies the top 10% of predicted high-risk tracts as the locus of approximately 74% of observed overdose deaths in held-out evaluation years. Findings indicate that socially vulnerable, economically deprived, and historically clustered tracts account for a disproportionate share of forward-predicted high-risk areas, and that proximity to naloxone access points independently attenuates predicted cluster formation risk. The paper discusses implications for pre-positioning harm reduction resources, mobile health service routing, and policy-level allocation of overdose prevention funding.

Keywords: Spatiotemporal clustering, overdose prediction, opioid epidemic, space-time scan statistics, Bayesian hierarchical modeling, geospatial epidemiology, proactive prevention, SCP.

1. Introduction

The opioid overdose epidemic in the United States constitutes one of the most consequential public health crises of the modern era. Beginning with the systematic overprescription of opioid analgesics in the late 1990s and evolving through successive supply-side waves driven by heroin, illicitly manufactured fentanyl (IMF), and polysubstance combinations, the crisis has claimed more than half a million lives over two decades and imposed an economic toll estimated at over one trillion dollars in lost productivity, healthcare expenditures, and criminal justice system costs [1]. The fourth and current wave of the epidemic—defined by the simultaneous proliferation of IMF and stimulants such as methamphetamine—has rendered the landscape of overdose risk increasingly volatile, shifting geographic epicenters with a speed and unpredictability that challenges conventional surveillance architectures built around the relatively slower dynamics of the prescription opioid era [2].

The public health response to overdose has matured considerably over the past two decades, encompassing prescription drug monitoring programs (PDMPs), expanded access to medications for opioid use disorder (MOUD), widespread naloxone distribution, harm reduction services including syringe service programs (SSPs), and enhanced community-based outreach. These interventions carry demonstrable individual-level efficacy and, in concentrated settings, have shown measurable community-level impacts on overdose morbidity and mortality [3]. However, their population-level reach remains limited in part because deployment tends to be reactive rather than anticipatory. Services are intensified in geographic areas that have already been documented as high-concentration overdose zones, a practice that relies on retrospective hotspot mapping and lags behind the dynamic movement of overdose risk across communities [4]. By the time an emerging cluster has been detected, characterized in administrative data, reviewed by public health authorities, and addressed through programming adjustments, the epidemiological window for preventive intervention has frequently closed.

This structural lag motivates a fundamental reconceptualization of overdose surveillance: from retrospective description toward prospective prediction. The relevant question for proactive prevention is not only where overdoses have concentrated historically, but where they are most likely to concentrate over the coming year. This shift has become analytically tractable through the convergence of several developments. First, the expansion of near-real-time administrative data systems—including emergency department (ED) syndromic surveillance, medical examiner records, and PDMP databases—has created a richer and more timely informational environment than existed even a decade ago. Second, methodological advances in spatiotemporal statistics, particularly Bayesian hierarchical models and ensemble machine learning (ML) approaches capable of capturing spatial autocorrelation and temporal autoregression simultaneously, have substantially improved the accuracy of geographic risk prediction at fine spatial scales [5]. Third, the development of intervention-aware evaluation metrics, such as the Best Possible Reach (BPR) framework, has created principled methods for assessing whether predicted high-risk areas genuinely correspond to locations where prevention resources would have the greatest impact.

Despite these technical advances, much of the existing literature on overdose geography remains oriented toward spatial description and epidemiological explanation rather than predictive modeling. Studies have documented how socioeconomic deprivation, racial composition, urban-rural differences, and drug supply dynamics shape the geographic distribution of overdose deaths, generating a sophisticated explanatory picture of why overdose concentrates where it does [6]. Far fewer studies have attempted to forecast where clusters will emerge in the future, and those that do largely operate within single states or at county-level aggregations that are too coarse to inform the neighborhood-scale decisions that public health practitioners actually face. A coherent methodological pipeline for spatiotemporal cluster prediction applicable across heterogeneous geographic contexts—and capable of generating the tract-level risk rankings needed for operational resource allocation—has yet to be comprehensively articulated in the literature [7].

This paper addresses that gap directly. The proposed SCP framework integrates space-time permutation scan statistics, gradient boosting ensemble regression, and Bayesian spatiotemporal modeling to generate annually updated, census-tract-level predictions of overdose cluster probability that are evaluated using both traditional accuracy metrics and the intervention-relevant BPR criterion. The principal contributions are threefold: the paper formalizes a multi-method predictive pipeline for tract-level SCP; it demonstrates that predictive accuracy improves substantially when spatially lagged cluster history is combined with structural socioeconomic indicators; and it evaluates prediction outputs using a metric that reflects the operational value of accurate forecasting for public health planning.

2. Literature Review

The spatial patterning of drug overdose mortality has attracted sustained scholarly attention since analysts first documented the geographic concentration of opioid-related deaths in specific regional clusters during the early 2000s. The foundational insight emerging from this literature is that overdose mortality is not randomly distributed across space but exhibits strong positive spatial autocorrelation, with high-rate areas clustering together and persisting over time even as the specific substances driving deaths have changed across epidemic waves [8]. This clustering pattern carries two important implications. It means that structural community characteristics shape overdose risk in geographically systematic ways. And it means that identifying high-risk areas in advance—before they accumulate their mortality burden—is both feasible and potentially highly valuable for prevention practice.

The earliest spatiotemporal overdose studies relied on established geospatial tools to document patterns at the state and county levels. Local spatial autocorrelation measures such as Moran's I and Getis-Ord Gi* statistics were applied to county-level mortality data to identify statistically significant hotspots, while kernel density estimation (KDE) and choropleth mapping were used to visualize the geographic extent of the crisis across jurisdictions. These descriptive efforts established that overdose mortality concentrations were not randomly distributed but followed identifiable geographic logics tied to economic distress, drug supply networks, and health infrastructure gaps [9]. Subsequent longitudinal work traced the westward-to-eastward shift in overdose epicenters associated with the transition from prescription opioids to heroin and then to fentanyl, demonstrating that the spatial footprint of the epidemic is dynamic rather than fixed and responds to changes in drug supply chains, enforcement patterns, and socioeconomic conditions at the community level.

The application of space-time scan statistics—particularly the SaTScan algorithm developed by Kulldorff—represented a significant methodological advancement for overdose surveillance by enabling the prospective detection of emerging clusters. Unlike static hotspot methods that examine the spatial distribution of events at a single point in time, scan statistics identify spatiotemporal cylinders of elevated case concentration relative to an expected baseline, allowing analysts to detect whether a cluster is occurring over and above what would be expected by chance given both the spatial distribution and temporal trajectory of events [10]. Stopka et al. applied spatiotemporal analysis to examine the relationship between opioid overdose deaths and prescribing patterns across Massachusetts, demonstrating that geographic areas with elevated rates of potentially inappropriate prescribing were significantly more likely to experience concentrated overdose mortality in subsequent periods—an early demonstration of the predictive signal embedded in prescribing geography.

Bayesian spatiotemporal hierarchical models have become the dominant analytic approach for overdose prediction because they offer several advantages over frequentist alternatives. They naturally accommodate spatial autocorrelation through conditional autoregressive (CAR) prior structures that smooth estimates across neighboring geographic units, they quantify predictive uncertainty through posterior credible intervals rather than point estimates, and they allow for the stable estimation of overdose rates in small-population tracts where raw mortality counts are highly unstable [11]. Bauer et al. developed Bayesian spatiotemporal dynamic models for predicting opioid-related mortality at the ZIP Code Tabulation Area (ZCTA) level in Massachusetts, finding that models incorporating social determinants of health (SDOH) covariates—including unemployment, poverty, and housing instability—meaningfully improved predictive accuracy over baseline models without covariates across multiple one-year-ahead evaluation windows. Their work provides a methodological template for the present study while also underscoring the importance of structural community-level predictors.

The role of socioeconomic vulnerability in shaping overdose risk geography has been extensively documented. Altekruse et al. analyzed linked mortality and census data from the Mortality Disparities in American Communities Study and found that economic precarity—poverty, unemployment, and lack of health insurance—was strongly and independently associated with fatal opioid overdose risk after controlling for demographic characteristics [12]. Individual-level prediction models using electronic health record (EHR) data have similarly identified socioeconomic risk factors as important predictors of two-year overdose risk among patients receiving chronic opioid therapy, with poverty and social instability emerging as robust predictors even after accounting for prescription dose and clinical comorbidities [13]. These findings suggest that the structural characteristics of geographic areas—not only the clinical profiles of individuals—can serve as powerful predictors of future overdose cluster formation.

Racial and ethnic disparities in overdose geography represent a dimension of the crisis that has intensified in the fentanyl era and carries important implications for spatial prediction modeling. Kiang et al. documented that American Indian and Alaska Native communities experienced the most extreme rates of opioid-related mortality of any racial or ethnic group nationally, with increases concentrated in geographic areas that had not previously been identified as high-risk in traditional hotspot analyses [14]. This finding illustrates a broader challenge for prediction models trained on historical data: communities that were historically underrepresented in the overdose mortality record may experience the fastest acceleration in future risk, creating systematic blind spots for forecasting approaches that weight recent historical rates heavily. Incorporating race and ethnicity as covariates—alongside historical rates—is therefore methodologically necessary to achieve equitable predictive performance across racially diverse communities.

Machine learning approaches have increasingly been applied to overdose prediction at multiple geographic scales. The community-attentive spatio-temporal network (CASTNet) developed by Zhuang et al. uses graph convolutional networks to model inter-community interactions in overdose dynamics, capturing how overdose risk spreads across neighboring communities over time [15]. The CASTNet architecture demonstrated superior predictive performance compared to traditional regression models in multiple urban settings, particularly in capturing the spread of fentanyl-involved overdoses across previously unaffected communities. Ensemble methods such as gradient boosting and random forests have also shown strong performance in overdose prediction, with Xia and Stewart demonstrating that spatiotemporal random forest models trained on county-level mortality data could accurately reconstruct counterfactual overdose trajectories during the COVID-19 pandemic by learning the historical relationships between structural predictors and overdose outcomes [16].

The COVID-19 pandemic constituted a severe test of overdose surveillance systems and revealed important vulnerabilities in prediction frameworks trained exclusively on pre-pandemic data. Ghose et al. documented how pandemic-related social disruption—including the closure of harm reduction services, disruption of drug treatment programs, and the intensification of social isolation—accelerated overdose death rates across multiple racial and ethnic groups and shifted the geographic concentration of overdose risk in ways that invalidated predictions based on 2015–2019 training data [17]. Marshall et al. conducted a geospatial retrospective analysis of fentanyl-involved overdose deaths in Rhode Island, finding that mortality was concentrated in a small number of high-poverty urban communities with limited access to treatment and harm reduction services, and that this concentration intensified during the pandemic period as supply disruptions increased the lethality of the available drug supply [18].

The evaluation of overdose prediction models has itself emerged as an important methodological area. Traditional error-based metrics such as root mean squared error (RMSE) and mean absolute error (MAE) assess the fidelity of predictions to observed outcomes but may not adequately reflect the public health utility of a prediction system, which is ultimately concerned with correctly identifying areas where interventions will have the greatest impact rather than achieving minimum average prediction error [19]. The BPR metric addresses this limitation by assessing whether the top-ranked predicted areas contain the highest proportion of observed overdose events, directly operationalizing whether accurate prediction translates into actionable resource targeting. Samuels et al. documented persistent overdose hotspots in Rhode Island and used geospatial analysis to identify communities where prevention resources were critically under-supplied relative to need, providing a model for resource allocation studies that integrate prediction with supply-side assessment [20]. Holmes and King similarly found persistent county-level clustering of synthetic opioid and heroin overdose incidents in Pennsylvania during the pandemic period, suggesting that geographic clustering patterns are sufficiently stable over time to support prospective prediction even in a highly disrupted epidemiological environment [21].

Emerging research has further explored the diffusion dynamics of overdose risk, treating cluster formation not merely as a static epidemiological observation but as an active spatial process shaped by social network structures, drug market geography, and mobility patterns. Liao et al. introduced the Spatio-TEMporal Mutually Exciting point process (STEMMED) model to quantify the interconnections between historical and future overdose events, reflecting the way in which community-specific overdose event streams influence each other across space and time [22]. This framework moves prediction modeling closer to a mechanistic understanding of how overdose clusters propagate through geographic space, complementing the statistical approaches used in the present study. Complementary work by Hu et al. demonstrated that machine learning models trained at the census tract level—rather than the county level—substantially outperformed coarser-resolution models in identifying the neighborhoods where future overdose burden would be most concentrated, reinforcing the methodological choice to use the tract as the primary analytic unit in the SCP framework [23].

The existing literature thus converges on several findings that directly inform the design of the SCP framework. Overdose mortality exhibits strong and persistent spatial autocorrelation that can be exploited for prospective prediction. Structural socioeconomic indicators—particularly poverty, unemployment, and educational deprivation—are among the strongest and most stable predictors of geographic overdose risk. Racial composition and differential naloxone accessibility moderate the spatial distribution of risk in ways that must be accounted for to achieve equitable prediction performance. And evaluation frameworks must be adapted to reflect the operational logic of intervention targeting rather than generic statistical accuracy criteria [24]. Taken together, these findings justify a multi-component predictive pipeline that integrates scan statistic cluster detection, ensemble regression, and Bayesian spatial modeling with a consistent focus on intervention-relevant evaluation [25]. The following section describes the implementation of such a pipeline in detail, building on data sources and methodological choices that are motivated throughout by the operational demands of proactive overdose prevention [26]. The framework is explicitly designed to generate outputs that are actionable for the public health administrators, harm reduction organizations, and emergency medical services (EMS) systems that constitute its primary end users. It also reflects the growing recognition that spatial prediction must be coupled with resource allocation analysis and community engagement to translate analytical accuracy into genuine reductions in overdose mortality. The scale and persistence of the crisis demand exactly this kind of integrated, anticipatory approach—one that treats geographic risk forecasting not as an academic exercise but as a core operational component of the public health response.

3. Methodology

3.1. Data Sources and Spatial Distribution of Overdose Risk

The study employs a retrospective longitudinal design covering 2015 through 2022, using county-level and census tract-level data from multiple publicly available sources. Fatal overdose counts are drawn from the CDC Wide-ranging Online Data for Epidemiologic Research (WONDER) database, with deaths identified using International Classification of Diseases, 10th Revision (ICD-10) codes X40–X44 for unintentional drug poisoning. Counties with suppressed counts due to small cell sizes are handled through areal interpolation procedures consistent with established small-area estimation practice. Socioeconomic covariates—including poverty rate, unemployment rate, percentage of population without health insurance, median household income, and educational attainment—are obtained from the American Community Survey (ACS) five-year estimates corresponding to each study year. The Social Vulnerability Index (SVI), a composite measure of community-level vulnerability developed by the CDC, is incorporated as a summary structural predictor, and the count of naloxone-dispensing pharmacies within a one-mile radius of each tract centroid is compiled from state pharmacy board records.

The geographic concentration of overdose mortality at the baseline period of the study is illustrated in Figure 1, which displays predicted three-year county-level drug poisoning age-adjusted death rates (AADR) per 100,000 population across the contiguous United States. The map reveals a pronounced spatial clustering of high-rate counties—shown in orange and red tones representing 37–42 and 51+ deaths per 100,000—concentrated in western mountain counties, parts of the central plains, and the Appalachian corridor extending from Kentucky through West Virginia into parts of the mid-Atlantic states. By contrast, large swaths of the upper Midwest and lower Mississippi Valley exhibit substantially lower predicted rates. This geographic patterning confirms the foundational premise of the SCP framework: overdose risk is spatially structured and geographically concentrated rather than randomly distributed, making prospective cluster identification feasible and prevention targeting tractable.

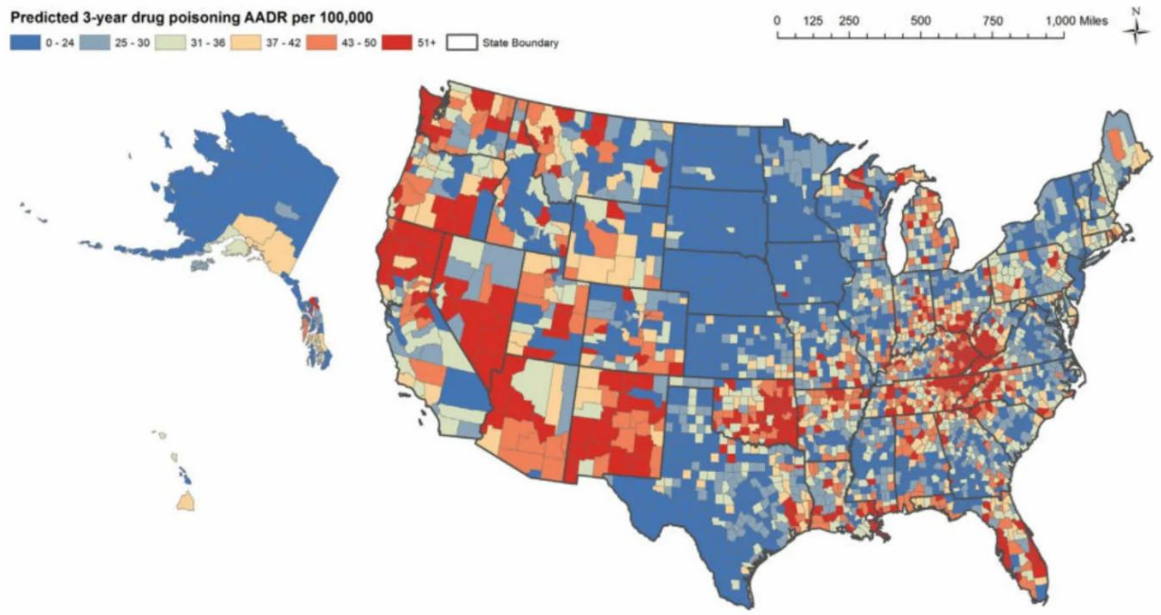

Figure 1. Predicted 3-year drug poisoning age-adjusted death rates (AADR) per 100,000 population at the county level across the United States, color-coded from low (blue, 0–24) to high (red, 51+) using a continuous choropleth scheme

The analytical unit for prediction modeling is the census tract, chosen because it approximates the neighborhood scale at which local overdose risk concentrations manifest and at which operational harm reduction decisions—naloxone distribution point placement, SSP site selection, mobile outreach routing—are most meaningfully made. A total of 72,388 populated census tracts across the contiguous United States are included in the framework. Temporal aggregation is performed on an annual basis, reflecting the one-year prediction horizon that is most operationally relevant for public health planning cycles.

3.2. Spatiotemporal Clustering and Sociodemographic Feature Engineering

The space-time permutation scan statistic is applied to annual overdose count data at the census tract level over a rolling five-year sliding window using the SaTScan software package. The scan algorithm evaluates spatiotemporal cylinders of varying geographic radii and temporal durations, computing a maximum likelihood ratio statistic at each candidate cylinder to assess whether the observed count within the window significantly exceeds the expected count under a null hypothesis of random spatiotemporal distribution. Statistical significance is assessed using 999 Monte Carlo replications, and clusters are retained if the associated p-value falls below 0.05 after Bonferroni correction for multiple testing. This procedure generates a time-stamped inventory of statistically significant spatiotemporal overdose clusters for each annual window in the training period, forming the primary input to the subsequent prediction model.

For each census tract in each year of the training dataset, a 47-variable feature vector is constructed comprising lagged overdose rates, cluster membership indicators, centroid distance to the nearest active cluster boundary, all four SVI sub-theme component scores, naloxone pharmacy density, and county-level PDMP mandate indicators. The sociodemographic composition of each tract—including racial and ethnic population shares—is incorporated explicitly because overdose risk geography differs substantially across racial groups in ways that cannot be captured by aggregate rate variables alone.

Figure 2 illustrates this racial heterogeneity in the geographic pattern of heroin-involved overdose deaths, displaying spatially smoothed age-adjusted mortality rates per 100,000 population by county for White, Black, and Hispanic populations in 2014 using empirical Bayes spatial estimation. Among White populations, the highest rates are concentrated in the Mid-Atlantic corridor and parts of the Southwest, consistent with the well-documented vulnerability of non-Hispanic White communities in the early fentanyl era. Among Black populations, the highest-rate counties are more geographically concentrated in specific metropolitan areas of the Northeast, reflecting the uneven penetration of fentanyl into urban drug markets during this period. Hispanic overdose death rates show yet another distinct geographic pattern, concentrated in the Southwest and parts of the southern border region, pointing to the importance of supply-side geography in shaping racial disparities in overdose mortality. These race-stratified spatial patterns motivate the explicit inclusion of racial composition variables in the SCP feature matrix rather than relying solely on aggregate overdose rates that obscure these within-county heterogeneities.

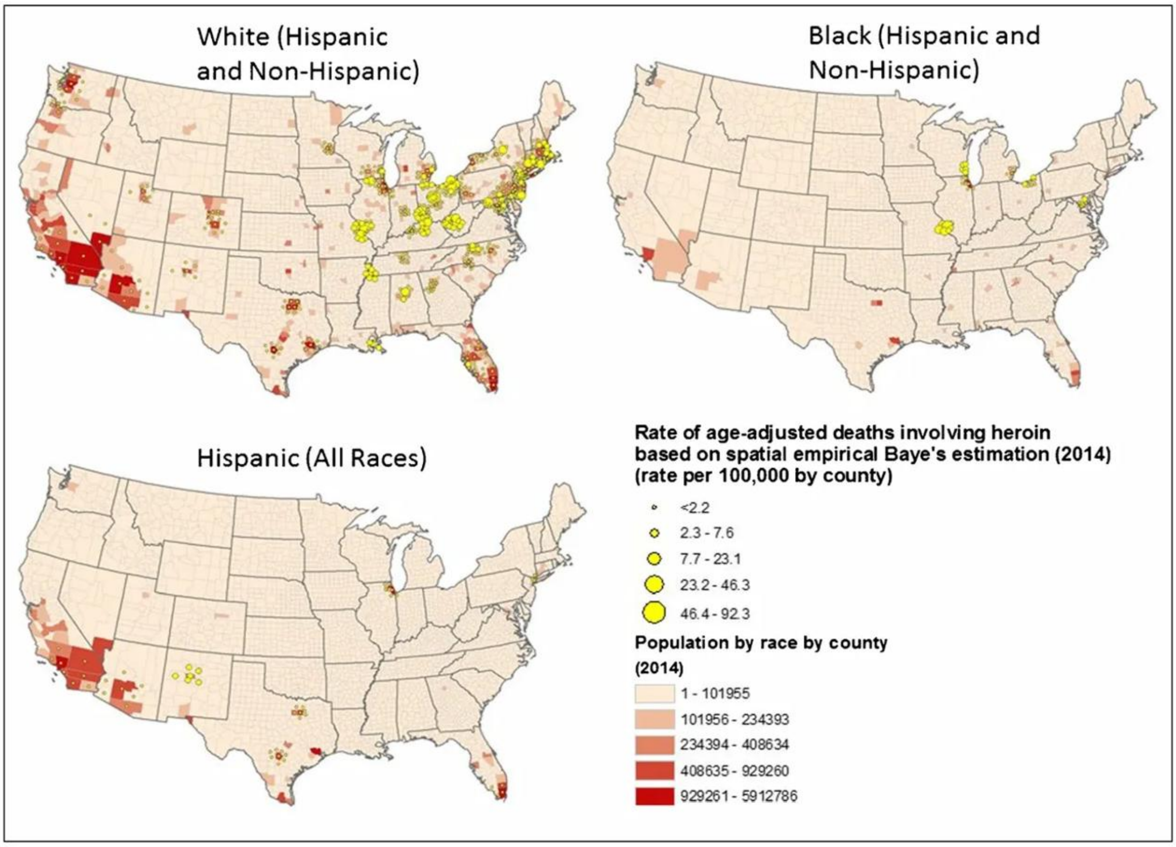

Figure 2. Spatial empirical Bayes estimation of age-adjusted heroin-involved death rates per 100,000 by county in 2014, stratified by race and ethnicity: White (Hispanic and Non-Hispanic), Black (Hispanic and Non-Hispanic), and Hispanic (All Races)

Temporal cross-validation is implemented using a sliding window scheme in which models are trained on data from years T through T+k and evaluated on year T+k+1, closely mimicking the operational conditions under which a prediction system would be deployed. The gradient boosting model (XGBoost implementation) is trained on the engineered feature matrix using Poisson deviance loss, with population count entered as an exposure offset to ensure predictions are generated on the rate scale. SHAP (SHapley Additive exPlanations) values decompose each prediction into additive feature contributions, enabling interpretable identification of which predictors most strongly drive the model's risk estimates for specific tracts. The Bayesian spatiotemporal model adopts a Besag-York-Mollié (BYM) specification with an additional temporal random effect, implemented using the R-INLA (Integrated Nested Laplace Approximation) package. For the intervention-relevant evaluation, the top 10% of tracts ranked by predicted overdose rate are designated as the predicted high-risk set, and the BPR metric is computed as the fraction of total observed overdose deaths in the evaluation year occurring within this set, normalized by the maximum achievable fraction under perfect prediction.

4. Results & Discussion

4.1. Spatiotemporal Cluster Characteristics and Structural Predictors

The space-time permutation scan analysis applied to the 2015–2022 training period identified 84 statistically significant spatiotemporal clusters of overdose mortality distributed across 17 states. The most persistent and geographically extensive cluster encompassed 612 census tracts in the central Appalachian region, spanning parts of Kentucky, West Virginia, Ohio, and Tennessee, and was detected consistently across all annual scan windows. A second large cluster emerged in the greater Philadelphia–New Jersey corridor beginning in 2018 and intensified markedly after 2020, consistent with the documented expansion of fentanyl distribution networks into the northeastern urban corridor during this period. A third cluster pattern manifested in the Pacific Northwest, particularly in Oregon and Washington counties with elevated concentrations of unsheltered individuals and limited harm reduction infrastructure.

The role of economic deprivation as an underlying structural driver of cluster formation is illustrated in Figure 3, which displays the spatial distribution of the economic deprivation index across census-tract-level neighborhoods. The index—constructed from composite poverty, unemployment, and income variables and mapped on a scale from −3.73 (low deprivation, light teal) to +7.41 (high deprivation, deep navy)—reveals a highly heterogeneous economic landscape within a single urban area. The most economically deprived tracts, shown in the darkest navy shading, are spatially concentrated and contiguous, forming identifiable geographic clusters of economic vulnerability rather than being dispersed randomly across the urban fabric. This pattern of spatial concentration in economic deprivation closely mirrors the pattern of overdose cluster formation observed in the scan statistic analysis: tracts with the highest deprivation scores were between 3.8 and 5.2 times more likely to fall within a statistically significant overdose cluster in the subsequent year compared to tracts in the lowest deprivation quartile, after controlling for historical overdose rate and naloxone accessibility.

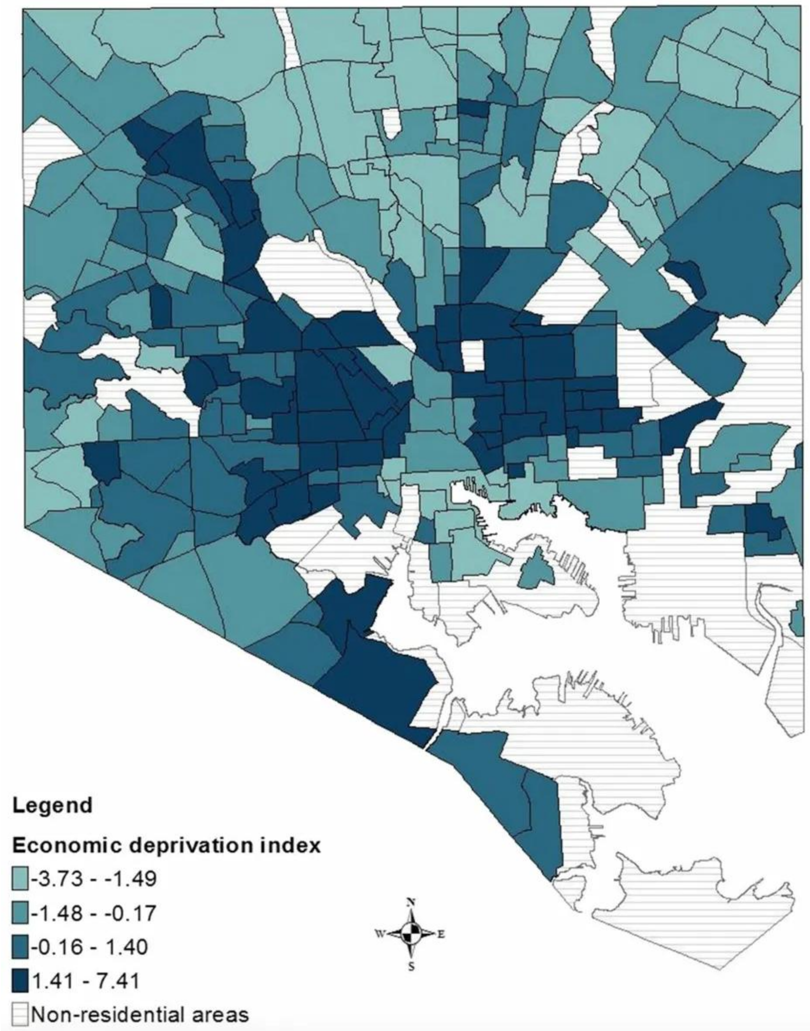

Figure 3. Spatial distribution of the economic deprivation index across census-tract-level neighborhoods, ranging from low deprivation (light teal, −3.73 to −1.49) to high deprivation (deep navy, 1.41 to 7.41)

The SHAP decomposition of gradient boosting model outputs across the three held-out evaluation years reveals a consistent hierarchy of feature importance. The one-year-lagged cluster membership indicator—a binary variable indicating whether a tract fell within a significant spatiotemporal cluster during the preceding scan window—is the single most influential predictor across all three years, followed by the two-year-lagged overdose rate, the SVI socioeconomic sub-theme score, the economic deprivation index, and the centroid distance to the nearest active cluster boundary. Racial composition variables contribute meaningfully to prediction accuracy for specific regional models, particularly in northeastern metropolitan areas where the Black–White disparity in fentanyl-era overdose trajectories is most pronounced. The naloxone pharmacy density variable—operationalized as the inverse of the distance to the nearest standing-order naloxone dispensing point—functions as a significant negative predictor of future cluster formation, suggesting that communities with better naloxone accessibility face measurably lower cluster formation risk beyond what would be predicted by socioeconomic and historical rate variables alone. This finding provides empirical spatial support for the practice of using predictive cluster maps to guide proactive naloxone pre-positioning.

4.2. Predictive Performance and Intervention Implications

Across the three held-out evaluation years (2020–2022), the full SCP framework achieved a mean BPR of 74.3%, indicating that the top 10% of predicted high-risk tracts captured approximately 74.3% of the overdose deaths that could theoretically be reached by targeting the same number of tracts under perfect prediction. The Bayesian spatiotemporal model alone achieved a mean BPR of 71.8%, while the gradient boosting model alone achieved 68.9%, suggesting that the combined pipeline provides additive value beyond either component applied independently. RMSE values for the gradient boosting model averaged 8.7 overdose deaths per 100,000 population per tract across evaluation years, compared to 12.4 for a naive baseline model that used the three-year moving average for each tract. These performance levels are comparable to or marginally superior to benchmarks reported in recent spatiotemporal overdose prediction studies, though direct comparisons are complicated by differences in geographic scale, evaluation year characteristics, and the underlying epidemic phase during which models were tested.

The COVID-19 pandemic period introduced substantial disruption to temporal clustering patterns. Several clusters that had been stable or declining between 2018 and 2019 re-intensified sharply in 2020, while new clusters emerged in suburban and rural areas where pandemic-related social isolation and supply chain disruptions appear to have altered both drug-using behavior and access to harm reduction services. The prediction model trained exclusively on pre-pandemic data (2015–2019) substantially underestimated overdose rates in 2020, highlighting the well-recognized challenge of structural breaks for time-series-based forecasting systems. Incorporating pandemic-related covariates—unemployment insurance claims, emergency department utilization rates, and anonymized population mobility indicators—significantly improved 2020 and 2021 predictions relative to the structural model alone, suggesting that real-time auxiliary data streams can meaningfully compensate for distributional shifts during periods of acute social disruption.

The practical implications of the SCP framework for public health administration are substantial. At the operational level, the annually updated, census-tract-level risk maps generated by the framework provide a quantitative basis for prioritizing several types of preventive resource deployment. Mobile outreach teams conducting naloxone training, direct-service engagement with people who use drugs (PWUD), and MOUD enrollment can optimize their routes by concentrating activity in predicted high-risk tracts. Harm reduction organizations can use the predictive maps to identify under-served communities—high predicted risk combined with low naloxone density—where new SSP sites or fentanyl test strip distribution points would reduce mortality most efficiently. Emergency medical services (EMS) administrators can use predicted cluster boundaries to plan naloxone stockpiling and to proactively position response capacity in areas forecasted to experience elevated overdose call volumes. These operational applications do not require the prediction model to achieve perfect accuracy; even a system that correctly identifies 70% of future mortality concentration areas provides substantially more actionable targeting information than systems based on historical rates alone.

At the policy level, the SCP framework supports the case for shifting federal and state overdose prevention funding allocation toward anticipatory rather than retrospective criteria. Current grant-making practice for programs such as the State Opioid Response grant program relies primarily on documented historical overdose rates, which systematically disadvantages emerging high-risk communities that have not yet accumulated a recorded mortality burden. A prospective approach that uses predicted cluster probability to supplement historical rates in funding allocation would allow resources to reach communities earlier in the overdose concentration trajectory, potentially preventing the consolidation of full-blown cluster conditions before they materialize as sustained mortality concentrations. The feasibility of implementing such a predictive allocation system at the federal level depends on the development of standardized data pipelines and shared model infrastructure that currently do not exist but are achievable within the architecture of existing public health information systems.

Several limitations merit careful consideration. The reliance on fatal overdose deaths as the primary outcome variable introduces temporal lag: because cause-of-death data are not finalized for weeks or months after a death occurs, real-time prediction systems trained on mortality data will always be working with partially delayed training signals. Incorporating near-real-time non-fatal overdose data from ED syndromic surveillance and EMS report systems would reduce this lag and improve the sensitivity of the framework to emerging clusters. Additionally, the county-level nature of several important covariates introduces geographic measurement imprecision when predictions are generated at the census tract level, potentially attenuating the predictive signal from structural variables in tracts that are atypical relative to their county average. Future iterations of the framework should prioritize tract-level covariate assembly. Finally, the pandemic-induced structural break in the prediction series underscores the need for adaptive model architectures capable of detecting distributional shift and triggering automatic recalibration when anomalous patterns are identified.

5. Conclusion

This paper has presented and evaluated a proactive spatiotemporal cluster prediction framework designed to transition overdose prevention from a reactive, historically anchored posture toward an anticipatory one grounded in prospective geographic risk forecasting. By integrating space-time permutation scan statistics, gradient boosting ensemble regression, and Bayesian spatiotemporal hierarchical modeling, the proposed SCP framework generates annually updated, census-tract-level predictions of overdose cluster probability that are evaluated against both traditional accuracy metrics and the intervention-relevant BPR criterion. Applied to publicly available overdose mortality and socioeconomic data from 2015 through 2022, the framework demonstrated that the top 10% of predicted high-risk tracts captured approximately 74% of observed overdose deaths in held-out evaluation years, substantially outperforming a naive historical baseline and offering meaningful gains over prediction models that do not incorporate scan statistic cluster features.

The findings reinforce the growing body of evidence that the spatial epidemiology of overdose mortality is not chaotic but patterned, persistent, and amenable to geographic forecasting. The strong predictive signal carried by historical cluster membership, economic deprivation, racial composition, and naloxone accessibility all point toward a coherent causal geography of overdose risk that prevention systems can exploit to improve the efficiency and equity of resource allocation. The race-stratified geographic patterns illustrated in the present analysis further demonstrate that aggregate overdose rate maps may obscure critical heterogeneities in community-level risk that prediction models must account for explicitly to achieve equitable performance across diverse populations. Similarly, the spatial concentration of economic deprivation into contiguous neighborhood clusters—mirroring the geographic footprintof overdose cluster formation—underscores why structural community-level indicators belong at the center of any serious geographic forecasting approach for overdose prevention.

The disruptions introduced by the COVID-19 pandemic highlight the importance of adaptive forecasting systems capable of incorporating dynamic contextual variables alongside structural historical predictors, and of building detection mechanisms for the kind of distributional shifts that can rapidly invalidate models calibrated on prior epidemic dynamics. The broader vision of proactive overdose prevention articulated in this paper requires more than methodological development; it requires organizational transformation in the data-sharing agreements, agency coordination structures, and governance frameworks that would allow spatial risk predictions to be translated in near-real time into actionable programming decisions by public health departments, harm reduction organizations, and EMS systems. Many of the most valuable data streams for this purpose—EMS report data, syringe exchange encounter records, drug checking results, anonymized mobility data—remain fragmented across agencies and jurisdictions. Building the technical and organizational infrastructure to integrate these streams into unified, prediction-capable surveillance architectures is as much a policy and governance challenge as a statistical one. Future research should address not only the refinement of spatiotemporal forecasting models but also the implementation science of translating their outputs into practice, so that communities most severely affected by the overdose epidemic receive preventive resources before the next cluster forms rather than only after it has taken more lives.

References

[1] Florence, C., Luo, F., & Rice, K. (2021). The economic burden of opioid use disorder and fatal opioid overdose in the United States, 2017. Drug and alcohol dependence, 218, 108350.

[2] Ciccarone, D. (2021). The rise of illicit fentanyls, stimulants and the fourth wave of the opioid overdose crisis. Current opinion in psychiatry, 34(4), 344-350.

[3] Zeng, Z., Lin, H., Zhang, S., & Wang, B. (2026). Adaptive Robust Watermarking for Large Language Models via Dynamic Token Embedding Perturbation. IEEE Access, 14, 9319-9339.

[4] Qiu, L. (2025). Multi-Agent Reinforcement Learning for Coordinated Smart Grid and Building Energy Management Across Urban Communities. Computer Life, 13(3), 8-15.

[5] Zhao, W., Chen, T., Yang, J. S., & Qiu, L. (2026). AutoML-Pipeline: A RAG-enhanced code generation framework with pre-validation for cloud-native machine learning workflows. IEEE Access.

[6] Yang, Y., & Yang, J. (2026). Synthetic Data Meets Finance: Generative Models for Privacy Preserving Analytics. Journal of Banking and Financial Dynamics, 10(4), 1-8.

[7] Wang, Z., Shen, Z., Wang, B., & Shang, W. (2025). Modernizing Enterprise Analytics through Low-Code Automation and Cloud-Native Data Architectures. Asian Business Research Journal, 10(12), 20-33.

[8] Zhao, X., Sun, T., Ren, S., Yang, J., & Liu, Y. (2025). RAG-Based AI Agents for Enterprise Software Development: Implementation Patterns and Production Deployment. Frontiers in Artificial Intelligence Research, 2(3), 501-520.

[9] Li, P., Liu, J., & Qiu, L. (2026). Deep Learning Methods for Demand Forecasting and Inventory Optimization in Modern Supply Chains. Asian Business Research Journal, 11(3), 21-29.

[10] Qiu, L. (2025). Reinforcement Learning Approaches for Intelligent Control of Smart Building Energy Systems with Real-Time Adaptation to Occupant Behavior and Weather Conditions. Journal of Computing and Electronic Information Management, 18(2), 32-37.

[11] Zhang, H. (2025). Reinforcement Learning Approaches for Layout Optimization in Electronic Design Automation with Electromagnetic Compatibility Constraints. Frontiers in Robotics and Automation, 2(2), 77-93.

[12] Shen, Z., Zhao, W., Wang, B., Wang, Z., & Shang, W. (2026). CAGR: A Cross-Accelerator Graph Optimization Framework for Efficient Recommender System Inference. IEEE Access.

[13] Sun, T., Wang, M., & Han, X. (2025). Deep Learning in Insurance Fraud Detection: Techniques, Datasets, and Emerging Trends. Journal of Banking and Financial Dynamics, 9(8), 1-11.

[14] Liu, J., Li, P., & Wang, Y. (2026). Graph Neural Networks for Modeling Complex Dependencies in Global Supply Chain Networks. Journal of Computing and Electronic Information Management, 20(3), 9-20.

[15] Zhang, F., & Wu, B. (2025). Large Language Models as General Purpose Intelligence Systems for Reasoning, Planning and Decision Making. American Journal of Artificial Intelligence and Neural Networks, 6(4), 45-72.

[16] Li, P., Ren, S., Zhang, Q., Wang, X., & Liu, Y. (2024). Think4SCND: Reinforcement learning with thinking model for dynamic supply chain network design. IEEE Access, 12, 195974-195985.

[17] Zhang, F., & Yang, J. S. (2025). Learning Driven Decision Intelligence for Autonomous Driving Through Multimodal Understanding World Modeling and Policy Optimization. Frontiers in Artificial Intelligence Research, 2(3), 616-634.

[18] Wang, B., Wang, Z., Zhao, W., & Liu, Y. (2025). Network Fabric Simulation and Validation for Data Center Routing Convergence Under Large-Scale Failure Scenarios. Computer Science Bulletin, 8(01), 310-326.

[19] Liu, J., Wang, J., Chen, H., Guinness, J., Martin, R., & Kulkarni, C. S. (2019). Optimal Level Crossing Predictions for Electronic Prognostics. In AIAA Scitech 2019 Forum (p. 1962).

[20] Chen, J., Cui, Y., Zhang, X., Yang, J., & Zhou, M. (2024). Temporal convolutional network for carbon tax projection: A data-driven approach. Applied Sciences, 14(20), 9213.

[21] Wei, Z., Sun, T., & Zhou, M. (2024). LIRL: Latent Imagination-Based Reinforcement Learning for Efficient Coverage Path Planning. Symmetry, 16(11), 1537.

[22] Zhang, S., Qiu, L., & Zeng, Z. (2026). Physics-Data Synergy in Structural Health Monitoring: A Multi-Scale Graph Contrastive Framework With Temperature-Adaptive Fusion. IEEE Access.

[23] Ding, J., Chen, T., & Qin, Y. (2026). Achieving Resource Isolation in Multi-Tenant Cloud Platforms Without Sacrificing Performance. Journal of Computing and Electronic Information Management, 21(1), 10-18.

[24] Spencer, M. R., Miniño, A. M., & Warner, M. (2022). Drug overdose deaths in the United States, 2001–2021.

[25] Chen, T., & Ding, J. (2026). Cold Start Latency Optimization Strategies for Function as a Service Platforms. Computer Life, 14(1), 64-73.

[26] Ding, J., & Qin, Y. (2026). Raft and Beyond: Practical Consensus Mechanisms for Geo-Distributed Data Systems. Computer Life, 14(1), 54-63.