Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Urban Complexity through the Spatial Structure of Amenities: An Information-Theoretic Approach

Joep Leeuwenburgh
University of Amsterdam
Rafiazka Millanida Hilman
r.m.hilman@uva.nl

University of Amsterdam

Research Article

Keywords:

Posted Date: November 11th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11676799/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Urban Complexity through the Spatial Structure of Amenities: An Information-Theoretic Approach

Abstract

Urban amenities represent the everyday infrastructure through which people organize their lives, from schooling and shopping to leisure, health care, and transport. The way these amenities are distributed across space shapes residential choice and neighborhood satisfaction, a relationship widely discussed in the literature on urban quality of life. Counting amenities nevertheless captures only part of this picture, since the complexity of a living space also depends on which kinds of amenities coexist (composition) and how they are arranged relative to each other (configuration). This study explores the extent to which information-theoretic measures can capture both dimensions and be translated into a comparable measure of urban liveliness. We combine OpenStreetMap points of interest with the official geographies of Statistics Netherlands to characterize 342 municipalities, 3,324 districts, and 13,943 neighborhoods, organizing amenities into 13 main and 53 sub-categories under three filter levels. Shannon entropy is compared with two spatial entropy measures, namely Leibovici and Altieri entropy, on synthetic point patterns and on the empirical data. Our findings reveal that Shannon entropy is insensitive to spatial arrangement, Leibovici entropy responds only when clustering follows amenity categories, and Altieri entropy responds to clustering in both cases. Within each filter level the three measures are strongly correlated, whereas the choice of which amenities to include alters the results considerably, especially for municipalities. Weighting spatial entropy by the degree of urbanity yields a measure of urban liveliness whose distribution becomes bimodal, separating urban from rural districts. Local Moran's I further reveals hot spots of liveliness in the large Dutch cities and cold spots across much of the north, with 36.83% of districts being statistically significant. The results are condensed into an interactive dashboard that identifies areas with comparable amenity profiles, offering a transparent tool for residential choice and spatial planning.

Keywords— urban amenities, spatial entropy, urban complexity, urban liveliness, spatial autocorrelation, OpenStreetMap

1 Introduction

Cities can be understood as complex systems in which people, activities, and infrastructure interact across multiple scales (Batty, 2013; Bettencourt, 2013). Among the elements that make up this system, urban amenities, such as schools, shops, restaurants, clinics, sports facilities, and transit stops, form the everyday infrastructure through which residents organize their lives. Their presence shapes the attractiveness of cities as places of consumption (Glaeser et al., 2001) and, at a finer scale, the attractiveness of individual neighborhoods as places to live. Choosing where to live therefore involves far more than housing prices and floor area. Residents also ask

what everyday life in a place looks like: some value a bustling neighborhood with restaurants and cafés within walking distance, while others prefer a quieter area with primary and secondary schools nearby.

Urban amenities play a significant role in the choice of residential area and in the satisfaction experienced there (Kelly, 2006). Most surveyed residents are willing to live in denser and more built-up areas in exchange for more and better amenities (Randall, 2008), and location, in particular proximity to parks and to recreational and educational amenities, is an essential factor in residential decisions (Saville-Smith, 2010; Allen, 2016). Amenities thus affect not only where people decide to live but also how satisfied they are with their neighborhood afterwards (McCrea et al., 2011; Mccrea and Walters, 2012). The same intuition underlies the recent policy interest in proximity-based planning, such as the "15-minute city", in which essential urban functions should be reachable within a short walk or bicycle ride (Moreno et al., 2021). Housing platforms typically list nearby schools, daycare, public transport, and supermarkets next to basic sociodemographic indicators. However, they rarely condense this information into a concise index, nor do they allow neighborhoods to be compared with one another, which makes the question of where to live particularly difficult for newcomers who are unfamiliar with an area.

Taking a broader view, the presence of amenities is also tied to the notion of urban vibrancy, referred to in this study as urban liveliness. It captures the dynamic interplay of human presence, activities, and interactions within urban space. The idea is rooted in the work of Jacobs (1961), who argues that lively, diverse, and attractive urban areas foster socioeconomic activity, and in urban design theory linking urbanity, vitality, and the mix of uses (Montgomery, 1998; Lynch, 1984). The built environment, and points of interest (POIs) such as urban amenities in particular, supports these activities and interactions (Yue et al., 2021). On top of that, a dense population is associated with urban vitality and quality of life (Jacobs, 1961; Wu and Niu, 2019). A growing body of work in computational urban science measures vitality with emerging data sources, such as mobile phone records, social media check-ins, and street-view imagery (Jin et al., 2017; Wu and Niu, 2019; Tu et al., 2020; Li et al., 2022; Ouyang et al., 2022). POIs have likewise been combined with human mobility to discover the functional regions of a city (Yuan et al., 2012) and to model flows between urban areas (Hilman, 2025). These approaches typically observe liveliness through human activity, which requires data that are often proprietary, costly, or unavailable at national scale. Rather than observing liveliness directly, one may ask how much of it can be inferred from the structure of the amenity landscape itself, using open data available for an entire country.

The complexity of an urban area cannot be reduced to the number of amenities it contains. Two districts with identical counts may differ in the variety of functions they host (composition) and in how these functions are arranged relative to each other (configuration). Entropy offers a natural way to distill both into a single continuous value, denoting the amount of information needed to describe a distribution (Saraiva, 2023). Shannon entropy summaries the diversity of a categorical distribution (Shannon, 1948) and is the most common basis of land use mix indices in urban research (Song et al., 2013), but it does not account for space. Leibovici (2009) builds on Shannon entropy by considering co-occurrences of categories at multiple orders, which opens the way from Shannon entropy towards measures that account for spatial dependence, broadly referred to as spatial entropy. Altieri et al. (2019b) compare several of these measures, namely the spatial entropy of Batty (1974), that of Karlström and Ceccato (2002), and spatial mutual information and residual entropy, and apply them to the measurement of urban sprawl. In a related contribution, Altieri et al. (2019a, 2023) decompose entropy into spatial mutual information, which captures the role of space in a distribution, and spatial residual entropy, which captures the heterogeneity that remains once space is accounted for. Entropy has been used in this way to trace urban sprawl (Cabral et al., 2013) and the expansion of Italian metropolitan

areas (Altieri and Cocchi, 2022).

Apart from describing single areas, comparing areas requires knowing whether similar values cluster in space. Spatial autocorrelation relates the value of a variable in one area to the values of the same variable in neighboring areas. Global Moran's I, ranging from -1 to 1, indicates whether values are clustered or dispersed. Anselin (1995) introduces local indicators of spatial association (LISA), which decompose global statistics such as Moran's I into local contributions and evaluate their significance through Monte Carlo simulation. This makes it possible to locate clusters rather than only to detect their presence. The framework was later extended to multivariate settings with the local Geary's C statistic, conceptualized as a weighted distance in multivariate attribute space between an observation and its neighbors (Anselin, 2019).

Given the lack of formulation on how entropy measures relate to the everyday amenity landscape of residential areas, the question of whether different measures capture the same notion of urban complexity remains unclear. Spatial entropy has so far been applied mainly to land cover and urban expansion (Cabral et al., 2013; Altieri et al., 2019b; Altieri and Cocchi, 2022), while land use mix indices rely almost exclusively on Shannon-type formulations that ignore configuration (Song et al., 2013). It is also unclear whether such measures can be translated into an interpretable measure of liveliness that is comparable across spatial scales, and how sensitive they are to the selection of amenities. In this research, we take a step forward by coupling open data on urban amenities with official geographies at national scale. OpenStreetMap (OSM) provides a free, volunteered map of POIs (OpenStreetMap contributors, 2017), while Statistics Netherlands (*Centraal Bureau voor de Statistiek*, CBS) provides boundaries and sociodemographic attributes at three nested scales (Centraal Bureau voor de Statistiek, 2023). Combining the two allows us to describe the spatial configuration of amenities and, at the same time, to contextualize it with the sociodemographic configuration of living space.

It is in our interest to establish the line between the composition and configuration of urban amenities and the liveliness of living space, as well as to measure the extent to which these can be used to compare areas. We specify three questions for this purpose: (1) How do Shannon, Leibovici, and Altieri entropy respond to compositional and configurational variation in the distribution of urban amenities? (2) Can spatial entropy, weighted by the degree of urbanity, yield a measure of urban liveliness that is comparable across areas? (3) To what extent can this measure, together with spatial autocorrelation, reveal clusters of liveliness and identify areas with comparable amenity profiles? We take the Netherlands as a case study, since official boundaries and attributes are available for every municipality, district, and neighborhood.

The contribution of this study is threefold. Conceptually, it frames urban complexity as the joint outcome of amenity composition and spatial configuration, and shows that the two are captured by different information-theoretic measures. Methodologically, it provides a reproducible pipeline based exclusively on open data that collects, categories, filters, and evaluates amenities for every spatial unit of a country at three scales, and benchmarks three entropy measures on controlled synthetic patterns before applying them empirically. Substantively, it delivers a national map of urban liveliness and its hot and cold spots for the Netherlands, together with an interactive dashboard that ranks comparable areas. The remainder of this paper is organized as follows. Section 2 describes the data. Section 3 introduces the analytical framework and presents the results for entropy, urban liveliness, spatial autocorrelation, and area similarity. Section 4 discusses the findings in light of the literature, outlines their implications and limitations, and concludes.

2 Data description

Our choice of the Netherlands as a case study is motivated by the fact that both administrative boundaries and amenity data are available at fine granularity for the whole country. The Netherlands is one of the most densely populated countries in Europe and combines a polycentric urban core with extensive rural regions, which provides a wide range of urban conditions within a single, consistently documented territory. It serves our purpose of stepping beyond amenity counts by synthesizing official geographies with volunteered geographic information. In the following, we present the administrative and amenity data, followed by the categorization and filtering of amenities.

Administrative data: We use the *Wijk- en buurtkaart 2023* published by CBS (Centraal Bureau voor de Statistiek, 2023). This publicly available resource contains shapefiles of the Netherlands at three nested scales, namely *gemeenten* (municipalities), *wijken* (districts), and *buurten* (neighborhoods). Each dataset contains the geometries of its units together with sociodemographic and geographic attributes, including the degree of urbanity (*stedelijkheid*). The latter is a five-class variable ranging from very highly urban (1) to not urban (5). In total, the data account for 342 municipalities, 3,324 districts, and 13,943 neighborhoods. Using three nested scales allows us to examine whether the results depend on the size and delineation of spatial units, a concern known as the modifiable areal unit problem (Openshaw, 1984). We refer to these scales by their English names hereafter.

Amenity data: Amenities are retrieved from OSM, an open-source project that builds a free and editable map of the world (OpenStreetMap contributors, 2017). Volunteers contribute data both by digitizing online resources and by surveying on the ground. OSM data are structured as nodes, ways, and relations. Nodes represent single points defined by latitude and longitude, ways are ordered lists of nodes representing lines and polygons, and relations describe connections between elements. Each element carries tags in the form of key–value pairs, which can be combined to query specific features. For example, looking up nodes tagged `amenity=restaurant` and `cuisine=italian` returns all Italian restaurants in an area of interest. Geometries are loaded into a GeoDataFrame with GeoPandas (Jordahl, 2014), and for each unit, the bounding box of its geometry is used to query shops, amenities, leisure, railway stations, healthcare, sports, crafts, public transport, and selected tourism features through the Overpass API. Data were retrieved in June 2024, and the full query is reported in Supplementary Section B. Ways and relations are reduced to their centroids, since entropy measures operate on point patterns rather than on lines or polygons. The completeness of OSM varies across regions and feature types (Haklay, 2010; Barrington-Leigh and Millard-Ball, 2017). We return to the implications of this variation in Section 4.3.

Categorization: For every element, a primary tag identifies the type of feature (e.g.: amenity, shop, or sport), and a secondary tag is extracted by using the primary tag as a key in the tag dictionary. Based on these tags, each amenity is assigned a main category (L0) and a subcategory (L1). There are 13 L0 categories, subdivided into 53 L1 categories (Table 1). The two levels allow us to evaluate complexity both at a coarse functional level (e.g.: education versus sustenance) and at a finer level (e.g.: restaurants versus cafés). An example of the category tree, together with the distribution of amenities by category and scale, is provided in Supplementary Section A.

Filtering: Not every amenity is equally relevant to every resident, and some categories can paint a distorted picture of an area. Private transportation, for instance, consists mainly of car parking spaces, which tend to be mapped individually. A neighborhood may therefore appear to

Table 1: **Amenity categories**. Number of sub-categories (L1) within each main category (L0).

L0 category	**No. of L1 categories**
Education	5
Entertainment, arts and culture	6
Facilities	1
Financial	2
Healthcare	9
Places of worship	1
Private transportation	4
Public service	5
Public transportation	3
Shopping	12
Sports	4
Sustenance	4
Waste management	1
Total	**53**

be dominated by parking even though this category is irrelevant to residents without a car. To assess the sensitivity of the results to such categories, three cumulative filter levels are constructed (Table 2). Filter level 0 removes only uncategorized elements. Filter level 1 additionally removes categories that relate to infrastructure rather than to everyday use, together with residual sub-categories. Filter level 2 further removes categories that are of interest to specific groups of residents only. The complete list is given in Supplementary Table S1.

Table 2: **Filter levels**. Each level includes the removals of the previous level. "Other" denotes the residual sub-category of an L0 category.

Level	**Removed amenities**
0	Uncategorized elements
1	Private transportation, facilities, and waste management; "other" healthcare, shopping, and public service
2	Financial; "other" entertainment, arts and culture; shopping for clothing and accessories, crafts, house and interior, media appliances and hardware, and mobility

3 Results

Our study is motivated by the need to characterize living space by the structure of its amenity landscape rather than by amenity counts alone. This structure is embedded in the composition and configuration of amenities and most likely intertwined with the degree of urbanity of an area. To incorporate both aspects, we propose a framework that moves from entropy, through urban liveliness, to spatial autocorrelation and area similarity (Figure 1). We first formalize the framework (Section 3.1), then evaluate the entropy measures (Section 3.2), transform them into

urban liveliness (Section 3.3), examine its spatial autocorrelation (Section 3.4), and finally derive area similarity and the dashboard (Section 3.5).

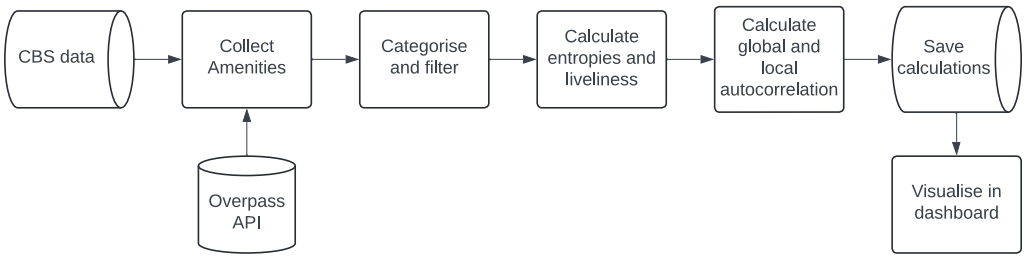


Figure 1: **Pipeline design**. For each municipality, district, and neighborhood, amenities are collected from OpenStreetMap within the CBS geometry, categorized into main (L0) and sub-categories (L1), and filtered at three levels. Shannon, Leibovici, and Altieri entropy are then computed and transformed into a measure of urban liveliness, followed by global and local spatial autocorrelation, area similarity, and the dashboard.

3.1 Analytical framework

Let $\mathcal{A}_s$ denote the set of spatial units at scale $s \in \{\text{municipality}, \text{district}, \text{neighborhood}\}$. Each area $a \in \mathcal{A}_s$ contains a marked point pattern

$$\mathcal{P}_a = \{(\mathbf{u}_j, c_j)\}_{j=1}^{N_a} \tag{1}$$

where $\mathbf{u}_j$ denotes the location of amenity j, c_j its category at level L0 or L1 after filtering, and N_a the number of amenities in area a. We distinguish two dimensions of this pattern. *Composition* refers to the distribution of categories $\{c_j\}$, irrespective of where the amenities are located. *Configuration* refers to how categories are arranged relative to each other in space, that is, which categories tend to co-occur at short distances and how this co-occurrence varies across the area. An area may thus be compositionally diverse but configurationally segregated, for instance when shops, schools, and restaurants are all present but concentrated in separate parts of the area.

Each entropy measure H maps $\mathcal{P}_a$ onto a scalar value. The measures differ in which dimension they capture, as summarized in Table 3 and evaluated empirically in Section 3.2.1. Urban liveliness L_a then combines the entropy of an area with its degree of urbanity, spatial autocorrelation relates L_a to the values of neighboring areas, and area similarity compares areas in the joint space of liveliness and autocorrelation.

3.2 Entropy of urban amenities

In information theory, entropy measures the degree of uncertainty of random events (Saraiva, 2023). This in turn corresponds to the amount of information needed to describe the distribution of a categorical variable. The more information we need, the more heterogeneous the distribution and the higher the entropy. In the context of this study, a higher entropy implies an area with

Table 3: **Entropy measures compared in this study**. Responses to clustering refer to the synthetic experiment in Section 3.2.1.

Measure	**Random variable**	**Spatial component**	**Response to clustering by category**	**Response to clustering of mixed categories**
Shannon (Shannon, 1948)	Category X	None	None	None
Leibovici (Leibovici et al., 2014)	Co-occurrence Z within distance d	Distance threshold	Responsive	Barely responsive
Altieri (Altieri et al., 2019a)	Co-occurrence Z and window W	Spatial windows	Responsive	Responsive

a more diverse and more evenly distributed set of amenities. Three measures are examined, namely Shannon, Leibovici, and Altieri entropy.

Shannon entropy (Shannon, 1948) is formalized as:

$$H(X) = -\sum_{i=1}^{n} p_i \log_2 p_i \tag{2}$$

where p_i denotes the probability of observing category $X = x_i$ among the n categories. Using base two ensures that the outcome can be interpreted as the number of bits needed to describe the distribution. $H(X)$ is zero when all amenities belong to a single category and reaches its maximum $\log_2 n$ when all categories are equally represented. Because Eq. (2) only considers category proportions, Shannon entropy captures composition, namely how many amenity types are present and how evenly they are distributed, but does not account for their spatial arrangement.

To account for configuration, Leibovici et al. (2014) introduce a variable Z that denotes co-occurrences of categories among points located within a distance d of each other. *Leibovici entropy* is then computed over the distribution of co-occurrences as:

$$H(Z \mid d) = \sum_{r=1}^{I^m} p(z_r \mid d) \log \frac{1}{p(z_r \mid d)} \tag{3}$$

where the sum runs over every ordered combination of categories, I denotes the number of categories present, and m denotes the size of the combination, resulting in I^m possible co-occurrence types.

Because co-occurrences are counted only within distance d, the measure is sensitive to whether amenities of the same or of different categories are located close to one another.

Altieri et al. (2019a) expand on this by introducing a variable W that represents a set of spatial windows over the study area. *Altieri entropy* splits the entropy of Z into spatial mutual information $MI(Z, W)$ and spatial residual entropy $H(Z)_W$:

$$H(Z) = \sum_{r=1}^{R} p(z_r) \log \frac{1}{p(z_r)} = MI(Z, W) + H(Z)_W \tag{4}$$

with

$$H(Z)_W = \sum_{k=1}^{K} p(w_k)\, H(Z \mid w_k), \qquad MI(Z, W) = \sum_{k=1}^{K} p(w_k) \sum_{r=1}^{R} p(z_r \mid w_k) \log \frac{p(z_r \mid w_k)}{p(z_r)} \tag{5}$$

where R is the number of co-occurrence types and K the number of windows. $MI(Z, W)$ measures how much the information from window w_k reduces the uncertainty of Z, while $H(Z)_W$ represents the remaining uncertainty within windows, averaged over all windows. In substantive terms, a high $MI(Z, W)$ indicates that the mix of amenities differs strongly between parts of an area, whereas a high $H(Z)_W$ indicates that amenities are well mixed within each part.

Shannon entropy is computed from the empirical category proportions with the entropy function of SciPy (Virtanen et al., 2020). Leibovici and Altieri entropy are computed with the SpatialEntropy package (Zheng, 2021), a Python adaptation of the R toolbox SpatEntropy (Altieri et al., 2023). All measures are computed for both L0 and L1 categories and for every filter level, resulting in 18 entropy values per area (three measures, two category levels, and three filter levels). For comparisons across areas and in the dashboard, values are normalized within each combination of scale, measure, category level, and filter level.

3.2.1 Behavior on synthetic point patterns

To build intuition about how the measures react to different point distributions, two sets of 300 synthetic points are constructed. Each point is given one of seven arbitrary labels representing amenity categories, and the labels are uniformly distributed in both sets. For each set, the degree of spatial clustering is gradually increased. In the first set, points gradually cluster into seven clusters of the same label. In the second set, the spatial clustering is identical but the labels are shuffled, resulting in seven clusters of random labels (Figure 2). This procedure is intended to mirror the empirical situation, in which amenities are not always clustered by category. Shopping streets, for example, concentrate shops, restaurants, and services in the same location, whereas business parks or school campuses concentrate a single category. Because the composition of labels is fixed by construction, any change in entropy can be attributed to configuration alone.

Figure 2 plots the three entropy values against the degree of clustering. For both sets, Shannon entropy stays the same at every degree of clustering. It is to be expected, since the distribution of labels does not change and Shannon entropy cannot account for spatiality. Leibovici entropy is higher overall and responds to clustering in the first set, as co-occurrences of the same label increase when points cluster. When the labels are shuffled, however, the measure becomes almost unresponsive because the composition of co-occurring pairs barely changes. Interestingly, Altieri entropy shows the same sensitivity for both sets, as indicated by the lines in Figure 3 running almost on top of each other. The additional window variable W and the decomposition into spatial mutual information and residual entropy make the measure more sensitive to clustering in general, which also explains its inconsistency around cluster sizes 14 and 20. Consequently, Shannon entropy reflects composition, Leibovici entropy reflects configuration only when it follows category labels, and Altieri entropy captures configuration regardless of labels (Table 3).

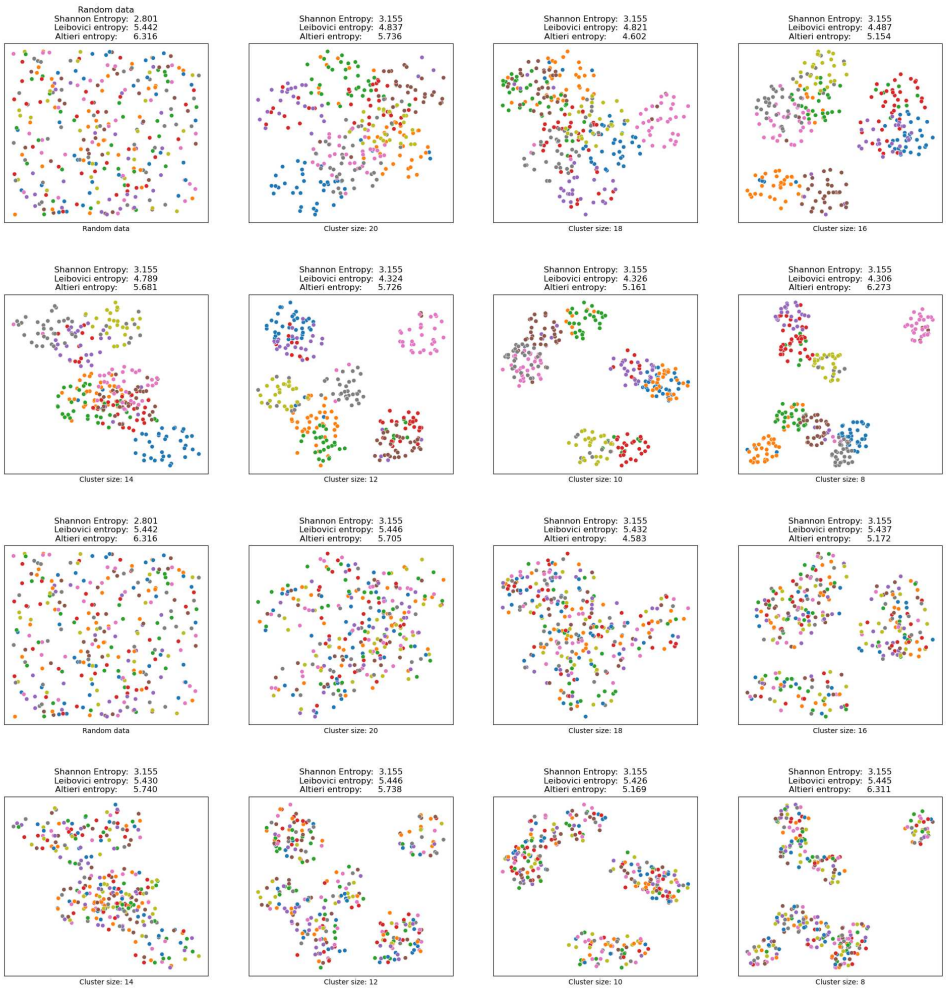


Figure 2: **Synthetic point patterns with increasing clustering**. Each panel shows 300 points with seven uniformly distributed labels (colors) and the corresponding entropy values. Rows one and two show clusters formed by a single label each, whereas rows three and four show the same spatial clusters with shuffled labels. Clustering increases from left to right.

Since empirical amenity patterns combine both types of clustering, Altieri entropy is retained as the primary measure in the remainder of the analysis.

3.2.2 Correlations between measures

After applying the three measures to all municipalities, districts, and neighborhoods, correlation matrices are computed between the normalized values for every filter level and for both L0 and L1 categories (Supplementary Figs. S4–S6). A recurring pattern can be observed at every scale. Each measure within a filter level correlates strongly with the other measures within that filter

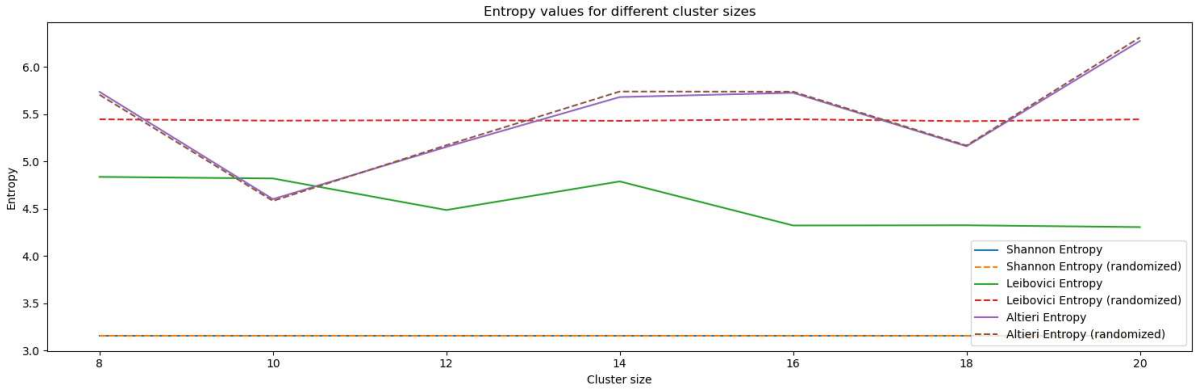


Figure 3: **Response of the entropy measures to clustering**. Shannon, Leibovici, and Altieri entropy for increasing degrees of clustering, for clusters formed by single labels and for clusters of shuffled labels. Shannon entropy stays constant, Leibovici entropy responds only when clusters follow labels, and Altieri entropy responds in both settings.

level. On top of that, filter levels 1 and 2 correlate strongly with each other, whereas filter level 0 tends to be more weakly correlated with the other two. This is most evident at the scale of municipalities, where the values at filter level 0 correlate only weakly with the rest. The pattern is consistent across L0 and L1 categories.

Two observations follow from this pattern. First, although the synthetic experiment reveals clear conceptual differences between the measures, their rankings of real areas largely agree once the set of amenities is fixed. This suggests that, in empirical amenity landscapes, composition and configuration tend to co-vary: areas with many categories also tend to mix them in space. Second, the categories removed at filter level 1, in particular individually mapped parking spaces, can dominate the amenity count of an area and thereby distort its entropy. The choice of which amenities to include therefore matters more than the choice of entropy measure when ranking areas. The weaker agreement at the municipal scale further suggests that large units aggregate heterogeneous sub-areas, so that the amenity structure of living space is better resolved at the finer scales of districts and neighborhoods, in line with the modifiable areal unit problem (Openshaw, 1984).

3.2.3 Amsterdam case

To ground the measures in concrete places, three districts are sampled from the municipality of Amsterdam with low, medium, and high Altieri entropy for L0 categories under filter level 1 (Figure 4). The first district, Tuindorp Buiksloot, has very few amenities and little variation among them, resulting in a low entropy of 2.114. The second district, Rijnbuurt, has many amenities, but their distribution is dominated by the entertainment, arts and culture category, resulting in a medium value. The third district, Oostelijke eilanden/Kadijken, does not have the most amenities of the three, but they are more varied and more evenly distributed, resulting in the highest value. Spatial entropy therefore enables us to contextualize space not only by the number and placement of amenities but also by the balance between amenity categories. The comparison between Rijnbuurt and Oostelijke eilanden/Kadijken in particular illustrates that a larger number of amenities does not guarantee a higher complexity when a single category dominates.

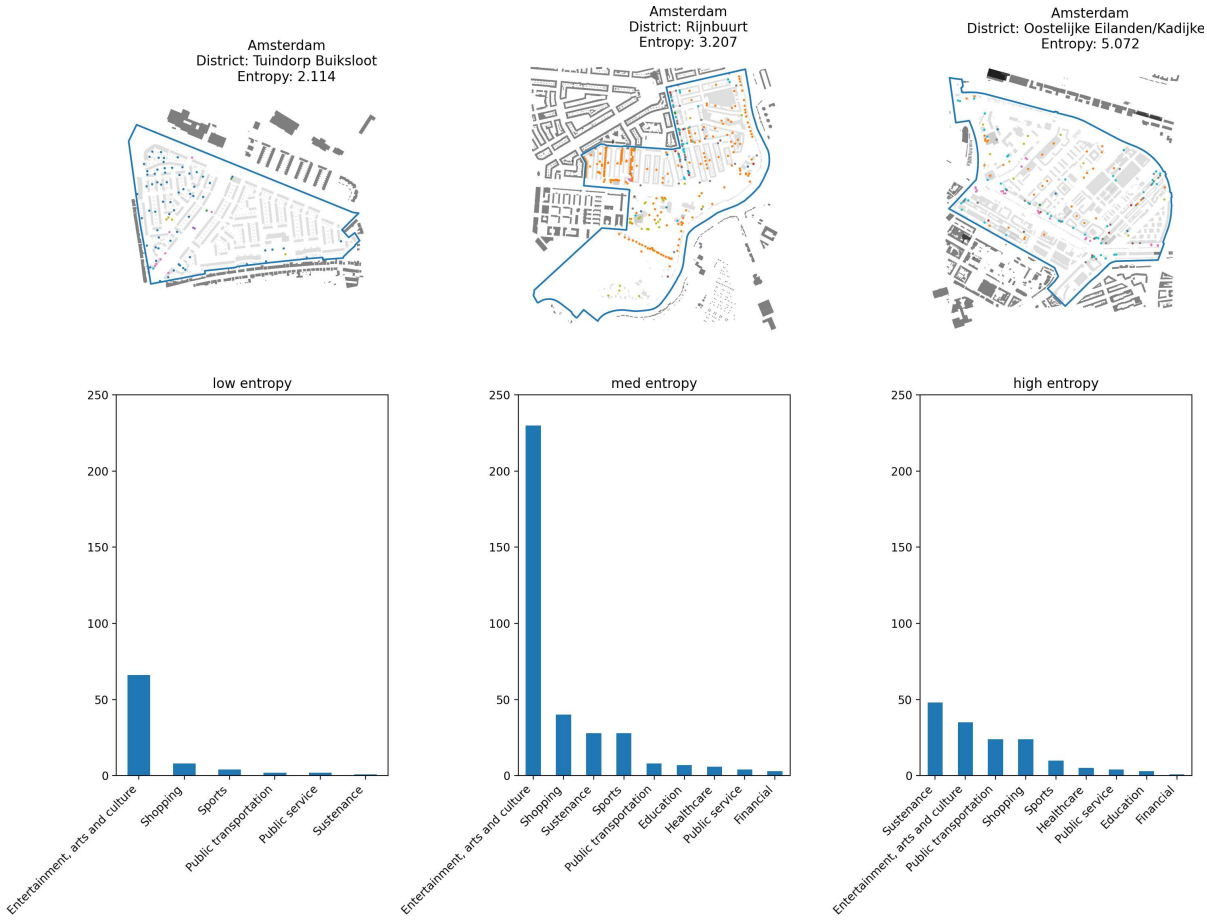


Figure 4: **Amenity landscapes of three Amsterdam districts**. Map of amenities and bar plot of amenities per L0 category for districts with low (Tuindorp Buiksloot, entropy 2.114), medium (Rijnbuurt), and high (Oostelijke eilanden/Kadijken) Altieri entropy (L0 categories, filter level 1). A high value requires amenities that are varied and evenly distributed across categories, not merely numerous.

3.3 Urban liveliness

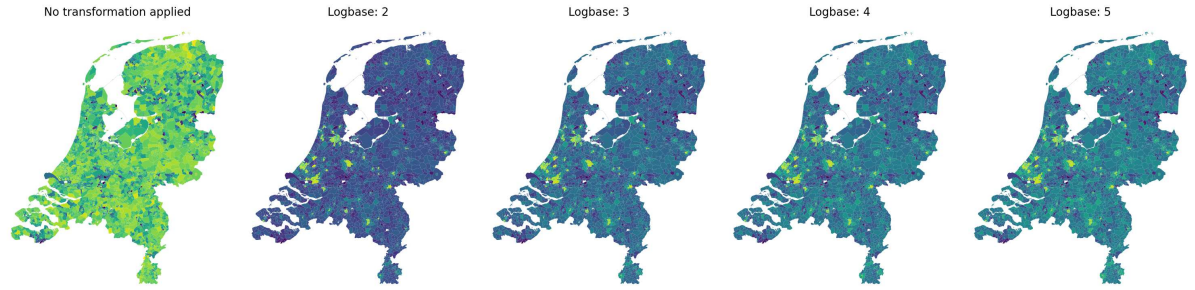
Entropy only conveys information about the distribution of amenities within an area. As a consequence, two areas with a similar distribution but very different degrees of urbanity or size end up with similar entropy values. A small village center with a supermarket, a school, a pharmacy, and a bus stop may thus obtain an entropy comparable to that of an inner-city district, although the intensity of urban life in the two areas differs considerably. To move from complexity to liveliness, the entropy value H_a of area a is weighted by its degree of urbanity:

$$L_a = \frac{H_a}{\log_b(\text{sted}_a) + 1} \tag{6}$$

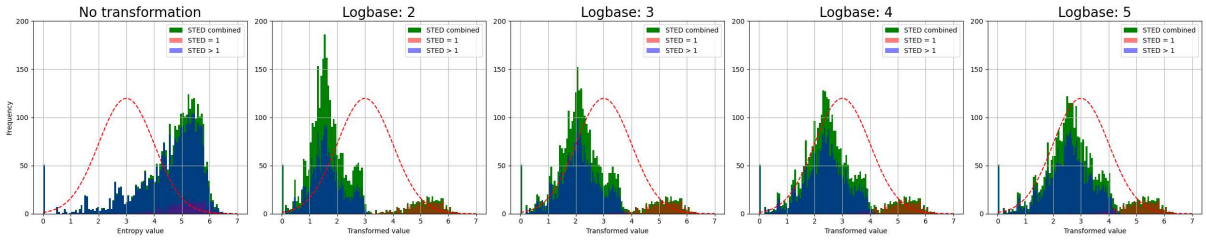
where $\text{sted}_a \in \{1, \ldots, 5\}$ denotes the CBS urbanity class (*stedelijkheid*), ranging from 1 (very highly urban) to 5 (not urban), and b denotes the logarithmic base. Adding one to the logarithm ensures that the most urban areas are not weighed down ($\log_b 1 = 0$), while the least urban areas are divided by at most two when $b = 5$. The concave shape of the logarithm penalizes

the step from very highly to highly urban areas more than the steps between the less urban classes (Supplementary Figure S7). We refer to L_a as the *urban liveliness* of area a. Eq. (6) thus captures the balance between the complexity of amenities and the density of urban features within a given space.

To find the appropriate base, Eq. (6) is applied with several bases to the Altieri entropy of districts (L0 categories, filter level 1), and the resulting maps and distributions are inspected (Figure 5). As seen in Figure 5b, increasing the base makes the distribution of transformed values increasingly resemble a bimodal distribution, the larger part of which approximates a normal distribution. The divide between the two modes reflects the shift from rural to urban areas. After the transformation, the large Dutch cities retain higher values than rural areas (Figure 5a). It indicates that weighting by urbanity separates urban from rural liveliness while preserving the variation in amenity complexity within each group. Accordingly, base 5 is used in the remainder of the analysis. A base equal to the number of urbanity classes also offers a convenient interpretation, since the weight then ranges exactly from one to two.



(a)



(b)

Figure 5: **From entropy to urban liveliness**. The transformation of Eq. (6) applied with different logarithmic bases to the Altieri entropy of districts (L0 categories, filter level 1). After the transformation, large cities retain higher values than rural areas (a). With increasing base, the distribution of values becomes bimodal, with the divide reflecting the transition from rural to urban districts (b).

3.4 Spatial autocorrelation of urban liveliness

Spatial autocorrelation relates values of a variable in specific areas to values of the same variable in other areas (Rey et al., 2023). Positive values indicate that similar values are close together, negative values that similar values are far apart, and zero values that no spatial pattern is present. Spatial weights are constructed with Queen contiguity (Getis and Aldstadt, 2004), meaning that

every area is weighted against the areas with which it shares a border or a vertex, and the weights of each area are row-standardized to sum to one. Using this pattern instead of a nearest-neighbor approach ensures that the outskirts of an area are not weighted against areas that lie several units away. All statistics are computed with PySAL (Rey and Anselin, 2007).

Global Moran's I captures the overall tendency within a study area (Rey et al., 2023). For a municipality or district, a high value implies that lively and quiet sub-areas form hot and cold spots, whereas a low value implies that they are more evenly spread. It is formalized as:

$$I = \frac{n}{\sum_i \sum_j w_{ij}} \frac{\sum_i \sum_j w_{ij}\, z_i\, z_j}{\sum_i z_i^2} \tag{7}$$

where n denotes the number of areas, z_i the deviation of the target variable from its mean in area i, and w_{ij} the spatial weight between areas i and j.

Global autocorrelation reveals whether values are clustered but not where the clusters are located (Rey et al., 2023). *Local Moran's I* (Anselin, 1995) therefore relates each area to its neighbors:

$$I_i = \frac{z_i}{m_2} \sum_j w_{ij} z_j, \qquad m_2 = \frac{\sum_i z_i^2}{n} \tag{8}$$

where $z_i = y_i - \bar{y}$ and m_2 denotes the variance of the target variable. With the spatial lag $\tilde{z}_i = \sum_j w_{ij} z_j$, each area is classified as high–high (HH) if $z_i > 0$ and $\tilde{z}_i > 0$, low–low (LL) if $z_i < 0$ and $\tilde{z}_i < 0$, high–low (HL) if $z_i > 0$ and $\tilde{z}_i < 0$, and low–high (LH) if $z_i < 0$ and $\tilde{z}_i > 0$. HH areas correspond to hot spots and LL areas to cold spots, while HL and LH areas are spatial outliers. Statistical significance is assessed with conditional random permutations, in which the values of the neighbors of area i are repeatedly reshuffled to obtain a reference distribution of I_i.

Local Moran's I is computed for districts using the urban liveliness derived from Altieri entropy (L0 categories, filter level 1, base 5). Figure 6 shows the local statistics divided into five quintiles, the quadrant classification of each district, and the significant districts. After masking non-significant values ($p < 0.05$), 36.83% of districts remain statistically significant.

What stands out is that the large Dutch cities emerge as hot spots of liveliness, while a substantial part of the north of the Netherlands forms a contiguous cold spot. It implies that, even though the urban landscape of the Netherlands might seem uniform, liveliness is spatially concentrated. Furthermore, hot spots are not confined to the districts of the city centers but extend into the surrounding districts, suggesting that the advantages of large cities in amenity complexity spill over municipal boundaries. Cold spots, in turn, form large contiguous regions, which indicates that low liveliness is a regional rather than a local condition and that residents of these regions have few nearby alternatives with a more complex amenity landscape.

3.5 Area similarity and dashboard

To gain insight into which areas resemble one another, each area a is described by a feature vector

$$\mathbf{x}_a = \left(L_a^{\mathrm{L0}},\ L_a^{\mathrm{L1}},\ I_a^{\mathrm{L0}},\ I_a^{\mathrm{L1}}\right) \tag{9}$$

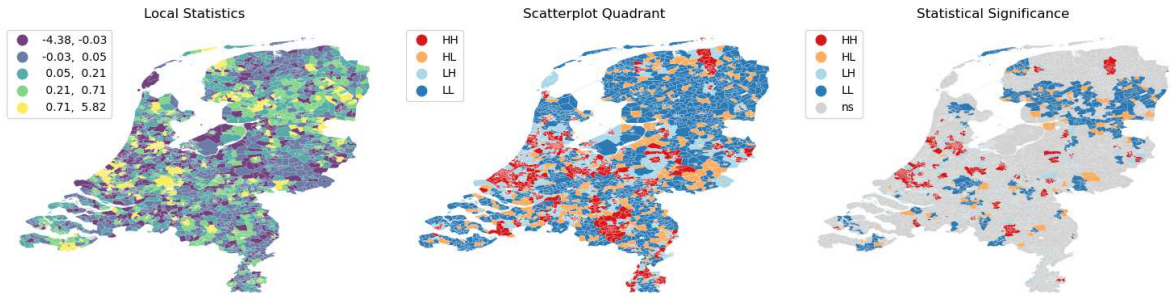


Figure 6: **Local spatial autocorrelation of urban liveliness across Dutch districts**. Local Moran's I divided into five quintiles (left), Moran scatterplot quadrants HH, LL, LH, and HL (middle), and quadrants of significant districts only, covering 36.83% of districts (right). Urban liveliness is derived from Altieri entropy (L0 categories, filter level 1, base 5). Large cities form hot spots, while much of the north forms a cold spot.

where L_a^{L0} and L_a^{L1} denote the urban liveliness derived from Altieri entropy for L0 and L1 categories, and I_a^{L0} and I_a^{L1} the corresponding local Moran's I values. The dissimilarity between two areas a and a' at the same scale is then given by the Euclidean distance:

$$D(a, a') = \|\mathbf{x}_a - \mathbf{x}_{a'}\|_2 = \sqrt{\sum_{f=1}^{4} \left(x_{a,f} - x_{a',f}\right)^2} \tag{10}$$

where f indexes the features. These features are chosen because they jointly contain information about the spatial distribution of amenities, the urbanity of an area, and its autocorrelation with the surrounding areas. Two areas are thus considered similar not only when their amenity landscapes are equally complex but also when they are embedded in similar surroundings. For any selected area, the five areas with the smallest distance are reported as its most similar counterparts.

The results are made accessible through an interactive dashboard, namely **"Urban Liveliness Observatory"**, built using the Dash framework (Plotly Technologies Inc., 2015) (Figure 7). The main view features a map that users can pan and zoom at the municipality, district, or neighborhood scale. Controls beneath the map set the entropy measure, the category level, the filter level, and whether normalized or transformed values are shown. When an area is selected, a side panel appears with a map of its amenities, a bar plot of amenities per L0 category, and a bar plot of the sociodemographic composition of its inhabitants in terms of migration background, age, and gender (Supplementary Figure S8). Lastly, a table lists the five areas most similar to the selected one according to Eq. (10). In this way, the dashboard translates the entropy-based measures into a form that can support residential search and local planning, and it allows users to inspect how the choice of measure and filter changes the picture of an area.

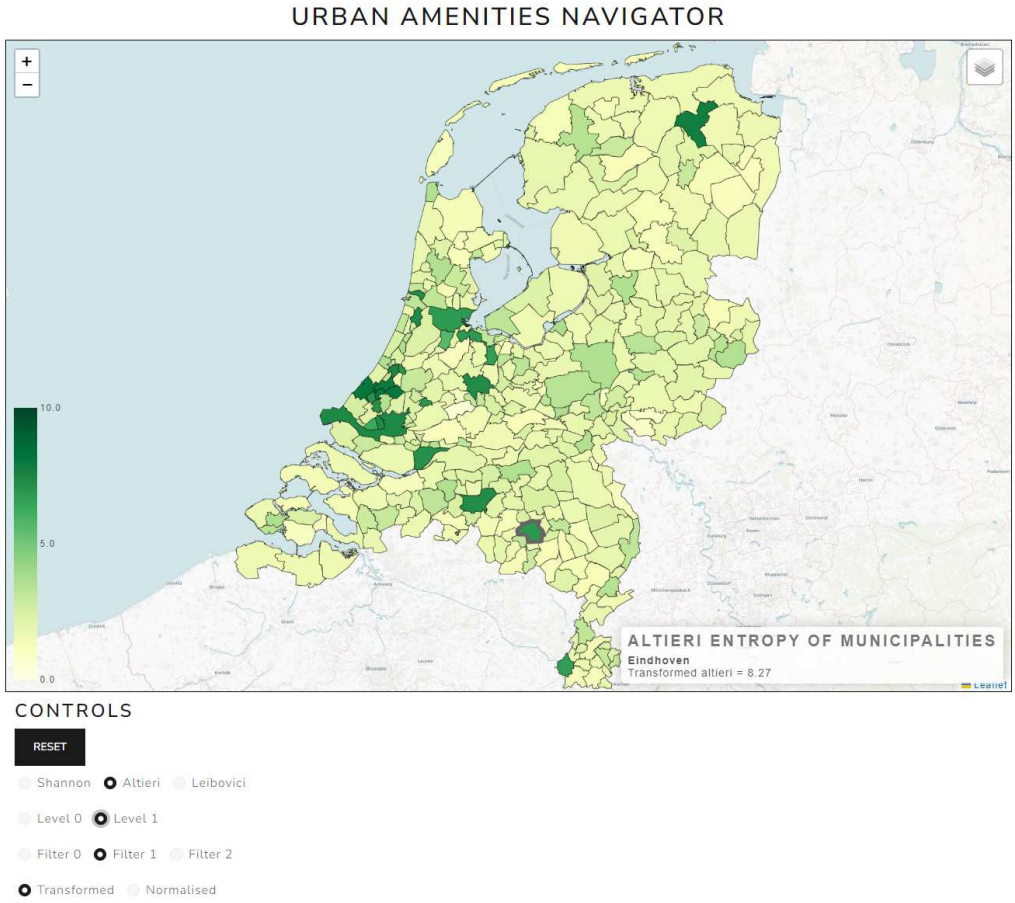


Figure 7: **Main view of the dashboard**. Users select the spatial scale (top right) and, below the map, the entropy measure, category level, filter level, and normalized or transformed values. Selecting an area opens an inspection panel with its amenities, sociodemographic composition, and the five most similar areas (Supplementary Figure S8).

4 Discussion and conclusions

This study aims to evaluate whether information-theoretic measures can capture the composition and configuration of urban amenities and whether they can be translated into a comparable measure of urban liveliness. We apply Shannon (Shannon, 1948), Leibovici (Leibovici et al., 2014), and Altieri entropy (Altieri et al., 2019a) to all municipalities, districts, and neighborhoods of the Netherlands, transform the resulting values into urban liveliness (Jacobs, 1961; Yue et al., 2021; Wu and Niu, 2019), and use them to compute spatial autocorrelation (Anselin, 1995). Through these methods, several insights are obtained, ranging from the spatial context of living space to similarity profiles between areas.

4.1 Composition, configuration, and urban complexity

Regarding the first question, the measures capture different facets of complexity. Shannon entropy reflects composition only, Leibovici entropy reflects configuration when clustering follows categories, and Altieri entropy responds to configuration regardless of how categories are clustered (Section 3.2.1). In the empirical data, the measures agree closely within a filter level, while the choice of filter changes the results considerably, especially at the municipal scale (Section 3.2.2). High entropy values are typically associated with areas that not only host many amenities but also distribute them well over multiple categories (Section 3.2.3).

Regarding the second question, weighting spatial entropy by urbanity with a logarithmic function proves to be an effective way of capturing urban liveliness and results in a refined delineation between urban and rural characteristics (Section 3.3). Regarding the third question, local Moran's I reveals hot spots of liveliness in the large cities and cold spots across much of the north (Section 3.4). Even in a country whose urban landscape might seem uniform, liveliness is spatially concentrated. Combining liveliness with local autocorrelation further allows areas to be matched to comparable counterparts, which the dashboard makes available to a wider audience (Section 3.5). Taken together, urban complexity emerges neither from the number of amenities nor from their diversity alone, but from the joint structure of what is present and where it is located.

These findings extend and qualify previous work in several ways. Studies of land use mix have relied predominantly on Shannon-type indices (Song et al., 2013). Our synthetic experiment shows that such indices are, by construction, blind to whether functions are mixed within an area or segregated into separate parts of it, a distinction that matters for walkability and for proximity-based planning concepts such as the 15-minute city (Moreno et al., 2021). At the same time, the strong empirical correlation between measures within a filter level qualifies this argument: in the Dutch amenity landscape, composition and configuration largely co-vary, so that Shannon entropy remains a reasonable first approximation when ranking areas. Previous applications of spatial entropy have focused on land cover and urban expansion (Cabral et al., 2013; Altieri et al., 2019b; Altieri and Cocchi, 2022). Applying these measures to point patterns of everyday amenities at national scale shows that they are also informative about the functional structure of residential areas. Finally, studies of urban vitality typically observe human activity through mobile phone, social media, or street-view data (Jin et al., 2017; Wu and Niu, 2019; Tu et al., 2020; Li et al., 2022; Ouyang et al., 2022). Our results suggest that a structural proxy derived exclusively from open data reproduces the expected contrast between large cities and rural regions, which makes it a transparent complement to activity-based measures where such data are unavailable.

4.2 Implications

The findings carry implications for different audiences. First, for residents, especially those unfamiliar with an area, entropy-based liveliness offers a concise and comparable summary of the amenity landscape that existing housing platforms do not provide. The similarity function further allows users to discover alternative areas with a comparable profile, which can broaden the search beyond familiar neighborhoods. Second, for planners and policymakers, maps of hot and cold spots point to regions where the everyday urban environment is less complex. Because cold spots form large contiguous regions rather than isolated districts, improving access to amenities in these regions may require coordination across municipal boundaries rather than interventions in single neighborhoods. Third, on the methodological side, the sensitivity of the results to the filter level shows that decisions about which amenities to include deserve as much attention as the choice of entropy measure. Researchers using OSM-based indicators should therefore report and test the sensitivity of their category selection, rather than treating all mapped features as equally relevant.

4.3 Limitations

Despite these results, there is no guarantee that the underlying data are complete. OSM data are predominantly collected by volunteer contributors, resulting in varying levels of documentation across Dutch cities and between urban and rural areas in general (Haklay, 2010; Barrington-Leigh and Millard-Ball, 2017). Part of the urban–rural contrast may therefore reflect differences in mapping effort rather than in the amenity landscape itself. Applying the same methods to a more comprehensive OSM dataset in the future, validating OSM against official registers of businesses and facilities, or classifying remote sensing imagery with deep learning could address this issue. In addition, amenities are collected within the bounding box of each spatial unit, so irregularly shaped units may include amenities located just outside their boundaries.

Apart from that, the urbanity weighting relies on a five-class categorical variable, and the logarithmic base is chosen by inspecting the resulting distributions rather than by optimizing an external criterion. The measure of urban liveliness should therefore be read as a structural proxy, and its validation against observed activity, such as footfall or mobility data, remains an open task. Local Moran's I is evaluated for each district separately, and the resulting significance levels are not corrected for multiple testing; applying a false discovery rate correction (Benjamini and Hochberg, 1995) would yield a more conservative set of hot and cold spots. Furthermore, the dashboard is restricted to three predetermined filter levels. Customizable filters would better meet the needs of individual users but are too computationally expensive at present; computing entropy on a remote server could make this feasible. Finally, the analysis is cross-sectional and descriptive, so it does not establish how amenity structure relates causally to residential choice or well-being.

4.4 Future research and concluding remarks

Future research could incorporate a temporal dimension by including OSM data from different years. This could shed light on how liveliness and its spatial clustering evolve, whether development policies leave a visible imprint on the urban landscape, and whether broader areas agglomerate or separate as cities grow. Another direction is to model the sociodemographic characteristics of an area from its amenity distribution, or vice versa. Such a model could reveal how

the built environment and its population co-evolve and guide policymakers in designing urban environments that enhance social cohesion and urban dynamics. Since the pipeline relies only on OSM and on official boundaries, it can also be transferred to other countries, which would allow testing whether the patterns observed in the Netherlands hold in urban systems with a different density and settlement structure.

In conclusion, information-theoretic measures provide a compact and interpretable description of the amenity landscape of cities. Distinguishing composition from configuration clarifies what each measure captures, weighting by urbanity turns complexity into a comparable measure of liveliness, and spatial autocorrelation reveals where liveliness is concentrated and where it is lacking. Built entirely on open data, this framework offers residents, planners, and researchers a transparent way to compare living spaces across scales.

Abbreviations

CBS: Statistics Netherlands (*Centraal Bureau voor de Statistiek*); HH: high–high; HL: high–low; LH: low–high; LISA: local indicators of spatial association; LL: low–low; OSM: OpenStreetMap; POI: point of interest.

Disclosure statement

No potential conflict of interest was reported by the authors.

References

NM Allen. *Quality of urban life and intensification: Understanding housing choices, trade-offs, and the role of urban amenities*. PhD thesis, ResearchSpace@ Auckland, 2016.

Linda Altieri and Daniela Cocchi. Understanding the expansion of italian metropolitan areas: A study based on entropy measures. *Environment and Planning B: Urban Analytics and City Science*, 49(2):447–463, 2022.

Linda Altieri, Daniela Cocchi, and Giulia Roli. Advances in spatial entropy measures. *Stochastic environmental research and risk assessment*, 33:1223–1240, 2019a.

Linda Altieri, Daniela Cocchi, and Giulia Roli. Measuring heterogeneity in urban expansion via spatial entropy. *Environmetrics*, 30(2):e2548, 2019b.

Linda Altieri, Daniela Cocchi, and Giulia Roli. Efficient computation of spatial entropy measures. *Entropy*, 25(12):1634, 2023.

Luc Anselin. Local indicators of spatial association—LISA. *Geographical analysis*, 27(2):93–115, 1995.

Luc Anselin. A local indicator of multivariate spatial association: extending geary's c. *Geographical Analysis*, 51(2):133–150, 2019.

Christopher Barrington-Leigh and Adam Millard-Ball. The world's user-generated road map is more than 80% complete. *PLoS ONE*, 12(8):e0180698, 2017.

Michael Batty. Spatial entropy. *Geographical analysis*, 6(1):1–31, 1974.

Michael Batty. *The New Science of Cities*. MIT Press, Cambridge, MA, 2013.

Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300, 1995.

Luís M. A. Bettencourt. The origins of scaling in cities. *Science*, 340(6139):1438–1441, 2013.

Pedro Cabral, Gabriela Augusto, Mussie Tewolde, and Yikalo Araya. Entropy in urban systems. *Entropy*, 15(12):5223–5236, 2013.

Centraal Bureau voor de Statistiek. Wijk- en buurtkaart 2023. `https://www.cbs.nl/nl-nl/dossier/nederland-regionaal/geografische-data/wijk-en-buurtkaart-2023`, 11 2023. URL `https://www.cbs.nl/nl-nl/dossier/nederland-regionaal/geografische-data/wijk-en-buurtkaart-2023`.

Arthur Getis and Jared Aldstadt. Constructing the spatial weights matrix using a local statistic. *Geographical analysis*, 36(2):90–104, 2004.

Edward L. Glaeser, Jed Kolko, and Albert Saiz. Consumer city. *Journal of Economic Geography*, 1(1):27–50, 2001.

Mordechai Haklay. How good is volunteered geographical information? A comparative study of OpenStreetMap and Ordnance Survey datasets. *Environment and Planning B: Planning and Design*, 37(4):682–703, 2010.

Rafiazka Millanida Hilman. Stratified mobility, segregated boundary, and socioeconomic mixing in New York. *Computational Urban Science*, 5(1):63, 2025.

Jane Jacobs. The death and life of american cities, the failure of town planning, 1961.

Xiaobin Jin, Ying Long, Wei Sun, Yuying Lu, Xuhong Yang, and Jingxian Tang. Evaluating cities' vitality and identifying ghost cities in china with emerging geographical data. *Cities*, 63:98–109, 2017.

K Jordahl. Geopandas: Python tools for geographic data. `https://github.com/geopandas/geopandas`, 2014.

Anders Karlström and Vania Ceccato. A new information theoretical measure of global and local spatial association. *The Review of Regional Research*, 22(1):13–40, 2002.

Andrew H Kelly. Securing urban amenity: does it coincide with biodiversity conservation at the local government level? *Australasian Journal of Environmental Management*, 13(4):243–253, 2006.

Didier G Leibovici. Defining spatial entropy from multivariate distributions of co-occurrences. In *Spatial Information Theory: 9th International Conference, COSIT 2009 Aber Wrac'h, France, September 21-25, 2009 Proceedings 9*, pages 392–404. Springer, 2009.

Didier G Leibovici, Christophe Claramunt, Damien Le Guyader, and David Brosset. Local and global spatio-temporal entropy indices based on distance-ratios and co-occurrences distributions. *International Journal of Geographical Information Science*, 28(5):1061–1084, 2014.

Yunqin Li, Nobuyoshi Yabuki, and Tomohiro Fukuda. Exploring the association between street built environment and street vitality using deep learning methods. *Sustainable Cities and Society*, 79:103656, 2022.

Kevin Lynch. *Good city form*. MIT press, 1984.

Rod Mccrea and Peter Walters. Impacts of urban consolidation on urban liveability: Comparing an inner and outer suburb in brisbane, australia. *Housing, theory and society*, 29(2):190–206, 2012.

Rod McCrea, Robert Stimson, and Robert W Marans. The evolution of integrative approaches to the analysis of quality of urban life. In *Investigating quality of urban life: theory, methods, and empirical research*, pages 77–104. Springer, 2011.

John Montgomery. Making a city: Urbanity, vitality and urban design. *Journal of urban design*, 3(1):93–116, 1998.

Carlos Moreno, Zaheer Allam, Didier Chabaud, Catherine Gall, and Florent Pratlong. Introducing the "15-minute city": Sustainability, resilience and place identity in future post-pandemic cities. *Smart Cities*, 4(1):93–111, 2021.

Stan Openshaw. *The Modifiable Areal Unit Problem*. Number 38 in Concepts and Techniques in Modern Geography. Geo Books, Norwich, 1984.

OpenStreetMap contributors. Planet dump retrieved from https://planet.osm.org. `https://www.openstreetmap.org`, 2017.

Jiani Ouyang, Hong Fan, Luyao Wang, Dongyu Zhu, and Mei Yang. Revealing urban vibrancy stability based on human activity time-series. *Sustainable Cities and Society*, 85:104053, 2022.

Plotly Technologies Inc. Collaborative data science. `https://plot.ly`, 2015.

Todd A Randall. Preferences of suburban residents in thunder bay, ontario towards neighbourhood intensification and rediversification. *Canadian Journal of Urban Research*, 17(2):28–56, 2008.

Sergio J. Rey and Luc Anselin. PySAL: A Python Library of Spatial Analytical Methods. *The Review of Regional Studies*, 37(1):5–27, 2007.

Sergio J. Rey, Dani Arribas-Bel, and Levi John Wolf. *Geographic data science with python*. Chapman and Hall/CRC, 2023.

Paulo Saraiva. On shannon entropy and its applications. *Kuwait Journal of Science*, 50(3): 194–199, 2023.

K Saville-Smith. Neighbourhoods and intensification: Measuring sustainability impacts of higher density and mixed use. *SB10 New Zealand*, 2010.

Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.

Yan Song, Louis Merlin, and Daniel Rodriguez. Comparing measures of urban land use mix. *Computers, Environment and Urban Systems*, 42:1–13, 2013.

Wei Tu, Tingting Zhu, Jizhe Xia, Yulun Zhou, Yani Lai, Jincheng Jiang, and Qingquan Li. Portraying the spatial dynamics of urban vibrancy using multisource urban big data. *Computers, Environment and Urban Systems*, 80:101428, 2020.

Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3):261–272, 2020.

Wanshu Wu and Xinyi Niu. Influence of built environment on urban vitality: Case study of shanghai using mobile phone location data. *Journal of Urban Planning and Development*, 145 (3):04019007, 2019.

Jing Yuan, Yu Zheng, and Xing Xie. Discovering regions of different functions in a city using human mobility and POIs. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 186–194, 2012.

Wenze Yue, Yang Chen, Pham Thi Mai Thy, Peilei Fan, Yong Liu, and Wei Zhang. Identifying urban vitality in metropolitan areas of developing countries from a comparative perspective: Ho chi minh city versus shanghai. *Sustainable Cities and Society*, 65:102609, 2021.

Y Zheng. Spatialentropy. `https://github.com/Mr-Milk/SpatialEntropy`, 2021.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- CUSAmenitiesSM.zip