SPATIAL STATISTICS · PUBLIC HEALTH · GEODA + ARCGIS PRO

Poverty and Opioid Mortality in Metropolitan Denver

An exploratory spatial data analysis of 620 census tracts, 2015–2019

Analyst	Jared Dumale
Question	How is poverty distributed across the Denver–Aurora–Lakewood metro area, and where does it coincide with opioid and heroin mortality?
Study area	Denver–Aurora–Lakewood, CO Metropolitan Statistical Area — 10 counties, 620 census tracts
Data	Colorado Department of Public Health & Environment Open Data: tract-level age-adjusted opioid/heroin mortality (2015–2019) and ACS 2015–2019 poverty rate
Software	GeoDa (spatial weights, Moran's I, LISA); ArcGIS Pro (cartography, reclassification, map algebra)
Methods	Queen contiguity weights, global Moran's I with 999 permutations, Local Indicators of Spatial Association, composite raster index
Headline finding	**Poverty is roughly twice as spatially structured as opioid mortality (Moran's I = 0.445 vs 0.184). Poverty forms a coherent urban core and suburban ring; overdose mortality is far more diffuse.**

Originally produced as a final project for a graduate spatial analytics course. Revised for portfolio presentation: framing, statistics reporting, and cartography corrected; underlying analysis unchanged.

Summary

Opioid and heroin overdose is one of the more spatially uneven public health burdens in the United States, and the intuition that it tracks poverty is widespread. This analysis tested how far that intuition survives contact with tract-level data for metropolitan Denver.

Two variables were mapped and analysed independently across all 620 census tracts in the Denver–Aurora–Lakewood MSA: age-adjusted mortality from prescription opioid analgesic or heroin poisoning for 2015–2019, and the share of residents living below the federal poverty line over the same period. Global Moran's I established whether each variable clusters in space; Local Indicators of Spatial Association identified where.

Both variables cluster significantly, but not to the same degree. Poverty returned a Moran's I of 0.445, with 255 of 620 tracts falling in a statistically significant local cluster. Opioid mortality returned 0.184, with only 121 tracts significant. Poverty has a clean, legible spatial structure: a high-poverty core through central Denver and a broad low-poverty suburban ring. Overdose mortality is markedly more scattered, with far fewer coherent hot spots and proportionally more spatial outliers.

That asymmetry is the finding worth carrying forward. If overdose mortality were simply a function of neighbourhood poverty, the two maps would be considerably more alike than they are.

What This Analysis Does and Does Not Establish

This is an exploratory analysis. It characterises the spatial structure of each variable on its own terms and identifies where high values of the two coincide. It is descriptive rather than inferential.

It does not measure the statistical association between poverty and opioid mortality. Global Moran's I and univariate LISA test whether a single variable is correlated *with itself* across space; neither is a test of the relationship between two different variables. The composite index in the final section shows where high values co-occur, which is a mapping operation, not a correlation coefficient. The appropriate next step is set out at the end of this report.

The distinction matters because co-occurrence at tract level would not, on its own, license a claim about individuals. A tract with both a high poverty rate and a high overdose rate does not establish that the residents who died were the residents in poverty. That inference is the ecological fallacy, and this analysis is not structured to avoid it.

Study Area and Data

The Denver–Aurora–Lakewood MSA comprises ten counties: Adams, Arapahoe, Broomfield, Clear Creek, Denver, Douglas, Elbert, Gilpin, Jefferson, and Park. The MSA was chosen over the City and County of Denver alone because it spans a wide range of neighbourhood income levels, from dense central tracts to affluent suburbs and sparsely populated mountain and plains tracts, and because CDPHE publishes both variables at census tract resolution for a matching period.

Income inequality is often cited as a reason to select Denver for work of this kind. It is worth being precise about what that justifies. The Gini coefficient measures dispersion of income within a population; it says nothing about whether high and low incomes are spatially segregated. A metro area can post a high Gini with incomes finely interleaved block by block, or with rich and poor districts cleanly separated. Inequality motivates the choice of study area. It does not predict the spatial pattern, which is what the analysis below actually measures.

Opioid and heroin mortality

The first dataset is a CDPHE feature layer of census tract polygons carrying the age-adjusted mortality rate per 100,000 persons for drug poisoning involving prescription opioid analgesics or heroin as the underlying cause of death, aggregated over 2015–2019. Age adjustment expresses the rate that would be observed if each tract's age-specific death rates applied to a population with a standard age distribution, which makes tracts with different age profiles comparable. The Colorado statewide age-adjusted rate of 9.25 per 100,000 serves as a reference value.

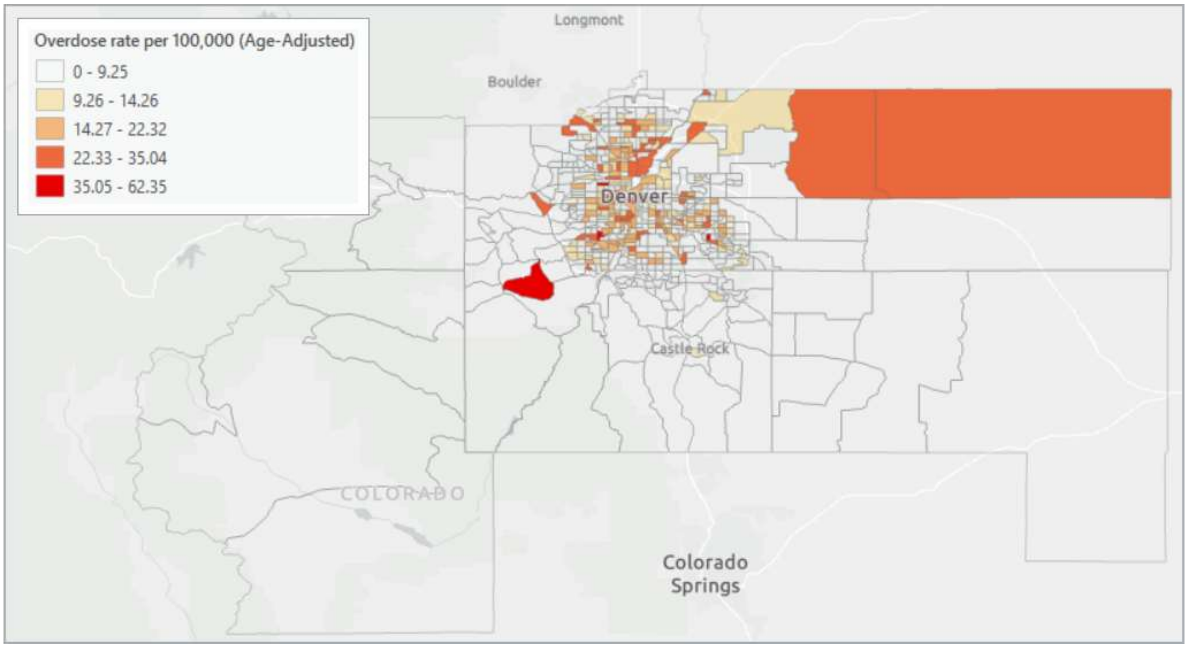

Figure 1. Age-adjusted opioid and heroin mortality per 100,000, 2015–2019. Tracts at or below the state rate of 9.25 are unshaded. The two most visually dominant features — the large orange tract on the eastern plains and the red tract in the western mountains — are among the least populated in the MSA.

Poverty

The second dataset, also from CDPHE, carries tract-level indicators from the 2015–2019 American Community Survey, including the percentage of all persons with annual income below the federal poverty line. This is an absolute threshold set nationally and adjusted for household size. It is not a measure of local inequality: a household's position relative to the line does not change because its neighbours earn more. A high-poverty tract bordering an affluent one registers as high-poverty, full stop.

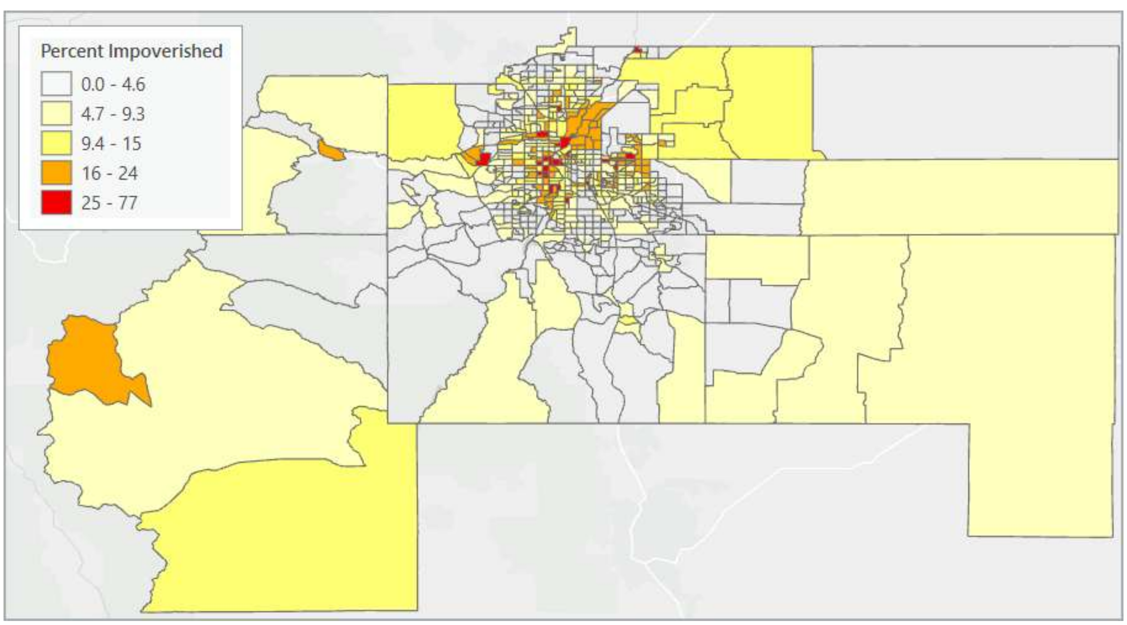

Figure 2. Percentage of persons with income below the federal poverty line by census tract, 2015–2019.

Method

Analysis followed a standard exploratory spatial data analysis sequence. A spatial weights matrix was built in GeoDa using queen contiguity, which treats tracts as neighbours if they share either an edge or a single corner point. Queen contiguity suits the small, irregular, tightly packed tracts of central Denver, where rook contiguity would sever genuine adjacencies at diagonal corners. Tract OBJECTID served as the unique identifier for the weights file.

Global Moran's I was then computed for each variable with 999 conditional permutations, producing a pseudo p-value with a floor of 0.001. Moran's I ranges from roughly −1 to +1: positive values indicate that similar values occupy neighbouring tracts, negative values indicate a checkerboard arrangement, and values near zero indicate spatial randomness. Local Indicators of Spatial Association then decomposed the global statistic to identify which individual tracts drive it, and in which direction.

LISA assigns each significant tract to one of four categories, and reading them correctly is essential to interpreting the maps that follow.

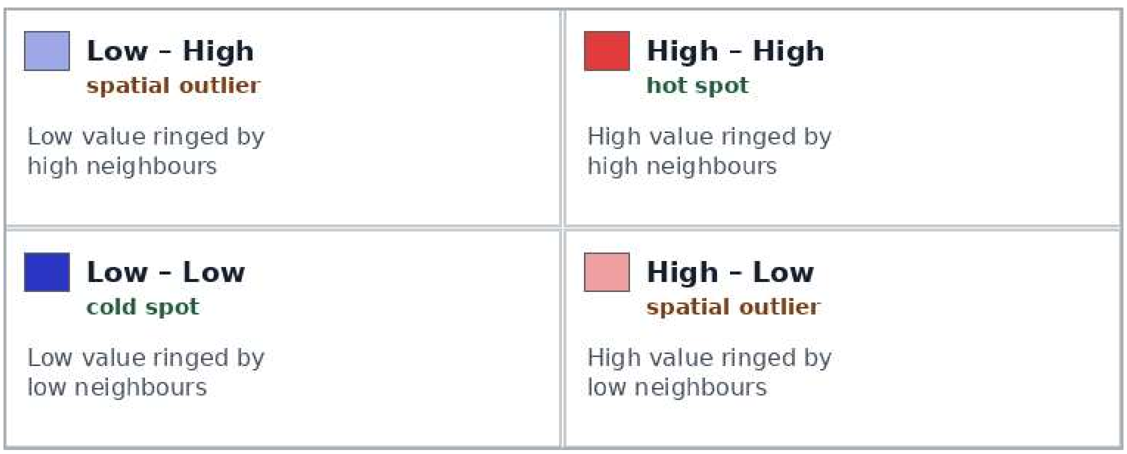

Figure 3. The four LISA categories. High–High and Low–Low are both clusters; the difference is whether the clustered values are high or low. Low–High and High–Low are spatial outliers — a tract that differs sharply from its surroundings. Tracts that reach no significance threshold are reported separately as Not Significant, and it is that category, not Low–Low, that indicates an absence of local clustering.

Results: Spatial Structure of Each Variable

Both variables show statistically significant positive spatial autocorrelation, but the strength differs substantially.

Table 1. Global Moran's I, queen contiguity, 999 permutations.

Variable	Moran's I	z-score	Pseudo p	Tracts in a significant cluster
Opioid and heroin mortality	0.184	8.24	≤ 0.001	121 of 620 (19.5%)
Poverty rate	**0.445**	19.85	≤ 0.001	**255 of 620 (41.1%)**

Statistical significance and effect size are separate matters here. With 620 tracts, a z-score of 8.24 makes the mortality result unambiguously significant, yet a Moran's I of 0.184 is a weak clustering effect: overdose deaths are more spatially organised than random, but only modestly so. Poverty at 0.445 is a different order of structure entirely, and the maps bear this out.

Table 2. LISA classification of all 620 census tracts.

LISA category	Mortality	Share	Poverty	Share
High–High (hot spot)	39	6.3%	83	13.4%
Low–Low (cold spot)	45	7.3%	154	24.8%
Low–High (outlier)	27	4.4%	14	2.3%
High–Low (outlier)	10	1.6%	4	0.6%
Not significant	499	80.5%	365	58.9%
Total	**620**	**100%**	**620**	**100%**

The local decomposition sharpens the contrast. Poverty places 237 tracts in a cluster of either kind against 84 for mortality, and produces half as many spatial outliers. Mortality's outlier count is high relative to its cluster count — 37 outliers against 84 clustered tracts — which is what a weakly organised pattern looks like at the local scale.

Opioid and heroin mortality

The Moran scatterplot plots each tract's standardised mortality rate against the average of its neighbours. The upward slope is the Moran's I value itself. The point cloud is tightly bunched near the origin with a small number of high-value tracts extending to the right, and no single observation exerts dominant leverage on the fit.

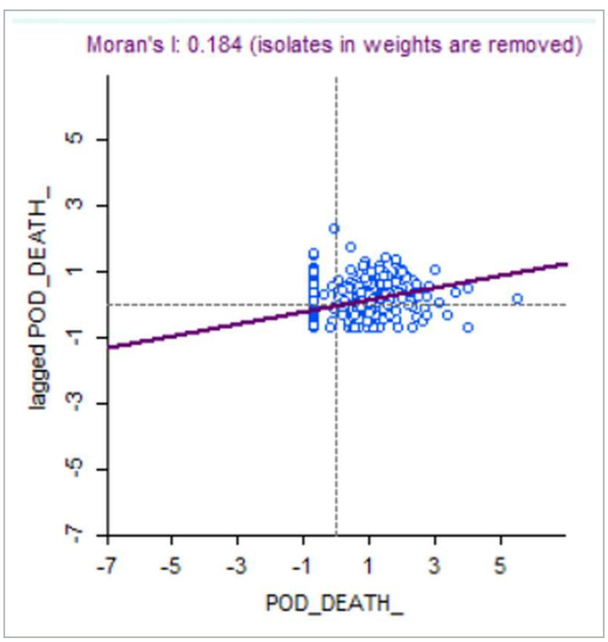

Figure 4. Moran scatterplot for age-adjusted opioid and heroin mortality (I = 0.184). Isolates — tracts with no queen-contiguous neighbours — are excluded from the weights.

The LISA cluster map locates the tracts driving that weak global figure. Only 39 tracts form high-mortality hot spots, concentrated in central and northeast Denver. A further 45 form cold spots, and 37 are spatial outliers, an unusually high proportion relative to the number of clusters. Nearly four fifths of the MSA reaches no significance threshold at all.

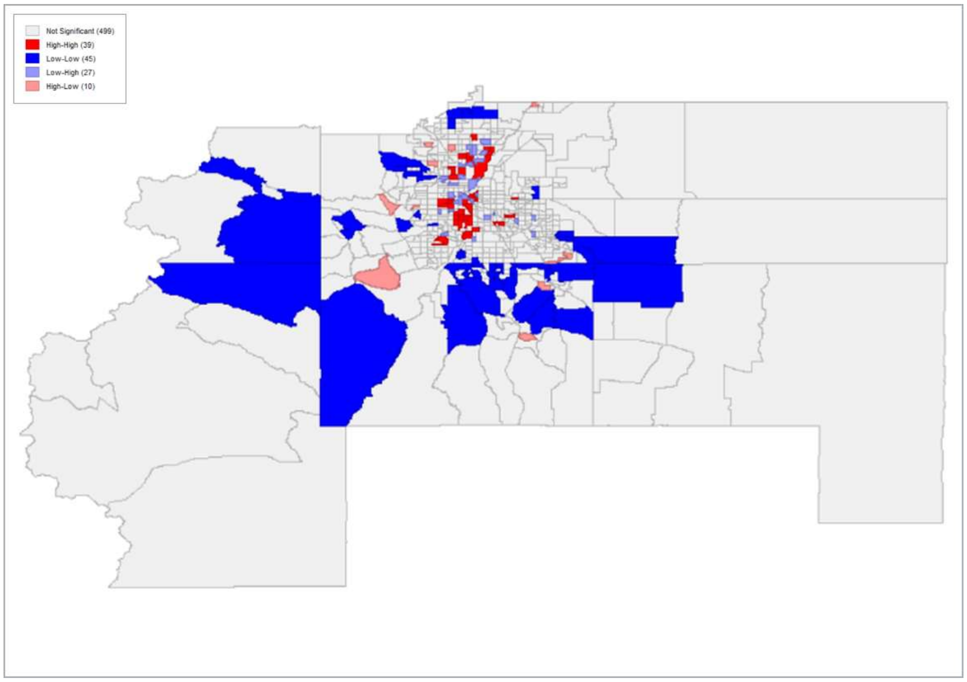

Figure 5. LISA cluster map for opioid and heroin mortality.

The significance map grades the same tracts by the strength of evidence from the permutation procedure rather than by cluster type. Only 14 tracts reach the 0.001 level, and 87 of the 121 significant tracts sit at the weakest 0.05 threshold — consistent with a pattern that is real but not strongly expressed.

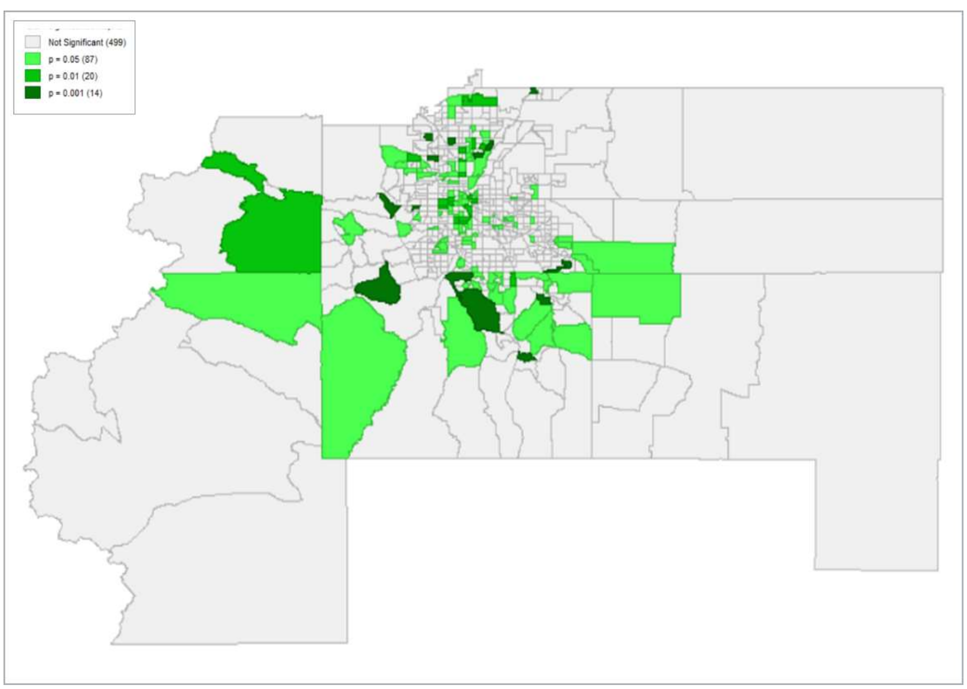

Figure 6. LISA significance map for opioid and heroin mortality, showing the strength of local evidence from the Monte Carlo permutation procedure.

Poverty

Poverty produces a visibly tighter scatterplot and a much steeper slope, with the point cloud tracking the fitted line closely across its full range.

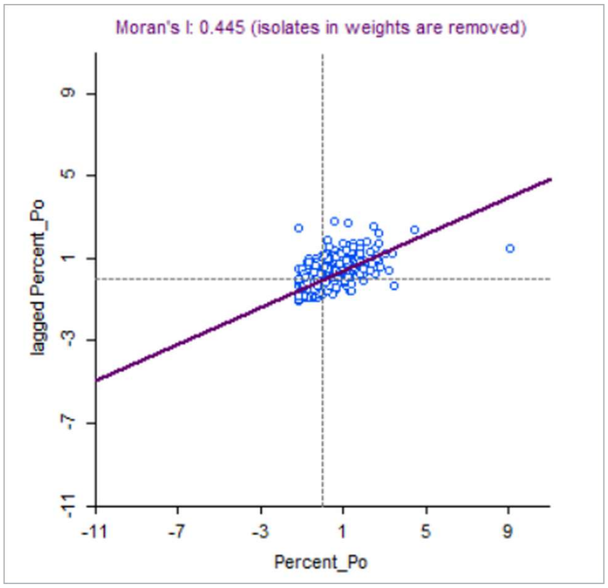

Figure 7. Moran scatterplot for percentage below the federal poverty line (I = 0.445).

The cluster map shows why. A contiguous band of 83 high-poverty tracts runs through central Denver, and an unusually large cold spot of 154 low-poverty tracts covers the suburban ring almost continuously. Only 18 tracts are spatial outliers, half the number seen for mortality. This is a strongly organised residential pattern: households of similar means occupy adjacent tracts.

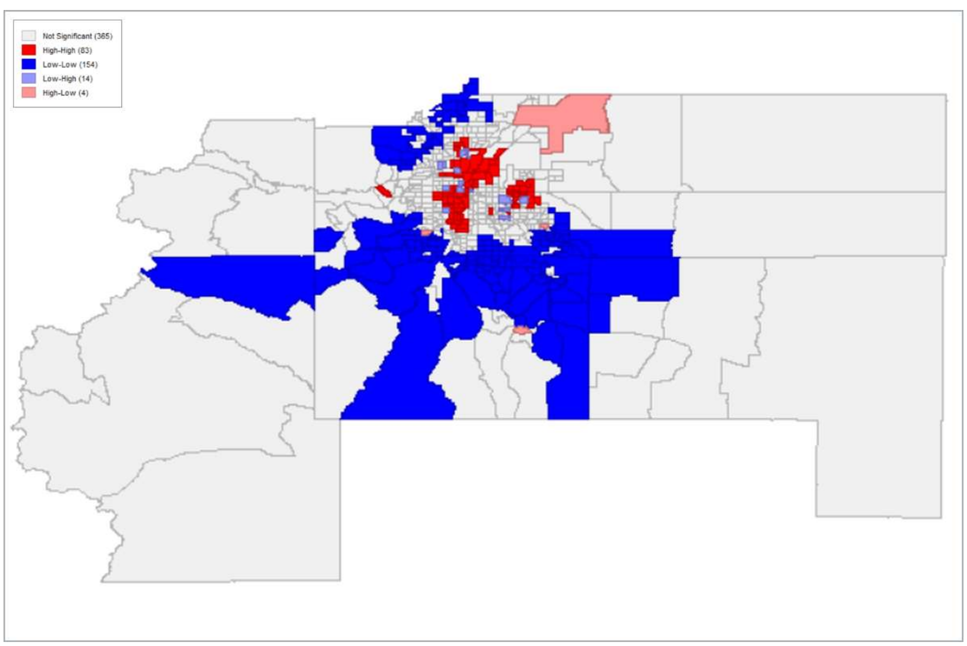

Figure 8. LISA cluster map for percentage below the federal poverty line.

Where the Two Patterns Coincide

To locate areas carrying both burdens, the two variables were reclassified onto a common one-to-five ordinal scale using each layer's own class breaks, converted to raster, and summed cell by cell with map algebra in ArcGIS Pro. The result is a composite index ranging from 2, where both variables sit in their lowest class, to 10, where both sit in their highest.

Table 3. Reclassification scheme. Each variable was binned to a 1–5 scale using its own distribution.

Class	Poverty (% below federal poverty line)	Mortality (age-adjusted deaths per 100,000)
1	0.0 – 4.6	0.00 – 9.25
2	4.7 – 9.3	9.26 – 14.26
3	9.4 – 15.0	14.27 – 22.32
4	15.1 – 24.0	22.33 – 35.04
5	24.1 – 77.0	35.05 – 62.35

The index is a co-occurrence surface, not a correlation. A value of 8 tells you both inputs were high; it does not tell you the two are statistically associated, and because the class breaks were drawn separately from two differently shaped distributions, a 4 in poverty is not equivalent in meaning to a 4 in mortality. Summing them applies *equal weight to two incommensurate scales*, which is a deliberate simplification and the principal caveat on the map below.

Read on those terms, the surface is informative. The highest values concentrate tightly in central Denver, where high-poverty tracts and high-mortality tracts overlap most consistently. The suburban ring falls into the lowest classes on both inputs. The large eastern plains tract reaching the top class does so on a small population base, discussed below.

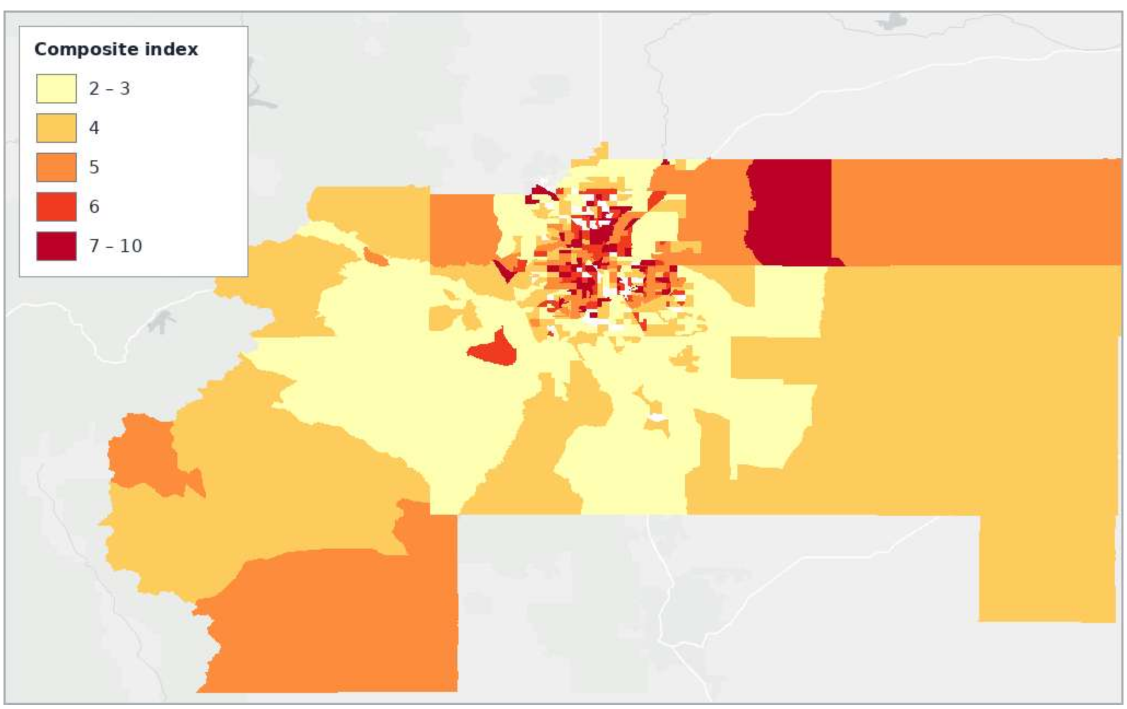

Figure 9. Composite index of poverty and opioid mortality classes, 2015–2019. Higher values indicate that both variables fall in higher classes within that tract.

Limitations

Six constraints bound how these results should be read.

- The bivariate question is not tested. Univariate Moran's I and LISA describe each variable separately. Nothing here quantifies the association between poverty and overdose mortality, and the composite index does not substitute for that measurement.
- Tract-level mortality rates are unstable. Overdose death is a rare event, and a rate computed over a small tract population swings widely on one or two additional deaths. The most visually striking features of Figure 1 — the eastern plains tract and the western mountain tract — are among the least populated in the MSA, so their extreme rates rest on very few events. Choropleth mapping compounds this by giving large, empty tracts the most ink.
- Findings are conditional on the census tract as the unit of analysis. Drawing the same data on different boundaries would produce different clusters and a different Moran's I. This is the modifiable areal unit problem, and census tracts do not resolve it — they are one arbitrary partition among many.
- Tract-level co-occurrence cannot support individual-level claims. The data cannot establish that the people who died of overdose were the people living below the poverty line.
- Isolates were dropped from the weights matrix. Tracts with no queen-contiguous neighbour are excluded from both Moran's I calculations, as noted on each scatterplot. The number excluded is small but the exclusion is not neutral, since isolated tracts are disproportionately rural.
- The MSA mixes structurally different geographies. Dense central Denver tracts, low-density suburbs, and vast mountain and plains tracts are analysed under one weights matrix and one set of class breaks. A separate analysis restricted to the urbanised core would likely produce a cleaner and more interpretable result.

Next Step: Testing the Relationship Directly

The obvious extension is a bivariate Moran's I and bivariate LISA, which correlate one variable in a tract against the spatially lagged values of a second variable in neighbouring tracts. That is the statistic this study's question actually calls for, and it is available in GeoDa on the same weights matrix already constructed here. It would return a single interpretable coefficient for the poverty–mortality relationship and a cluster map showing where high poverty is surrounded by high mortality, and vice versa.

Two further refinements would strengthen the result. Empirical Bayes smoothing of the mortality rates would shrink the unstable small-population estimates toward the regional mean and remove the rural artefacts that currently dominate Figure 1. Restricting the analysis to the urbanised core would remove the geographic heterogeneity noted above. Together these would convert a descriptive exploration into a defensible test.

References

1. Anderson, R. N., & Rosenberg, H. M. (1998). Age standardization of death rates: Implementation of the year 2000 standard. *National Vital Statistics Reports, 47*(3). National Center for Health Statistics.
2. Anselin, L. (1995). Local indicators of spatial association — LISA. *Geographical Analysis, 27*(2), 93–115.
3. Anselin, L., Syabri, I., & Kho, Y. (2006). GeoDa: An introduction to spatial data analysis. *Geographical Analysis, 38*(1), 5–22.
4. Coffman, C. H. (2019). *Colorado substance abuse trend and response task force: Thirteenth annual report*. Colorado Attorney General's Office. https://coag.gov/app/uploads/2020/02/ThirteenthAnnualSATFR_2019.pdf
5. Colorado Department of Public Health and Environment. (2022). *Drug poisoning or overdose involving Rx opioid analgesic or heroin mortality rate (census tract)* [Data set]. CDPHE Open Data. https://data-cdphe.opendata.arcgis.com
6. Colorado Department of Public Health and Environment. (2022). *Income / poverty (census tracts)* [Data set]. CDPHE Open Data. https://data-cdphe.opendata.arcgis.com
7. Moran, P. A. P. (1950). Notes on continuous stochastic phenomena. *Biometrika, 37*(1/2), 17–23.