THE
MILBANK QUARTERLY
A MULTIDISCIPLINARY JOURNAL OF POPULATION HEALTH AND HEALTH POLICY

Perspective

Stemming the Tide of the US Overdose Crisis: How Can We Leverage the Power of Data Science and Artificial Intelligence?

MAGDALENA CERDÁ,[*] DANIEL B. NEILL,[†,‡,§,‖]
ELLICOTT C. MATTHAY,[*] JOHNATHAN A. JENKINS,[#]
BRANDON D. L. MARSHALL,[**] and KATHERINE M. KEYES[††]

[*]*Center for Opioid Epidemiology and Policy, NYU Langone Health;* [†]*Courant Institute, New York University;* [‡]*Wagner School of Public Service, New York University;* [§]*Center for Urban Science and Progress, New York University;* [‖]*Tandon School of Engineering, New York University;* [#]*Arcos Technologies, Inc.;* [**]*School of Public Health, Brown University;* [††]*Mailman School of Public Health, Columbia University*

Policy Points:

- We can leverage data science and artificial intelligence to inform state and local resource allocation for overdose prevention.
- Data science and artificial intelligence can help us answer four questions: (1) What is the impact of laws on access to interventions and overdose risk? (2) Where should interventions be targeted? (3) Which types of demographic subgroups benefit the most and the least from interventions? and (4) Which types of interventions should they invest in for each setting and population?
- Advances in data science and artificial intelligence can accelerate the pace at which we can answer these critical questions and help inform an effective overdose prevention response.

Keywords: overdose, data science, artificial intelligence, epidemiology, population health, prevention.

The Overdose Crisis: Epidemiologic Profile and Potential Solutions

People in the United States are dying at record numbers from overdose.[1] Overdose deaths increased from fewer than 17,000 deaths in 1999 to an estimated 100,000 deaths approximately 25 years after, with a peak of almost 108,000 deaths

The Milbank Quarterly, Vol. 103, No. S1, 2025 (pp. 124-140)

in 2022.[2] Racial/ethnic minoritized groups are now particularly affected: in 2023, the highest rates of overdose were among non-Hispanic Black and American Indian/Alaska Native Americans.[3] Although overdoses increasingly involve both opioids and stimulants, opioids contribute to over three-quarters of all overdose deaths, primarily driven by illegally manufactured synthetic opioids like fentanyl.[1]

Provisional data indicate that we have seen an overall decline in US overdose deaths in the past 2 years, with a decline of 4% in 2022–2023 and 17% in 2023–2024.[4] Despite this substantial decline, the number of deaths per year remain higher than they were prior to the start of the COVID-19 pandemic. Further, this decline has been unequally experienced across racial and ethnic groups. According to the latest available race-specific data, White populations saw a 7% decline in 2022–2023 compared with a 3% increase among Black Americans, a 39.4% increase among Native Hawaiian or Other Pacific Islander populations, and no change among other racial/ethnic groups.[3] These patterns signal a need to develop targeted and tailored interventions that effectively reduce both total overdose deaths and disparities in overdose.

Federal overdose prevention strategies call for investment at four levels of action to address the overdose crisis: primary prevention, harm reduction, evidence-based treatment, and recovery.[5] Primary prevention focuses on strategies to prevent misuse of opioids and other drugs, and the onset of opioid use disorder (OUD) and other substance use disorders (SUDs), including regulation of the legal drug supply and reduction of the unregulated drug supply. Harm reduction seeks to minimize the negative consequences of drug use by meeting people where they are—whether that is active drug use, treatment, or recovery—and providing evidence-based services in a compassionate manner, including drug-checking equipment, syringe service programs (SSPs), and naloxone, among other services. Evidence-based treatment includes treatment modalities for which research has documented improved outcomes across the spectrum of OUD and other SUD health and recovery. Pharmacological treatment with medications for OUD including methadone, buprenorphine, and naltrexone meets these criteria.[6] Finally, recovery refers to overall improvements in health and social functioning, increased quality of life, and deepened social integration associated with effective treatment of OUD and other SUDs.

Investment across these four levels has strong potential to reduce overdose deaths, yet large fractions of people who use drugs cannot access the services they need. Primary prevention efforts have emphasized supply control by enacting laws and programs to regulate opioid prescribing, including prescription drug monitoring programs, prescription limits, and pain-management clinic laws, among others. The lack of concomitant investment in primary prevention efforts that address other drivers of demand, however, have meant that the unregulated drug market has responded to supply regulation efforts by shifting to cheaper and more potent synthetic products, hindering the potential decline in overdose risk (and causing an increase in risk caused by the increased lethality of the new drug products). Although harm reduction services have expanded substantially in the past 10 years and may have contributed to the decline in overdose deaths in the past 2 years, access to these services remains limited in communities across the nation. For example, although a survey of SSPs found a shift from 55% to 94% of SSPs that implemented overdose education and naloxone distribution between 2013 and 2019, the bulk of naloxone distribution was carried out by only 6% of SSPs.[7] Further, although substantial advances have been made in the distribution and prescribing of buprenorphine and methadone,[8] fewer than one in four people with OUD receive treatment with buprenorphine, methadone, or naltrexone.[9] The harm reduction and treatment gaps are particularly acute for racially/ethnically minoritized groups.[10–14]

The Role of Public Health in Informing an Effective Response

In the next 4 years, opioid settlement funds will likely be one of the largest sources of funding for state and local governments to address the overdose crisis. This represents a unique opportunity for localities to make a substantial and equitable investment in evidence-based programs and services across the four levels of the overdose prevention strategy. More than $50 billion in settlement funds from pharmaceutical companies are being paid out over the next 18 years to state and local governments, and 85% of funding must go to treatment and prevention.[15]

Public health scientists, in partnership with experts in data science and artificial intelligence (AI), have an important role to play in informing state and local governments about the allocation of resources for overdose prevention. They can help governments answer four critical questions: (1) What impact do state and local laws have on access to evidence-based interventions and overdose risk? (2) In which geographic areas should we target interventions? (3) What type of heterogeneity exists across demographic subgroups in overdose risk and in the reach and effectiveness of evidence-based interventions? and (4) What types of interventions will work best and be most cost-effective for each setting and population, given our understanding of population-level effects, geographic targeting, and subgroup heterogeneity in intervention reach and effectiveness?

Data Science and AI: Important Tools for an Epidemiologic Response to the Overdose Crisis

Substantial barriers in data availability and analytic capacity limit our ability to answer these four critical questions. However, we can leverage tools from data science and AI to overcome these barriers and accelerate the pace at which we can help inform an effective overdose prevention response. Below, we describe some of the challenges associated with answering each of the four questions and provide some suggestions about the role that data science and AI can play in addressing these challenges. We illustrate each case with an example from existing research.

Evaluating the Impact of State and Local Laws on Intervention Access and Overdose Risk: Tracking State and Local Laws

State and local (i.e., municipal, county) laws are critical levers that have the power to increase or restrict access to evidence-based primary prevention, harm reduction, treatment, and recovery services. Although there has been a concerted effort to collect data on state laws that can affect overdose prevention, harm reduction, treatment, and recovery, available state law data sources are often outdated by several years,[16] and municipal and county law data sources are virtually nonexistent. Comprehensive and up-to-date data on ever-evolving state laws are limited by the time-intensive effort required to manually extract and code laws across states and over time. This issue is compounded when the number of governments is expanded to consider counties and cities. Simple keyword searches may also miss relevant laws that affect access to services, especially in cases in which services are regulated under broad protections or restrictions (e.g., local ordinances that restrict loitering or congregation in public spaces, which are used to target methadone clinics with early morning opening hours). Finally, manual data collection does not allow for changes to the measures included in the legal data collection protocol midway through the process, as doing so would require coders to go back and repeat months of work to accommodate the changes in a consistent way.

Developments in AI—in particular, the current generation of text-embedding models and large language models—offer a potential breakthrough in the way we review legal materials for epidemiology. Text embeddings encode blocks of text as vectors of numbers representing the content of the text; these embeddings can then be used to implement better retrieval over large legal documents compared with traditional keyword searches.[17] Further, large language models (such as OpenAI's GPT-4o, Google's Gemini 1.5, or Anthropic's Claude 3.5 Sonnet) make it feasible to extract information from retrieved legal text and generate quantitative variables capturing the presence or absence and content of prespecified provisions, together with excerpts of the legal text for human review. The broad category of techniques following this approach is known as retrieval-augmented generation (RAG).[18–20]

An example application relates to the retrieval and coding of local harm reduction laws. Using traditional methods, a team of analysts working at the NYU Center for Opioid Epidemiology and Policy, under the supervision of legal experts from the Network for Public Health Law, spent 2 years hand coding data about local laws relevant to harm reduction services and supplies for several hundred counties and municipalities. We are now developing a prototype RAG pipeline with the aim of evaluating the new approach against the traditional hand-coding method, updating the legal data set as new laws are introduced or revised and extending the data collection to other domains of law. Our team is also tailoring RAG techniques for legal research (e.g., taking care to segment the source texts to preserve the structure of the document, preserving section headings for additional context and accurate citations, pulling in related definitions and other terms from across the documents when answering questions, and providing retrieved citations as validation information for review and verification by human analysts).

The application of AI to retrieve and code laws is still in development, and questions remain about whether the models will be able to accurately identify and code laws as well as humans can. Limited access to legal databases without a paywall (e.g., LexisNexis) constrains the development of automated retrieval systems. The lack of a uniform structure for legal texts increases complexity and potential error rates of AI approaches. Further, the quality of validation depends on how much human-coded data there are for comparison. At the same time, this approach offers several potential benefits, including the capacity to identify relevant legal code that could be missed by a simple keyword search by humans, substantially lower cost and time, and introduces greater flexibility to add new laws and variables to the resulting legal database with the simple change in lines of code rather than through repetition of a years' long legal review. If successful, the application of large language models and other forms of AI to legal research could transform the pace of discovery of the effects of laws on the overdose crisis.

Geographic Targeting of Interventions

In a context of scarce resources, policymakers and community stakeholders need data to guide their choices about where to allocate limited resources to prevent overdoses from escalating in their community. The overdose crisis has evolved at a rapid pace so that governments are often forced to make choices about where to allocate resources in a reactive fashion using surveillance data that are often outdated or institutional or individual knowledge. Yet, because of the dynamic nature of the crisis, past overdose burden may not reflect current community overdose risk.[21] Hence, new methods are needed to proactively and accurately forecast how the crisis will evolve in communities so that local stakeholders can optimally target interventions to prevent future spikes in overdose. Spatiotemporal machine learning models offer a tool to predict which geographic areas have the highest future risk of overdose.

The Preventing Overdose Using Information and Data from the Environment (PROVIDENT) study[22–25] illustrates the potential use of spatiotemporal machine learning to forecast community-level overdose risk and to use these predictions to allocate resources. The study developed ensemble-based machine learning models to predict the 20% of census block groups (CBGs) in Rhode Island towns that would experience at least 40% of state overdose deaths in the next 6 months and evaluated these models with respect to various measures of accuracy and equity. The final model successfully and consistently identified the top 20% of CBGs in which more than 40% of overdose deaths subsequently occurred.[22–25] The team also developed a dashboard tool to inform community-based harm reduction organizations about CBGs prioritized by the model so that the organizations could use this information to steer outreach and other services to prioritized areas. An evaluation is now underway to test whether fine-grained targeting of interventions at the community level (based on prediction from machine learning models) can reduce overdose rates and increase equity in allocating limited prevention resources.[22–25] This type of model can be adapted to new settings to help local governments decide where to proactively allocate overdose prevention funds given their own relevant community characteristics.

The application of machine learning as a decision support tool to inform the allocation of community-level overdose prevention interventions requires the development of new, practice-based model evaluation criteria that are tailored to the concerns of local practitioners and that supplement traditional model performance metrics (e.g., mean squared error). As part of PROVIDENT, Allen and colleagues proposed four guiding questions that can help researchers work with practitioners to define the parameters of a neighborhood-based forecasting model and to evaluate model performance.[23,25] First, what capacity do stakeholders have to distribute prevention resources across a government? In PROVIDENT, resource constraints expressed by local partners led us to focus on 20% of CBGs that could be prioritized for overdose prevention per 6-month period. Second, what is the target preventive potential of focused intervention deployment? In the case of PROVIDENT, our target was to prevent 40% of overdose deaths in 5 years by targeting 20% of CBGs, so we evaluated whether the model could predict at least 40% of deaths, assuming 20% implementation capacity. Third, based on model predictions, how will resources be allocated across different intersecting dimensions of equity (e.g., geographic, racial/ethnic, and socioeconomic)? In PROVIDENT, we evaluated the racial/ethnic and socioeconomic profile of CBGs prioritized by the model to ensure that equity was a key consideration in both prediction and resource allocation. Fourth, how would local partners like to deploy resources within jurisdictions? In the case of PROVIDENT, practitioner concerns about ensuring that all towns had at least one prioritized CBG led us to set a model constraint to select at least one CBG per town. Application of such evaluation criteria will ensure that machine learning is deployed in partnership with local stakeholders so that the model fits the specific use case, priorities, and practical constraints relevant to the partners and the setting.

Machine learning approaches to forecasting overdose risk and informing intervention allocation do pose their own set of challenges. First, prediction accuracy is limited by the number of observation time points and communities available to train the model.[26] Second, the type of data available to inform predictions may induce selection bias. For example, data on patients receiving treatment for OUD may only reflect risk among service-involved populations, whereas data that do not vary over time may only help identify communities with high, stable risk of overdose as opposed to more dynamic patterns of risk. Third, predictions are only useful to the extent that they are used by local partners. Hence, the types of questions that will be answered using machine learning tools, the criteria against which model performance is evaluated, and the meaning of the predictions should all be decided in close partnership with local stakeholders who will put the information into practice.

Measuring Heterogeneity in Risk, Reach, and Effectiveness Across Intersectional Subgroups

Escalating disparities in overdose rates across racial and ethnic groups[3] point to a potentially unequal impact of current efforts to address the overdose crisis. For example, although they experience the highest rates of overdose deaths, non-Hispanic Black people are less likely to access harm reduction services such as overdose education and naloxone distribution,[10–12] and evidence-based SUD treatment than non-Hispanic White people.[13,14] Epidemiologists working with local partners must thus ask key questions about the heterogeneity of intervention reach and effectiveness. That is, which population groups do evidence-based overdose prevention policies, services, and interventions reach, and which population groups do they fail to reach? Among which population groups do we see a reduction in overdoses and related harms following the implementation of a policy, service, or intervention, and in which groups do we fail to see an effect? To better understand sources of inequity, we must go beyond only comparing racial/ethnic groups to investigate variation among subgroups of the population defined across intersecting demographic, social, contextual, geographic, and behavioral dimensions. An intersectional lens reveals that the impacts of intersecting social statuses are often not simply additive; for example, Black women may experience barriers to accessing overdose prevention services that are not fully explained by racism and sexism.[27–29] An intersectional approach is therefore critical to ensure true equity in investment of resources to reduce the overdose crisis.

Quantifying heterogeneity across multiple dimensions presents significant computational and statistical challenges, particularly in identifying subpopulations with the greatest inequities. These challenges stem from the need to search across numerous subgroups while accounting for multiple hypothesis testing, which can increase the risk of false discoveries and inflate effect size estimates. Recent advances in machine learning and causal inference,[30–37] some of which we describe below, offer novel solutions to these issues. By leveraging the linear-time subset scanning property,[31,34] these methods efficiently identify the most significant population subgroups, reducing the computational burden from an exponential to a linear number of subset evaluations and dramatically decreasing processing time from years to seconds. Furthermore, randomization testing provides robust significance estimates while correcting for multiple comparisons, and data-splitting techniques ensure unbiased effect size estimation.

Two types of machine learning tools are helpful to quantify heterogeneity across intersectional subgroups. First, tools such as Bias Scan[38,39] and Conditional Bias Scan[40] can help us quantify heterogeneity in the reach of policies, services, and interventions. Bias Scan[38,39] identifies intersectional subgroups with significantly higher or lower probability of receiving an intervention than expected based on the overall population. Instead, Conditional Bias Scan[40] identifies intersectional subgroups for whom the probability of an outcome is significantly different for individuals in a protected class, compared with individuals in a nonprotected class. For example, Bias Scan could be used to discover intersectional subgroups for whom the reach of services is significantly lower than the population as a whole (e.g., service reach among Black men living in rural areas may be significantly lower than among the general US population). In contrast, Conditional Bias Scan could be used to discover contexts in which and subgroups, defined in terms of social identities other than race, for whom a service is less likely to reach a certain racial group compared with their counterparts in other racial groups (e.g., reach among Black men living in rural areas may be significantly lower than among non-Black men living in rural areas).

To date, Bias Scan and Conditional Bias Scan have primarily been used for assessing systematic biases in predictive models rather than inequities in the reach of services.[38–40] However, Bias Scan has been applied to identify intersectional subgroups most impacted by overpolicing, analyzing over 760,000 pedestrian stops made under the New York Police Department's "stop and frisk" policy.[39]

Second, tools such as Heterogeneous Treatment Effect (HTE)-Scan[41] can be used to quantify heterogeneity in the effectiveness of policies, services, or interventions. This approach identifies intersectional subgroups for whom the treatment effect is larger or smaller than the overall average treatment effect. We can use such an approach to detect subgroups for whom a policy, service, or intervention is significantly more, or less, effective in preventing overdose. As a concrete example of the use of HTE-Scan to discover heterogeneity (applied to an exposure rather than an intervention), we analyzed the impact of the COVID-19 pandemic on overdose among New York State (NYS) Medicaid enrollees.[41] During March-December 2020, overdoses increased 21% compared with prepandemic levels (March-December 2019) for the

NYS Medicaid population as a whole. However, HTE-Scan identified two large intersectional subgroups with statistically significant heterogeneity: overdoses increased by 54% for Black and Hispanic men aged 45–64 years old with no documented OUD and by 57% for blind or disabled individuals aged 45–64 years old with no documented OUD, compared with prepandemic levels. This example demonstrates the ability of HTE-Scan to identify intersectional subpopulations (including race, gender, age, and/or disability status) who were most significantly impacted by the confluence of the COVID-19 and overdose crises.

The use of machine learning tools such as Bias Scan, Conditional Bias Scan, and HTE-Scan to search over intersectional subgroups faces several challenges. First, although the use of randomization testing to determine the statistical significance of detected subgroups correctly adjusts for the search procedure and thus avoids the problem of high false-positive rates, this procedure is computationally expensive, and larger search spaces correspond to reduced detection power. Second, these tools work best when the number of data records is large; for smaller data sets, there may not be enough data records for an intersectional subgroup to identify whether that group is differentially impacted. Third, Bias Scan is a descriptive method, identifying differences among subgroups, and results should not be interpreted causally. In contrast, HTE-Scan explicitly measures the causal effect of treatment, and Conditional Bias Scan can be implicitly interpreted as measuring the causal effect of membership in the protected class, controlling for observed variables that may affect selection into treatment or protected class membership. These methods require standard assumptions for unbiased causal inference,[42–44] and any unmeasured confounders can lead to biased estimates of treatment effects. Approaches for addressing unmeasured confounders include sensitivity analysis[45,46] as well as new approaches developed by our team for automated discovery of natural experiments,[47–49] which produce unbiased local treatment effect estimates even in the presence of unmeasured confounding and extrapolate these estimates to the remainder of the population.

What Types of Interventions Will Work Best for Each Setting, and How Much to Invest

Once intervention target geographic areas or population subgroups are identified, the next step is to work with local stakeholders to identify the types of interventions that will have the greatest potential to reduce population rates of overdose and other harms in prioritized sites and populations given resource constraints and cost benefit tradeoffs associated with alternative investment strategies. This is particularly salient at the current time, with the funds promised to states, cities, and counties as a result of lawsuits against pharmaceutical manufacturers, distributors, and retailers.[15] Many governments are engaging with public health experts to make data-driven decisions about how, where, and when to spend funds to ensure the highest impact immediately and into the future. The lack of localized evidence from intervention trials that would be suited to answer these questions highlights the need for new methods.

Computational simulation models such as systems dynamics models[49] and agent-based models[50,51] are ideally suited to answer these "what if" questions. By relying on local data (e.g., on overdose rates, risk factors, and currently deployed services and interventions) and a methodological toolkit customized to a community's priority concerns, simulation models can be used to advise local governments on optimal investments to achieve their objectives. Robust, mathematical models of opioid use, overdose, and other opioid-related outcomes have been in development for more than 10 years.[52,53] These models use empirical data to simulate dynamics and consequences of drug use within designated geographic areas. An advantage of systems' science simulation modeling is that it incorporates empirical data and assumptions about connections between key inputs and outputs and can incorporate nonlinear dynamics, feedback loops, and emergent properties that shape overdose dynamics in specific communities and subpopulations. Table 1 provides examples of the types of questions models can answer, which models are best suited to answer these questions, and examples of studies that have applied these models.

One example of the utility of simulation models is the Healing Communities Study (HCS).[68] The HCS aimed to reduce opioid-related overdoses by 40% in selected counties using a multisite, parallel-group, cluster randomized wait-list controlled trial. The HCS was implemented in 67 counties in New York, Kentucky, Massachusetts, and Ohio. The study tested an implementation science framework to assist counties in selecting and adopting three types of evidence-based interventions: naloxone distribution, medication for OUD (MOUD) initiation and retention, and safe opioid prescribing. In parallel to conducting an evaluation of the impact of the HCS, the HCS team used a range of microsimulation, systems dynamics, and agent-based models to inform local governments about the types of investments that would help them achieve the 40% overdose reduction target. For example, the NYS team used the agent-based model Simulation of Community-Level Overdose Prevention Strategy (SiCLOPS)[61] to advise eight New York counties about the amount by which delivery of MOUD and naloxone would have to increase to reduce overdose deaths by 40%. County-specific models calibrated agent characteristics using data from the county on opioid use and use disorder, treatment availability, access and uptake, and overdose. The study used model visualization and infographic materials in a public-facing web portal to help counties to directly engage with the data and simulation findings.[61] The results from SiCLOPS demonstrated the need for locally tailored interventions to reduce overdose deaths, as the level of investment in MOUD and naloxone distribution required depended on county urbanicity and on each county's preexisting treatment and harm reduction service infrastructure.

Table 1. Example Questions Addressed in Computational Simulation Models

Question	Model
Given the current set of circumstances and resources, which of a defined set of potential interventions (e.g., harm reduction, treatment, recovery support) is likely to have the biggest impact on outcomes of interest in this community in 1, 2, and 3 years?	Systems dynamics models with large simulated populations to approximate social networks and mixing, forecast anticipated changes in overdose; simpler and focused compartmental models to address statewide change[54–58]
Within a county, are there geographic hotspots that would most benefit from a particular intervention (e.g., harm reduction, treatment, recovery support)? How can we leverage social networks to distribute resources and expand the impact of interventions?	Agent-based models at a local, granular level that accommodate data on geography and distance traveled; microsimulation models that simulate individual trajectories over multiple years under different scenarios[52,59–66]
In which geographic areas should we place interventions to have maximum impact on our intended outcomes? How could the distribution of unprescribed buprenorphine in networks affect opioid use and heroin/fentanyl mortality?	Agent-based models at a local, granular level that accommodate data on geography, service location, and distance traveled[62–66]
How do organizational networks, administrative structures, and contextual factors affect the adoption of evidence-supported behavioral health policies within an agency?	Network-based, agent-based models[67]

The application of simulation models to identify intervention priorities must be considered in light of current methodological challenges. First, data sources used for model calibration will define model-based outcomes and, therefore, the accuracy and utility of estimated intervention impacts. Yet, despite investigating different contexts and populations, most current studies using simulation approaches rely on the same data sources. Hence, any limitations of such data will be reflected in the estimates obtained from simulation models. Second, in most cases, researchers rely on published estimates of intervention effects to calibrate model parameters.[52] However, intervention-effect estimates may not be transportable across populations, periods, and geographic areas.[69,70] Third, the quality of simulation models directly depends on calibration and validation; established guidelines for calibration and validation of simulation models should be followed in future studies using these methods.[71,72]

Conclusions

Extraordinary developments in AI and data science offer new opportunities for researchers to partner with policymakers and practitioners to inform evidence-based approaches to address the overdose crisis. As public health scientists, we can leverage data science and AI to accelerate the pace of research on urgent public health problems such as overdose. This paper covered just four examples of the types of applications of these methods to inform overdose prevention, including automating the extraction and coding of data on policy levers with greater frequency and fewer resources, forecasting emerging trends in key public health problems in small geographic areas to optimally target interventions, quantifying inequities in intervention reach and effectiveness across intersectional subgroups, and advising governments about the expected outcomes associated with alternative investments.

As applications of data science and AI in public health accelerate, we must bring together experts with different disciplinary perspectives and different methodological and modeling expertise and integrate a range of data sources that offer different strengths. With such a consortium, we can work with local stakeholders to identify and answer important questions communities would like answered, offering local governments a "modeling toolbox" they can leverage to inform intervention targets and to devise optimal responses to this continued overdose crisis.

References

1. Drug overdose deaths: facts and figures. National Institute on Drug Abuse, National Institutes of Health. August 21, 2024. Accessed July 25, 2023. https://nida.nih.gov/research-topics/trends-statistics/overdose-death-rates
2. Spencer MR, Garnett MF, Miniño AM. *Drug Overdose Deaths in the United States, 2002–2022*. NCHS Data Brief No 491. National Center for Health Statistics; 2024.
3. Garnett MF, Miniño AM. *Drug Overdose in the United States, 2003–2023*. NCHS Data Brief No 522. National Center for Health Statistics; 2024.
4. Ahmad FB, Cisewski JA, Rossen LM, Sutton P. Provisional drug overdose death counts. National Center for Health Statistics. Updated May 14, 2025. Accessed January 1, 2025. https://dx.doi.org/10.15620/cdc/20250305008
5. Cerdá M, Krawczyk N, Keyes K. The future of the United States overdose crisis: challenges and opportunities. *Milbank Q*. 2023;101(S1):478-506.
6. Bell J, Strang J. Medication treatment of opioid use disorder. *Biol Psychiatry*. 2020;87(1):82-88. https://doi.org/10.1016/j.biopsych.2019.06.020

7. Lambdin BH, Bluthenthal RN, Wenger LD, et al. Overdose education and naloxone distribution within syringe service programs – United States, 2019. *MMWR Morb Mortal Wkly Rep*. 2020;69(33):1117-1121. doi: 10.15585/mmwr.mm6933a2
8. Kennalley AL, Fanelli JL, Furst JA, et al. Dynamic changes in methadone utilisation for opioid use disorder treatment: a retrospective observational study during the COVID-19 pandemic. *BMJ Open*. 2023;13(11):e074845. https://doi.org/10.1136/bmjopen-2023-074845
9. Jones CM, Han B, Baldwin GT, Einstein EB, Compton WM. Use of medication for opioid use disorder among adults with past-year opioid use disorder in the US, 2021. *JAMA Netw Open*. 2023;6(8):e2327488. https://doi.org/10.1001/jamanetworkopen.2023.27488
10. Khan MR, Hoff L, Elliott L, et al. Racial/ethnic disparities in opioid overdose prevention: comparison of the naloxone care cascade in White, Latinx, and Black people who use opioids in New York City. *Harm Reduct J*. 2023;20(1):24. https://doi.org/10.1186/s12954-023-00736-7
11. Nolen S, Zang X, Chatterjee A, et al. Evaluating equity in community-based naloxone access among racial/ethnic groups in Massachusetts. *Drug Alcohol Depend*. 2022;241:109668. https://doi.org/10.1016/j.drugalcdep.2022.109668
12. Zang X, Walley AY, Chatterjee A, et al. Changes to opioid overdose deaths and community naloxone access among Black, Hispanic and White people from 2016 to 2021 with the onset of the COVID-19 pandemic: an interrupted time-series analysis in Massachusetts, USA. *Addiction*. 2023;118(12):2413-2423. https://doi.org/10.1111/add.16324
13. Lagisetty PA, Ross R, Bohnert A, Clay M, Maust DT. Buprenorphine treatment divide by race/ethnicity and payment. *JAMA Psychiatry*. 2019;76(9):979-981. https://doi.org/10.1001/jamapsychiatry.2019.0876
14. Magee T, Peters C, Jacobsen SM, et al. Inequities in the treatment of opioid use disorder: a scoping review. *J Subst Use Addict Treat*. 2023;152:209082. https://doi.org/10.1016/j.josat.2023.209082
15. Minhee C. Tracking opioid settlements & spending plans since 2019. Updated March 23, 2025. Accessed January 1, 2025. https://www.opioidsettlementtracker.com/
16. Prescription Drug Abuse Policy System. Prescription Drug Abuse Policy System. 2025. Accessed January 1, 2025. pdaps.org
17. Chitturi K. Understanding vector embeddings in AI search: a technical deep dive. *Int J Comput Eng Technol*. 2024;15(6):1725-1733.
18. Gao Y, Xiong Y, Gao X, et al. Retrieval-augmented generation for large language models: a survey. *arXiv*. Preprint revised online March 27, 2024. doi: 10.48550/arXiv.2312.10997
19. Gupta S, Ranjan R, Narayan Singh S. A comprehensive survey of retrieval-augmented generation (RAG): evolution, current landscape and future directions. *arXiv*. Preprint posted online October 3, 2024. doi: 10.48550/arXiv.2410.12837

20. Feuerriegel S, Maarouf A, Bär D, et al. Using natural language processing to analyse text data in behavioural science. *Nat Rev Psychol*. 2025;4:96-111. https://doi.org/10.1038/s44159-024-00392-z
21. Davis CS, Green TC, Zaller ND. Addressing the overdose epidemic requires timely access to data to guide interventions. *Drug Alcohol Rev*. 2016;35(4):383-386. https://doi.org/10.1111/dar.12321
22. Schell RC, Allen B, Goedel WC, et al. Identifying predictors of opioid overdose death at a neighborhood level with machine learning. *Am J Epidemiol*. 2022;191(3):526-533. https://doi.org/10.1093/aje/kwab279
23. Allen B, Neill DB, Schell RC, et al. Translating predictive analytics for public health practice: a case study of overdose prevention in Rhode Island. *Am J Epidemiol*. 2023;192(10):1659-1668. https://doi.org/10.1093/aje/kwad119
24. Marshall BDL, Alexander-Scott N, Yedinak JL, et al. Preventing Overdose Using Information and Data from the Environment (PROVIDENT): protocol for a randomized, population-based, community intervention trial. *Addiction*. 2022;117(4):1152-1162. https://doi.org/10.1111/add.15731
25. Allen B, Schell RC, Jent VA, et al. PROVIDENT: development and validation of a machine learning model to predict neighborhood-level overdose risk in Rhode Island. *Epidemiology*. 2024;35(2):232-240.
26. DeCamp M, Lindvall C. Latent bias and the implementation of artificial intelligence in medicine. *J Am Med Inform Assoc*. 2020;27(12):2020-2023. https://doi.org/10.1093/jamia/ocaa094
27. Crenshaw K. Mapping the margins: intersectionality, identity politics, and violence against women of color. *Stan Law Rev*. 1991;43(6):1241-1299. https://doi.org/10.2307/1229039
28. Crenshaw K. Race, gender, and sexual harassment. *S Cal L Rev*. 1992;65:1467-1476.
29. Collins P. *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment*. Routledge; 1990.
30. Neill DB. Subset scanning for event and pattern detection. In: Shekhar S, Xiong H, Zhou X, eds. *Encyclopedia of GIS*. Springer; 2017:2218-2228.
31. Neill DB. Fast subset scan for spatial pattern detection. *J R Stat Soc Series B Stat Methodol*. 2012;74(2):337-360. https://doi.org/10.1111/j.1467-9868.2011.01014.x
32. Neill DB, McFowland E III, Zheng H. Fast subset scan for multivariate event detection. *Stat Med*. 2013;32(13):2185-2208. https://doi.org/10.1002/sim.5675
33. McFowland E, Speakman S, Neill D. Fast generalized subset scan for anomalous pattern detection. *J Mach Learn Res*. 2013;14:1533-1561.
34. Speakman S, Somanchi S, McFowland E, Neill D. Penalized fast subset scanning. *J Comput Graph Stat*. 2016;25(2):382-404. https://doi.org/10.1080/10618600.2015.1029578
35. Fitzpatrick D, Ni Y, Neill D. Support vector subset scan for spatial pattern detection. *Comput Stat Data Anal*. 2020;157(55):107149. https://doi.org/10.1016/j.csda.2020.107149

36. Herlands W, McFowland E, Wilson A, Neill D. Gaussian process subset scanning for anomalous pattern detection in non-iid data. *Proc Mach Learn Res*. 2018;84:425-434.
37. McFowland E III, Somanchi S, Neill D. Efficient discovery of heterogeneous quantile treatment effects in randomized experiments via anomalous pattern detection. *arXiv*. Preprint revised online May 10, 2023. doi: 10.48550/arXiv.1803.09159
38. Zhang Z, Neill DB. Identifying significant predictive bias in classifiers. *arXiv*. Preprint revised online July 4, 2017. doi: 10.48550/arXiv.1611.08292
39. Ravishankar P, Mo Q, McFowland E III, Neill DB. Provable detection of propagating sampling bias in prediction models. *Proc AAAI Conf Artif Intell*. 2023;37(8):9562-9569.
40. Boxer K, McFowland E III, Neill DB. Auditing predictive models for intersectional biases. *arXiv*. Preprint posted online June 22, 2023. doi: 10.48550/arXiv.2306.13064
41. Pamplin JR II, Wheeler-Martin K, Shroff R, Cerdá M, Neill D. Identifying demographic predictors of increased opioid overdose risk among New York State Medicaid enrollees following the COVID-19 pandemic: an analysis of heterogeneous treatment effects. PowerPoint presented at: World Congress of Epidemiology 2024; September 27, 2024; Cape Town, South Africa. Accessed January 1, 2025. https://www.wce2024.org/wp-content/uploads/2024/10/John-Identifying-heterogeneous-treatment-effects-of-the-COVID-19-pandemic-on-non-fatal-opioid-overdose-amount-New-York-State-Medicaid-enrollees.pdf
42. Igelström E, Craig P, Lewsey J, Lynch J, Pearce A, Katikireddi SV. Causal inference and effect estimation using observational data. *J Epidemiol Community Health*. 2022;76(11):960-966. https://doi.org/10.1136/jech-2022-219267
43. Hernán MA, Robins JM. *Causal Inference: What If?* Chapman & Hall; CRC Press; 2020.
44. Imbens G, Rubin D. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press; 2015.
45. Lash TL, Fox MP, MacLehose RF, Maldonado G, McCandless LC, Greenland S. Good practices for quantitative bias analysis. *Int J Epidemiol*. 2014;43(6):1969-1985. https://doi.org/10.1093/ije/dyu149
46. Rosenbaum P, Rubin D. Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *J R Stat Soc Series B Stat Methodol*. 1983;45(2):212-218.
47. Herlands W, McFowland Iii E, Wilson A, Neill D. Automated local regression discontinuity design discovery. *Proc ACM SIGKDD Int Conf Knowl Discov Data Min*. 2018:1512-1520.
48. Jakubowski B, Somanchi S, McFowland Iii E, Neill D. Exploiting discovered regression discontinuities to debias conditioned-on-observable estimators. *J Mach Learn Res*. 2023;24(133):1-57.
49. Ip EH, Rahmandad H, Shoham DA, et al. Reconciling statistical and systems science approaches to public health. *Health Educ Behav*. 2013;40(1 Suppl):123S-131S. https://doi.org/10.1177/1090198113493911

50. Marshall BD, Galea S. Formalizing the role of agent-based modeling in causal inference and epidemiology. *Am J Epidemiol*. 2015;181(2):92-99. https://doi.org/10.1093/aje/kwu274
51. Pearce N, Merletti F. Complexity, simplicity, and epidemiology. *Int J Epidemiol*. 2006;35(3):515-519. https://doi.org/10.1093/ije/dyi322
52. Cerdá M, Jalali MS, Hamilton AD, et al. A systematic review of simulation models to track and address the opioid crisis. *Epidemiol Rev*. 2022;43(1):147-165. https://doi.org/10.1093/epirev/mxab013
53. Beaulieu E, DiGennaro C, Stringfellow E, et al. Economic evaluation in opioid modeling: systematic review. *Value Health*. 2021;24(2):158-173.
54. Linas BP, Savinkina A, Madushani R, et al. Projected estimates of opioid mortality after community-level interventions. *JAMA Netw Open*. 2021;4(2):e2037259. https://doi.org/10.1001/jamanetworkopen.2020.37259
55. Flam-Ross JM, Marsh E, Weitz M, et al. Economic evaluation of extended-release buprenorphine for persons with opioid use disorder. *JAMA Netw Open*. 2023;6(9):e2329583. https://doi.org/10.1001/jamanetworkopen.2023.29583
56. Chatterjee A, Weitz M, Savinkina A, et al. Estimated costs and outcomes associated with use and nonuse of medications for opioid use disorder during incarceration and at release in Massachusetts. *JAMA Netw Open*. 2023;6(4):e237036. https://doi.org/10.1001/jamanetworkopen.2023.7036
57. Savinkina A, Madushani R, Eftekhari Yazdi G, et al. Population-level impact of initiating pharmacotherapy and linking to care people with opioid use disorder at inpatient medically managed withdrawal programs: an effectiveness and cost-effectiveness analysis. *Addiction*. 2022;117(9):2450-2461. https://doi.org/10.1111/add.15879
58. Adams JW, Savinkina A, Fox A, et al. Modeling the cost-effectiveness and impact on fatal overdose and initiation of buprenorphine-naloxone treatment at syringe service programs. *Addiction*. 2022;117(10):2635-2648. https://doi.org/10.1111/add.15883
59. Keyes KM, Hamilton A, Swanson J, Tracy M, Cerdá M. Simulating the suicide prevention effects of firearms restrictions based on psychiatric hospitalization and treatment records: social benefits and unintended adverse consequences. *Am J Public Health*. 2019;109(S3):S236-S243. https://doi.org/10.2105/ajph.2019.305041
60. Tracy M, Cerdá M, Keyes KM. Agent-based modeling in public health: current applications and future directions. *Annu Rev Public Health*. 2018;39:77-94. https://doi.org/10.1146/annurev-publhealth-040617-014317
61. Cerdá M, Hamilton AD, Hyder A, et al. Simulating the simultaneous impact of medication for opioid use disorder and naloxone on opioid overdose death in eight New York counties. *Epidemiology*. 2024;35(3):418-429. https://doi.org/10.1097/ede.0000000000001703
62. Adams JW, Duprey M, Khan S, Cance J, Rice DP, Bobashev G. Examining buprenorphine diversion through a harm reduction lens: an agent-based modeling study. *Harm Reduct J*. 2023;20(1):150. https://doi.org/10.1186/s12954-023-00888-6

63. Bobashev G, Goree S, Frank J, Zule W. Pain Town, an agent-based model of opioid use trajectories in a small community. In: Thomson R, Dancy C, Hyder A, Bisgin H, eds. *Social, Cultural, and Behavioral Modeling*. Springer International Publishing; 2018:274-285.
64. Bobashev G, Eggleston B, Morris R, et al. Impact of recovery interventions on opioid users. A simulation study. *Drug Alcohol Depend*. 2017;171:e21-e22. https://doi.org/10.1016/j.drugalcdep.2016.08.074
65. Hoffer L, Bobashev G, Morris RJ. Simulating patterns of heroin addiction within the social context of a local heroin market. In: Gutkin B, Ahmed SH, eds. *Computational Neuroscience of Drug Addiction*. Springer Science & Business Media; 2012:313-331.
66. Hoffer LD, Bobashev G, Morris RJ. Researching a local heroin market as a complex adaptive system. *Am J Community Psychol*. 2009;44(3-4):273-286. https://doi.org/10.1007/s10464-009-9268-2
67. Combs T, Nelson KL, Luke D, et al. Simulating the role of knowledge brokers in policy making in state agencies: an agent-based model. *Health Serv Res*. 2022;57(Suppl 1):122-136. https://doi.org/10.1111/1475-6773.13916
68. Healing Communities Study Consortium; Samet JH, El-Bassel N, Winhusen TJ, et al. Community-based cluster-randomized trial to reduce opioid overdose deaths. *N Engl J Med*. 2024;391(11):989-1001. https://doi.org/10.1056/NEJMoa2401177
69. Murray EJ, Robins JM, Seage GR, Freedberg KA, Hernan MA. A comparison of agent-based models and the parametric g-formula for causal inference. *Am J Epidemiol*. 2017;186(2):131-142. https://doi.org/10.1093/aje/kwx091
70. Keyes KM, Smith GD, Koenen KC, Galea S. The mathematical limits of genetic prediction for complex chronic disease. *J Epidemiol Community Health*. 2015;69(6):574-579. https://doi.org/10.1136/jech-2014-204983
71. Briggs AH, Weinstein MC, Fenwick EA, et al. Model parameter estimation and uncertainty analysis: a report of the ISPOR-SMDM Modeling Good Research Practices Task Force Working Group-6. *Med Decis Making*. 2012;32(5):722-732. https://doi.org/10.1177/0272989X12458348
72. Pitman R, Fisman D, Zaric GS, et al. Dynamic transmission modeling: a report of the ISPOR-SMDM Modeling Good Research Practices Task Force Working Group-5. *Med Decis Making*. 2012;32(5):712-721. https://doi.org/10.1177/0272989X12454578

Conflict of Interest Disclosures: Drs. Cerdá and Keyes have participated in opioid litigation as expert witnesses. Dr. Marshall is chair of the state of Rhode Island's opioid settlement advisory committee. The perspectives shared in this article are those of the authors alone and do not represent the views of the state of Rhode Island nor other committee members.

Address correspondence to: Magdalena Cerdá, NYU Langone Health, 180 Madison Ave, 5th Floor, New York, NY 10016 (email: magdalena.cerda@nyulangone.org).