

Municipal environmental exposure profiles for climate adaptation in the Netherlands

Joshua Osejigbe Aguele

joshuaaguele.research@gmail.com

Independent Researcher

Research Article

Keywords: environmental indicators, multivariate profiles, principal component analysis, k-means clustering, parallel analysis, composite indicators, Netherlands

Posted Date: October 14th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11245737/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Municipal adaptation planning requires information about combinations of environmental pressures, but separate indicators and composite rankings provide different and incomplete summaries of those combinations. This study examines which recurring environmental profiles distinguish Dutch municipalities and whether their membership persists under resampling, drought preprocessing and alternative flood aggregation denominators. Five population-weighted municipal indicators from the Klimaateffectatlas Buurtdashboard were analysed: canopy cover, drought depth, heat exposure, flood exposure and heavy-rain exposure. Of 342 municipalities, 337 had complete data. Principal component analysis (PCA), permutation-based parallel analysis and k-means clustering were applied to standardized indicators. Two components exceeded the 95th-percentile eigenvalues from 1,000 permuted datasets and explained 70.53% of variance. A four-cluster solution accounted for 57.86% of total standardized variation, with an average silhouette width of 0.328. It distinguished high-canopy/low-heat municipalities ($n = 52$), low-flood/rain municipalities ($n = 133$), a flood-exposure group ($n = 136$), and a drought-heavy-rain extreme group ($n = 16$). Bootstrap mean Jaccard similarities ranged from 0.816 to 0.978. Agreement with the original partition remained high after log-transforming drought depth (adjusted Rand index = 0.932) and excluding its highest value (0.997). Flood sensitivity comparisons on matched municipal samples yielded ARIs of 0.912–1.000 across thresholds for the population-weighted share of neighbourhoods with recorded inundation, despite substantial changes in some flood estimates. These profiles support further local assessment. They complement individual indicators; they do not establish causal effects, comprehensive vulnerability or the superiority of clustering over a composite index.

1 Introduction

Municipal adaptation planning must address different combinations of heat, drought, heavy rainfall and flood conditions. The substantive problem is how to identify places with similar environmental configurations so that further assessment can focus on the pressures most relevant locally. A single ranking offers limited information about which combination underlies a municipality's position. Composite indices also introduce choices about weighting, aggregation and robustness: assigned weights need not equal an indicator's effective influence, and compensatory aggregation allows favourable conditions in one domain to offset unfavourable conditions in another (OECD and European Union 2008; Becker et al. 2017; Greco et al. 2019). A profile approach can help retain these distinctions, although its usefulness must be assessed rather than assumed.

Indicator-based climate assessments commonly distinguish exposure, sensitivity and adaptive capacity, yet the boundaries among these concepts must be explicit if indicators are to support defensible decisions (Birkmann 2007; Cutter et al. 2003; Hinkel 2011; IPCC 2022). The Dutch Climate Impact Atlas Buurtdashboard provides comparable municipal environmental indicators, including heat, drought, flooding, heavy rain and canopy. These variables characterize environmental exposure and, in the case of canopy, a potentially protective environmental condition; they do not directly measure household

sensitivity, health, income, housing quality or institutional capacity. The present analysis therefore constructs environmental exposure profiles rather than comprehensive vulnerability classes.

European research provides a substantive basis for examining environmental configurations across places. Tapia et al. (2017) assessed vulnerability across 571 European cities and five climate-threat impact chains, showing how indicator frameworks can support comparison across multiple threats. The European Climate Risk Assessment emphasizes risks across sectors and the need for coordinated adaptation (European Environment Agency 2024). Heat adaptation also depends on environmental context: an analysis of 293 European cities found that urban trees were associated with lower land surface temperatures, with effectiveness varying across climatic regions (Schwaab et al. 2021). Together, these studies motivate attention to hazard-specific conditions and protective environmental features. Here, the empirical question is whether five indicators from one Dutch municipal dashboard reveal interpretable configurations that remain recognizable under resampling, drought sensitivity checks and alternative flood aggregation.

Profile construction complements separate indicator lists and one-number rankings by describing recurring combinations. Principal component analysis (PCA) summarizes correlated variables through orthogonal dimensions (Jolliffe and Cadima 2016). Cluster analysis groups municipalities with similar indicator configurations. The two methods answer related but distinct questions: PCA describes continuous gradients, while clustering imposes a discrete partition. K-means still aggregates squared distances across indicators and is not a fully non-compensatory decision rule. Component retention and cluster membership therefore require explicit assessment, and stability alone does not establish external validity (Hennig 2007; Horn 1965; Rousseeuw 1987).

The contribution is an applied assessment of municipal environmental configurations, rather than a new statistical algorithm. The workflow combines a common municipal dataset with permutation-based component retention, explicit comparison of alternative cluster counts, bootstrap stability, drought and flood sensitivity checks and interpretation in original units. This makes the evidence for the selected typology inspectable and helps distinguish a recurring profile from a partition dependent on one extreme observation, one transformation or a change in the flood aggregation denominator.

The main research question is: which recurring combinations of environmental exposure and canopy cover distinguish Dutch municipalities for adaptation screening, and how robust are these combinations to resampling, drought preprocessing and alternative flood aggregation denominators? Two analytical questions support it: (1) which dimensions summarize covariance among the five indicators, and (2) which municipal profiles emerge from their joint distribution? The resulting typology is intended to guide further investigation of local priorities. It does not measure adaptation outcomes or replace neighbourhood evidence, social-sensitivity measures and local judgement.

2 Materials and Methods

2.1 Data source and indicators

Data were obtained from the Klimaateffectatlas Buurtdashboard and associated metadata workbook (Klimaateffectatlas 2026). Neighbourhood indicators were aggregated to municipalities using neighbourhood population as weights, retaining available indicator values with positive, non-missing population weights. Thus, the reported percentages are population-weighted averages of neighbourhood environmental percentages, not direct estimates of the percentage of municipal residents affected. The derived dataset contained 342 municipalities with no duplicate municipality codes. Five municipalities lacked drought-depth observations; PCA and clustering used the 337 complete cases (98.54%). The other four derived indicators were complete. Descriptive statistics are reported in Table 1.

Canopy denotes tree-canopy coverage. Drought depth uses the source field lg3max_huidig, describing extreme low groundwater in the current climate, in metres; it should not be interpreted as drought duration. Heat combines source apparent-temperature classes of at least 41°C. Flood combines maximum inundation-depth classes greater than 50 cm under the source's extremely low-probability flood scenario. Heavy-rain exposure combines surface-water-depth classes of at least 20 cm according to the metadata, despite legacy source field names containing “mm”. These indicators describe different mapped conditions rather than simultaneous observed events. The source provider confirmed that blank entries in the four coastal and fluvial flood-depth categories indicate no flooding in the corresponding category, and that each percentage uses total neighbourhood area as its denominator (A. Koekoek, Stichting CAS, personal communication, 2 October 2026). Blank category entries were therefore coded as zero before summing, including records with all four categories blank. The resulting municipal flood indicator is a population-weighted average of neighbourhood area percentages; it is not the percentage of residents exposed or an area-weighted percentage of total municipal land.

Table 1
Descriptive statistics for the environmental indicators

Indicator	N	Mean	SD	Minimum	Median	Maximum
Drought depth (m)	337	3.95	6.71	0.90	2.17	59.80
Canopy cover (%)	342	13.31	4.76	4.83	12.68	34.90
Heat exposure ≥ 41°C (%)	342	65.62	6.11	45.15	66.05	83.57
Flood exposure over 50 cm (%)	342	40.06	35.91	0.00	35.11	98.89
Heavy-rain water depth ≥ 20 cm (%)	342	2.95	1.44	0.48	2.63	9.77
Note. Descriptive statistics use all available values. PCA and clustering use 337 complete cases.						

The indicators differ in units and dispersion. Standardization was therefore required before PCA and clustering. Indicator direction was not reversed because the aim was to describe empirical configurations rather than construct a normative risk score. Canopy was retained as an environmental feature that may contribute to heat protection; it is not a measure of institutional adaptive capacity.

2.2 Standardization and principal component analysis

For municipality i and indicator j , each complete-case observation was standardized to zero mean and unit sample variance:

$$z_{ij} = \frac{x_{ij} - \bar{x}_j}{s_j}$$

1

Here x_{ij} is the observed indicator, $\bar{x}_j$ is its complete-case sample mean, and s_j is its sample standard deviation. PCA was applied to the correlation structure. The score on component k for municipality i was:

$$PC_{ik} = \sum_{j=1}^5 a_{jk} z_{ij} \quad (2)$$

In Eq. (2), a_{jk} is the eigenvector coefficient for indicator j on component k , and z_{ij} is the standardized value for municipality i . Score coefficients differ from the variable–component correlation loadings reported in Table 4; those loadings equal a_{jk} multiplied by the square root of the component eigenvalue. Component retention used parallel analysis: observed eigenvalues were compared with the 95th percentile of the ordered eigenvalues from 1,000 datasets formed by independently permuting each indicator without replacement, preserving marginal distributions while disrupting cross-indicator associations (Horn 1965). A component was retained when its observed eigenvalue exceeded that benchmark. Component signs are arbitrary.

2.3 Cluster analysis, selection and stability

K-means clustering was applied directly to the standardized five-indicator matrix, not to the retained PCA scores. Candidate solutions from two to ten clusters were examined using within-cluster sum of squares and average silhouette width. The final two- and four-cluster comparisons used seed 2026, 100 random starts and a maximum of 100 iterations. Multiple starts reduce dependence on initialization (Lloyd 1982). Silhouette width compares within-cluster cohesion with separation from the nearest competing cluster (Rousseeuw 1987):

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}$$

3

In Eq. (3), $a(i)$ is the mean distance from observation i to other members of its cluster, and $b(i)$ is the minimum mean distance to another cluster. Higher silhouette values indicate clearer separation. The

final cluster count balanced parsimony, average silhouette width, between-cluster sum of squares as a fraction of total standardized sum of squares, and substantive distinctiveness. This fraction describes partition fit; it is not the PCA variance percentage. Two- and four-cluster solutions were reported explicitly. Names were assigned after examining original-unit means and medians and describe relative profiles rather than absolute safety or vulnerability.

Cluster stability was evaluated with 500 bootstrap samples. Each original cluster was compared with its closest resampled match, and mean Jaccard similarity summarized membership overlap (Hennig 2007). Values of at least 0.75 were interpreted as stable and values of at least 0.85 as highly stable. Dissolution and recovery frequencies were also recorded. These are internal resampling assessments rather than validation against outcomes or an independent population. The original analysis and sensitivity checks were conducted in R (R Core Team 2026).

2.4 Drought sensitivity checks

Drought depth was strongly right-skewed, so two sensitivity analyses retained four clusters while changing this aspect of the input. First, natural-log drought depth replaced its original values for all 337 complete cases; all five indicators were then standardized again. All drought values were positive, so no offset was needed. Second, Landgraaf (GM0882), the municipality with the highest drought value (59.80 m), was excluded, and all indicators were re-standardized across the remaining 336 municipalities. PCA and k-means were re-estimated with the same four-cluster settings. Agreement with the original assignments was measured using the adjusted Rand index (ARI), which corrects pairwise partition agreement for chance and is invariant to cluster numbering (Hubert and Arabie 1985). The exclusion comparison used the original labels for the same remaining 336 municipalities. The first two PCA variance percentages were reported descriptively; parallel analysis was not repeated for these altered datasets.

2.5 Flood aggregation denominator sensitivity checks

The source workbook contained 14,574 neighbourhoods, of which 4,240 had blank entries in all four flood-depth classes. Among neighbourhoods with positive finite population weights, these records represented 28.11% of population. Following the source-provider clarification, these blanks represent absent coastal or fluvial inundation in the mapped depth categories, rather than missing flood observations. The baseline therefore averages the combined greater-than-50-cm indicator across all eligible neighbourhoods, retaining zero values. A previously computed alternative omits neighbourhoods with all four categories blank from the numerator and denominator. It is retained solely to illustrate sensitivity to conditioning the average on neighbourhoods with recorded inundation in at least one category. Partially blank categories are coded as zero in both versions. The alternative changes the target summary and is not a missing-data correction or an equally supported estimate of the baseline quantity.

The population-weighted share of neighbourhoods with recorded inundation was defined as the share of eligible population weights assigned to neighbourhoods with at least one nonblank flood-depth category.

This share does not measure data completeness, the percentage of residents flooded, or the percentage of neighbourhood area inundated. Exploratory minimum-share thresholds of 0%, 80%, 90% and 100% were applied within the original 337 complete cases, additionally requiring a finite conditional flood estimate. At each threshold, both versions used identical municipalities; only the flood indicator changed. Each five-indicator matrix was standardized separately, and PCA and four-cluster k-means were re-estimated using seed 2026, 100 starts and 100 maximum iterations. ARI compared the paired partitions. First-two-component variance percentages were descriptive; parallel analysis was not repeated. At a 100% share, equality of the input matrices was checked within a tolerance of 10^{-6} . These thresholds describe the denominator sensitivity exercise and do not justify excluding municipalities from the baseline analysis.

ChatGPT (OpenAI) assisted with R code development and troubleshooting and manuscript preparation. The author reviewed the code, checked the outputs, and takes responsibility for the analyses and manuscript; further details are provided in Statements and Declarations.

3 Results

3.1 Correlation structure

The indicators were related but not interchangeable. The strongest positive correlation linked drought depth and heavy-rain exposure ($r = 0.602$); the strongest negative correlation linked canopy and heat exposure ($r = -0.578$). Canopy was also negatively correlated with flood exposure ($r = -0.405$). Heat exposure was nearly unrelated to flood exposure ($r = -0.037$) and only weakly related to heavy-rain exposure ($r = 0.074$). These patterns indicate contrasting associations and motivate retaining domain-specific information. The full correlation matrix is shown in Table 2.

Table 2
Pearson correlations among standardized indicators

Indicator	Canopy	Drought	Flood	Heat	Heavy rain
Canopy cover	1.000	0.145	-0.405	-0.578	0.268
Drought depth	0.145	1.000	-0.269	0.125	0.602
Flood exposure	-0.405	-0.269	1.000	-0.037	-0.266
Heat exposure	-0.578	0.125	-0.037	1.000	0.074
Heavy-rain exposure	0.268	0.602	-0.266	0.074	1.000
Note. N = 337 complete cases. Coefficients are descriptive Pearson correlations.					

3.2 Principal components

Parallel analysis retained two components. PC1 had an eigenvalue of 2.009 and explained 40.18% of variance; PC2 had an eigenvalue of 1.517 and explained 30.35%. Their cumulative explained variance

was 70.53%. PC3 explained a further 16.38%, but its eigenvalue (0.819) did not exceed the random 95th-percentile benchmark (1.036). Eigenvalues and retention decisions appear in Table 3 and Fig. 1.

Table 3
PCA eigenvalues and retention decision

Component	Eigenvalue	Variance (%)	Cumulative (%)	Random 95th	Retained
PC1	2.009	40.18	40.18	1.221	Yes
PC2	1.517	30.35	70.53	1.113	Yes
PC3	0.819	16.38	86.91	1.036	No
PC4	0.404	8.08	94.99	0.975	No
PC5	0.250	5.01	100.00	0.909	No
Note. Random benchmarks are 95th-percentile eigenvalues from 1,000 permuted datasets.					

PC1 loaded positively on heavy-rain exposure (0.753), drought depth (0.694) and canopy (0.694), and negatively on flood exposure (-0.653). It therefore captured a drought/heavy-rain/canopy versus flood contrast. PC2 loaded positively on heat exposure (0.868) and negatively on canopy (-0.623), with smaller positive contributions from drought and heavy rain. It represented a heat-versus-canopy gradient. The axes describe covariance patterns; they do not establish mechanisms or value judgements. Loadings are reported in Table 4 and Fig. 2.

Table 4
Loadings on the retained principal components

Indicator	PC1	PC2
Drought depth	0.694	0.484
Canopy cover	0.694	-0.623
Heat exposure	-0.232	0.868
Flood exposure	-0.653	-0.020
Heavy-rain exposure	0.753	0.377
Note. Entries are correlations between standardized indicators and component scores, not the eigenvector coefficients used in Eq. (2). Values are rounded to three decimals; component signs are arbitrary.		

3.3 Selection of the municipal typology

The two-cluster solution had the highest average silhouette width (0.340), but its between-cluster sum of squares accounted for only 25.88% of total standardized variation. The four-cluster solution had a slightly lower silhouette width (0.328) and accounted for 57.86%. It distinguished a flood-exposure

profile and a small drought-heavy-rain extreme profile. Four clusters were retained to provide greater substantive resolution, rather than because they maximized separation. The modest average silhouette indicates appreciable overlap and cautions against treating the groups as sharply separated natural classes. Table 5 reports within-cluster sums of squares for the two solutions, making the reduction in unexplained variation explicit. Table 5 compares the two- and four-cluster summaries.

Table 5 Comparison of candidate summary solutions				
Solution	Average silhouette	Variance explained (%)	Cluster sizes	Within-cluster SS
Two clusters	0.340	25.88	85, 252	1245.18
Four clusters	0.328	57.86	52, 133, 136, 16	707.99
Note. Candidate k values from 2 to 10 were reviewed; the table reports the two-cluster and selected four-cluster solutions. Total standardized sum of squares is 1,680.00. The between-cluster fraction equals $100 \times (\text{total SS} - \text{within-cluster SS}) / \text{total SS}$; it measures partition fit rather than predictive performance.				

3.4 Four municipal environmental profiles

Cluster 1 contained 52 municipalities (15.43%) and combined the highest canopy mean (20.81%) with the lowest heat mean (57.78%); it was labelled high canopy, low heat. Cluster 2 contained 133 municipalities (39.47%) and had low flood exposure (11.28%) and below-average heavy-rain exposure (2.77%), although heat exposure was relatively high (68.15%). Cluster 3 was the largest group, containing 136 municipalities (40.36%), and was dominated by flood exposure (78.96%).

Cluster 4 contained 16 municipalities (4.75%) and had mean drought depth of 28.99 m, mean heavy-rain exposure of 6.59%, and mean heat exposure of 70.78%. Drought and heavy-rain medians were also high (26.24 m and 6.35%). These summaries suggest a group-level pattern; the exclusion sensitivity check below directly assesses dependence on the highest-drought municipality. Its flood mean was the lowest of the four groups (8.62%). Table 6 and Fig. 3 summarize the four profiles.

Table 6
Original-unit means for the four municipal profiles

Profile	N	Drought (m)	Canopy (%)	Heat (%)	Flood (%)	Heavy rain (%)
1 High canopy, low heat	52	5.30	20.81	57.78	19.18	3.83
2 Low flood/rain	133	2.54	12.94	68.15	11.28	2.77
3 Flood exposure	136	1.86	10.76	65.24	78.96	2.31
4 Drought-heavy-rain extreme	16	28.99	13.34	70.78	8.62	6.59
Note. Values are cluster means. Labels summarize relative configurations and do not imply total vulnerability. Standardized indicator centroids are reported in Appendix A. Table 10 in Appendix A reports the standardized centroids.						

3.5 Geographic distribution

Figure 4 describes the spatial distribution of profile membership. Flood-exposure municipalities occupy broad contiguous areas, while the drought-heavy-rain extreme group occurs in fewer locations. This is a descriptive map; no test of spatial randomness or geographic clustering was performed. Shared landscapes and hydrological regimes are plausible influences on these configurations, but the map does not establish those mechanisms. Geographic context should therefore inform local review rather than be treated as independent validation of the typology.

3.6 Bootstrap stability

All four profiles exceeded the 0.75 mean-Jaccard stability threshold. Clusters 2 and 3 were especially reproducible (0.936 and 0.978), and Cluster 4 remained highly stable despite its small size (0.879). Cluster 1 had the lowest mean Jaccard similarity (0.816) but was still stable. Only Cluster 4 dissolved in more than one bootstrap run, with nine dissolutions among 500 samples (1.8%). Cluster 1 was recovered in 357 runs (71.4%), compared with at least 86.2% for the other groups, indicating a less consistently reconstructed boundary despite its mean Jaccard similarity exceeding 0.75. Table 7 and Fig. 5 report bootstrap stability.

Table 7
Bootstrap stability of the four-cluster typology

Cluster	Mean Jaccard	Dissolved runs	Recovered runs	Assessment
1	0.816	1 (0.2%)	357 (71.4%)	Stable
2	0.936	0 (0.0%)	499 (99.8%)	Highly stable
3	0.978	0 (0.0%)	500 (100.0%)	Highly stable
4	0.879	9 (1.8%)	431 (86.2%)	Highly stable
Note. Results are based on 500 bootstrap samples. Recovered runs indicate strong reconstruction of the corresponding original profile.				

3.7 Sensitivity to drought skewness and the highest value

The log-drought partition had high agreement with the original four-cluster solution (ARI = 0.932). The profile with high drought and heavy-rain exposure persisted, with original-unit means of 24.55 m and 6.26%, respectively, although membership shifted and its size increased from 16 to 21 municipalities. The first two PCA components accounted for 74.96% of variance, compared with 70.53% originally. Removing Landgraaf produced ARI = 0.997: one of the remaining 336 municipalities changed assignment after matching cluster labels. The first two components then accounted for 70.76% of variance. These checks support broad partition robustness to this transformation and this single-observation exclusion, while showing that the detailed boundaries depend on preprocessing. They do not prove that every loading, the retained component count, or the optimal cluster count is unchanged. Table 8 summarizes the drought sensitivity checks.

Table 8
Drought sensitivity checks

Analysis	N	ARI	PC1–PC2 variance (%)
Original	337	—	70.53
Log-transformed drought	337	0.932	74.96
Highest drought excluded	336	0.997	70.76
Note. ARI compares each altered four-cluster partition with the original labels for the same municipalities. It is not the percentage of municipalities unchanged. PCA percentages describe the first two components without renewed parallel-analysis retention decisions.			

3.8 Sensitivity to the flood aggregation denominator

The conditional flood aggregation yielded finite estimates for 314 of the 337 original complete cases. In the remaining 23 municipalities, all eligible neighbourhoods had zero coastal or fluvial inundation across the recorded depth categories; the conditional denominator was therefore zero. These municipalities remain valid baseline observations and were omitted only from the paired denominator comparison. The

paired four-cluster partitions had ARI = 0.912 across the 314 municipalities and 0.944 among the 184 meeting the 80% minimum share of population weights assigned to neighbourhoods with recorded inundation. Agreement was complete among the 166 meeting 90% and the 133 meeting 100% (Table 9). The 100% comparison also passed the input-equality control. Each ARI compares re-estimated partitions on identical municipalities within that row; it does not validate all 337 original assignments. Table 9 reports the matched-sample comparisons.

Individual flood summaries changed substantially under the conditional denominator. The largest increase was in Venray (GM0984): the baseline population-weighted average was 5.12%, compared with 65.01% after omitting neighbourhoods with zero inundation in all categories, a difference of 59.88 percentage points. Only 7.88% of eligible population weights were assigned to neighbourhoods with recorded inundation. The increase reflects conditioning on that subset, rather than recovering missing observations. Given the source-provider clarification, the baseline retains the zero-inundation neighbourhoods; the conditional result does not replace it.

Table 9
Flood aggregation denominator sensitivity on matched municipal samples

Minimum share with recorded inundation (%)	Municipalities	ARI	PC1 + PC2 baseline (%)	PC1 + PC2 conditional (%)
0	314	0.912	71.30	70.60
80	184	0.944	65.53	65.46
90	166	1.000	67.23	67.21
100	133	1.000	66.80	66.80
Note. The minimum share refers to the population-weighted share of neighbourhoods with at least one nonblank flood-depth category, indicating recorded inundation. It is not data coverage or a resident-exposure rate. Each row compares the baseline retaining zero-inundation neighbourhoods with a conditional average omitting all-four-blank neighbourhoods, using identical municipalities. ARI is adjusted pairwise agreement, not the percentage of unchanged assignments. PCA percentages alone do not demonstrate unchanged loadings or component retention. Threshold-based sample restriction changes the population represented.				

Table 10
Cluster means on the original standardized five-indicator scale

Cluster	Canopy	Drought	Heat	Flood	Heavy rain
1	1.586	0.202	-1.282	-0.570	0.645
2	-0.075	-0.210	0.440	-0.789	-0.115
3	-0.534	-0.311	-0.043	1.091	-0.443
4	0.010	3.732	0.876	-0.863	2.626
Note. Each indicator was standardized using the mean and sample standard deviation of the 337 complete cases. Entries are cluster means in standard-deviation units, not PCA score centroids. Positive values are above the complete-case sample mean; negative values are below it. Cluster numbering follows Table 6.					

4 Discussion

4.1 A multidimensional municipal indicator structure

The five indicators describe more than one dimension of municipal environmental conditions. PC1 contrasted drought, heavy rain and canopy with flood exposure, while PC2 contrasted heat with canopy. This is relevant to adaptation screening because municipalities with similar overall summaries may still have different dominant pressures. However, no composite index was constructed or compared directly in this study. The findings therefore motivate preserving domain information and evaluating profile summaries; they do not demonstrate that a particular composite ranking fails.

The inverse canopy–heat association is consistent with vegetation's potential contribution to heat mitigation. Schwaab et al. (2021) found lower land surface temperatures associated with urban trees across European cities, with variation across climate contexts. Their surface-temperature outcome differs from the apparent-temperature indicator used here, so that study provides context rather than validation of the present association. Our cross-sectional covariance does not identify a causal canopy effect. Canopy also loaded positively on PC1 alongside drought and heavy rain; this illustrates that principal components summarize observed relationships rather than isolated mechanisms.

4.2 Value of the four-profile typology

The four profiles distinguish configurations relevant to further adaptation assessment. The high-canopy/low-heat group has a relatively favourable thermal configuration but non-zero flood and heavy-rain indicators. The low-flood/rain group combines comparatively low water-exposure indicators with relatively high heat. The flood-exposure group has a dominant flood indicator, whereas the drought-heavy-rain extreme group combines high values in those two domains despite low flood exposure. The groups are neither a complete vulnerability classification nor evidence that the represented hazards occur simultaneously. No cluster is uniformly best or worst across the five variables.

These contrasts can guide the questions asked in local assessment. Flood-exposure municipalities warrant review of flood scenarios, consequences, land use and protective infrastructure; heavy-rain exposure requires attention to drainage and surface-water accumulation. Heat-exposed municipalities may warrant examination of shade and cooling access, while drought depth requires interpretation alongside local groundwater and soil conditions. European risk assessment similarly stresses context and cross-sector adaptation (European Environment Agency 2024). The profiles do not identify which intervention is effective or cost-efficient, and municipal averages cannot identify the neighbourhoods or populations facing the greatest consequences.

Bootstrap and drought sensitivity results support internal reproducibility of the broad profiles. Clusters 2 and 3 had especially high mean Jaccard similarity. Cluster 1's lower similarity and 71.4% recovery suggest a softer boundary; overlap with another profile is plausible but is not established by these summary statistics alone. The log transformation preserved broad agreement while expanding the high-drought/heavy-rain group to 21 municipalities, showing that stable overall structure can coexist with changed boundaries. Excluding the highest drought value changed only one remaining assignment. Robustness to these two changes is useful evidence, but neither resampling nor sensitivity checks validate adaptation outcomes. The flood denominator exercise showed high partition agreement within matched subsets, but the conditional averages could differ substantially from the baseline. The source clarification supports retaining zero-inundation neighbourhoods. Results at higher minimum shares describe selected municipal subsets and cannot be generalized to municipalities omitted from those comparisons.

4.3 Implications for indicator design

Composite indicators simplify communication, but their weighting and aggregation rules can conceal effective influence and compensation across domains (Becker et al. 2017; Greco et al. 2019; OECD and European Union 2008). Profile summaries offer another view of cross-domain configurations. The present typology distinguishes flood exposure from a drought-heavy-rain configuration and can accompany the underlying indicators. It does not eliminate scaling choices or aggregation: standardized k-means distances give equal nominal treatment to five inputs. The relative merits of profiles and composite rankings require a direct comparison against a defined decision task.

This study complements the multi-threat city assessment of Tapia et al. (2017) by examining a compact set of population-weighted indicators within one Dutch municipal system. Its applied contribution is a documented sequence of dimensional assessment, cluster-count comparison, bootstrap stability and sensitivity analysis, with original-unit interpretation. The aim is to make environmental configurations easier to inspect, while retaining the distinction between exposure, protective environmental features and comprehensive vulnerability (Hinkel 2011).

For dashboard use, the results favour presenting the underlying domains alongside any profile label, documenting preprocessing and classification choices, and reporting uncertainty in membership. Labels should be revised if future data or sensitivity checks reveal unstable boundaries. Environmental profiles

also require separate information on social sensitivity and adaptive capacity before they can inform vulnerability assessment. A practical next step is to use contrasting profiles to select municipalities for more detailed investigation, then test whether the grouping improves an explicit planning task.

4.4 Limitations and future research

Several limitations constrain interpretation. The data are cross-sectional and combine mapped indicators with different scenarios and source definitions. Population-weighted neighbourhood averages conceal variation within municipalities and are not municipal resident-exposure rates. Five incomplete drought cases were excluded. Although the source provider confirmed the flood blank-cell convention and neighbourhood-area denominator, municipal aggregation remains sensitive to the weighting scheme and spatial scale. The conditional flood denominator exercise changes the target summary and restricts the samples; it is not evidence of missing flood coverage or a correction to the baseline. PCA summarizes linear covariance, while k-means favours compact clusters and depends on scaling and cluster count. The drought and flood checks address specified transformations and denominator choices, not all sources of uncertainty. Geographic patterning was mapped descriptively, without a spatial clustering test or spatially structured resampling. Conventional bootstrap stability therefore should not be treated as validation against spatial dependence or future observations.

Future work should compare alternative population and area weighting schemes and municipal summaries with neighbourhood-level evidence. It should then compare alternative cluster counts and algorithms, spatially structured validation and subsequent dashboard releases. Incorporating social sensitivity, adaptive capacity, health and housing data would broaden the assessment beyond environmental conditions. Direct comparisons with composite scores should specify a substantive decision criterion, such as identifying municipalities that require different kinds of local assessment, rather than assume one summary method is superior.

5 Conclusions

Five environmental indicators from 337 Dutch municipalities were summarized by two retained components and four interpretable profiles. The components explained 70.53% of variance, and the four-cluster partition distinguished high-canopy/low-heat, low-flood/rain, flood-exposure and drought-heavy-rain extreme configurations. Bootstrap mean Jaccard similarities ranged from 0.816 to 0.978; agreement remained high after log-transforming drought ($ARI = 0.932$) and excluding its highest value (0.997). Flood denominator comparisons yielded ARIs of 0.912–1.000 on matched subsets, while showing substantial changes in conditional averages. The source provider confirmed that blank flood categories represent zero inundation, supporting their retention in the baseline. The substantive contribution is to identify contrasting combinations that can guide further local adaptation assessment. The modest silhouette and sensitivity of some boundaries support using profile labels together with their underlying indicators. The typology is an exploratory screening aid, requiring source-data checks, local assessment and social information before it can support claims about vulnerability or intervention priorities.

Declarations

Competing interests:

The author declares no competing interests.

Ethics statement

Not applicable. The study used aggregated, publicly available municipality-level data and involved no human participants or personal data.

CRedit authorship contribution statement

The author was responsible for conceptualization, methodology, software, formal analysis, visualization, writing—original draft, and writing—review and editing.

Declarations

of generative AI and AI-assisted technologies: The author used ChatGPT (OpenAI) for assistance with R code development and troubleshooting, manuscript structure, language editing and interpretation of reported outputs. The author executed the R analyses and reviewed the resulting outputs. The author is responsible for verifying the code, data, references and final manuscript and takes full responsibility for the content.

Funding:

The author received no specific funding for this work.

Author Contribution

J.O.A. conceived and designed the study, prepared the data, developed and checked the analysis code, performed the statistical analyses, prepared all figures and tables, and wrote and revised the manuscript. J.O.A. reviewed and approved the final manuscript and takes responsibility for the work.

Data Availability

The study used data from the Klimaateffectatlas Buurtdashboard workbook, version 9 March 2026. The processed municipal dataset and R scripts used for the analyses are available from the corresponding author upon reasonable request. Redistribution of source data is subject to the data provider's applicable terms and conditions.

References

1. Becker W, Saisana M, Paruolo P, Vandecasteele I (2017) Weights and importance in composite indicators: Closing the gap. *Ecol Ind* 80:12–22. <https://doi.org/10.1016/j.ecolind.2017.03.056>
2. Birkmann J (2007) Risk and vulnerability indicators at different scales: Applicability, usefulness and policy implications. *Environ Hazards* 7(1):20–31. <https://doi.org/10.1016/j.envhaz.2007.04.002>
3. Cutter SL, Boruff BJ, Shirley WL (2003) Social vulnerability to environmental hazards. *Soc Sci Q* 84(2):242–261. <https://doi.org/10.1111/1540-6237.8402002>
4. European Environment Agency (2024) European climate risk assessment (EEA Report No. 01/2024). Publications Office of the European Union. <https://www.eea.europa.eu/en/analysis/publications/european-climate-risk-assessment>
5. Greco S, Ishizaka A, Tasiou M, Torrisi G (2019) On the methodological framework of composite indices: A review of the issues of weighting, aggregation, and robustness. *Soc Indic Res* 141(1):61–94. <https://doi.org/10.1007/s11205-017-1832-9>
6. Hennig C (2007) Cluster-wise assessment of cluster stability. *Comput Stat Data Anal* 52(1):258–271. <https://doi.org/10.1016/j.csda.2006.11.025>
7. Hinkel J (2011) Indicators of vulnerability and adaptive capacity': Towards a clarification of the science–policy interface. *Glob Environ Change* 21(1):198–208. <https://doi.org/10.1016/j.gloenvcha.2010.08.002>
8. Horn JL (1965) A rationale and test for the number of factors in factor analysis. *Psychometrika* 30:179–185. <https://doi.org/10.1007/BF02289447>
9. Hubert L, Arabie P (1985) Comparing partitions. *J Classif* 2:193–218. <https://doi.org/10.1007/BF01908075>
10. IPCC (2022) Climate change 2022: Impacts, adaptation and vulnerability. Cambridge University Press. <https://doi.org/10.1017/9781009325844>
11. Jolliffe IT, Cadima J (2016) Principal component analysis: A review and recent developments. *Philosophical Transactions of the Royal Society A* 374(2065):20150202. <https://doi.org/10.1098/rsta.2015.0202>
12. Klimaateffectatlas (2026) Buurtdashboard and associated metadata workbook. <https://www.klimaateffectatlas.nl/nl/buurtdashboard>
13. Lloyd S (1982) Least squares quantization in PCM. *IEEE Trans Inf Theory* 28(2):129–137. <https://doi.org/10.1109/TIT.1982.1056489>
14. OECD and European Union (2008) Handbook on constructing composite indicators: Methodology and user guide. OECD Publishing. <https://doi.org/10.1787/9789264043466-en>
15. R Core Team (2026) R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>
16. Rousseeuw PJ (1987) Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J Comput Appl Math* 20:53–65. [https://doi.org/10.1016/0377-0427\(87\)90125-7](https://doi.org/10.1016/0377-0427(87)90125-7)

17. Schwaab J, Meier R, Mussetti G, Seneviratne S, Bürgi C, Davin EL (2021) The role of urban trees in reducing land surface temperatures in European cities. *Nat Commun* 12:6763. <https://doi.org/10.1038/s41467-021-26768-w>

18. Tapia C, Abajo B, Feliu E, Mendizabal M, Martinez JA, Fernández JG, Laburu T, Lejarazu A (2017) Profiling urban vulnerabilities to climate change: An indicator-based vulnerability assessment for European cities. *Ecol Ind* 78:142–155. <https://doi.org/10.1016/j.ecolind.2017.02.040>

Figures




Figure 1

Parallel analysis comparing observed eigenvalues with random 95th-percentile benchmarks




Figure 2

Indicator loadings on the two retained principal components





Figure 3

Environmental profiles indexed to the complete-case sample mean. A value of 100 equals the sample mean. Cluster 4's drought index exceeds 700 because its mean drought depth is more than seven times the overall mean; these indexed values are neither probabilities nor comparable measures of harm.

Environmental Profiles of Dutch Municipalities
Four-cluster solution based on five standardised indicators



Source: Author's calculations from 337 Dutch municipalities; grey areas represent incomplete observations

Figure 4

Geographic distribution of the four municipal environmental profiles




Figure 5

Mean Jaccard stability by cluster. The dashed line marks the predefined 0.75 stability threshold

Supplementary Files

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

- [AppendixAStandardizedprofilecentroids.docx](#)