

Spatial Dependence in Municipal Heat Exposure Across the Netherlands: Evidence from Environmental Indicators and Spatial Error Modelling

Joshua Osejigbe Aguele

joshuaaguele.research@gmail.com

Independent Researcher

Research Article

Keywords: municipal heat exposure, spatial error model, spatial autocorrelation, canopy cover, climate adaptation, Netherlands

Posted Date: October 8th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11252675/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

Spatial Dependence in Municipal Heat Exposure Across the Netherlands: Evidence from Environmental Indicators and Spatial Error Modelling

Joshua Oseijgbe Aguele | Independent Researcher, Netherlands

ORCID: 0009-0008-1214-4008

3rd October 2026

Abstract

This study examines whether associations between municipal heat exposure and four environmental indicators change when spatial dependence is modelled explicitly. Neighbourhood indicators from the Dutch Climate Impact Atlas Buurtdashboard were aggregated using population weights, yielding 337 complete municipal observations. Ordinary least squares (OLS), spatial-lag and spatial-error models were compared using row-standardized queen-contiguity weights. A five-nearest-neighbour specification assessed sensitivity to the weights definition. OLS explained 45.9% of variation in the heat-exposure indicator, but its residuals exhibited positive spatial autocorrelation (Moran's $I = 0.356$, $p < 0.001$). The spatial-error model had the lowest AIC (1881.9, compared with 1896.3 for the spatial-lag model and 1970.7 for OLS). Its residual Moran test was not significant at the 5% level ($I = -0.062$, $p = 0.090$), although this does not establish the absence of residual dependence. Canopy cover remained negatively associated with heat exposure ($B = -0.818$, $p < 0.001$), whereas the groundwater-depth indicator was positively associated ($B = 0.160$, $p = 0.002$). Flood and heavy-rain coefficients were not statistically distinguishable from zero. These coefficient patterns also occurred under five-nearest-neighbour weights. The findings illustrate the importance of spatial diagnostics in municipal environmental comparisons. They support a consistent canopy association within the specifications examined, but do not establish causal effects. Post hoc sensitivity checks retained the positive groundwater-depth association after excluding the maximum value ($B = 0.167$, $p = 0.003$), but the transformed-depth association was not statistically significant after a $\log(1 + \text{depth})$ transformation ($B = 1.025$, $p = 0.185$). Canopy coefficients remained negative and significant in these checks. The groundwater-depth association is therefore specification dependent.

Keywords: municipal heat exposure; spatial error model; spatial autocorrelation; canopy cover; climate adaptation; Netherlands

1 Introduction

Heat is an increasingly important climate risk for European cities and municipalities. Climate projections indicate that warming and heat extremes will intensify adaptation demands, while local land cover, urban form and hydrological conditions shape the distribution of exposure (European Environment Agency, 2024; van Dorland et al., 2023). In the Netherlands, national adaptation policy places increasing emphasis on climate-resilient spatial development, but municipal conditions vary sharply and administrative boundaries do not isolate environmental processes (Climate Adaptation Platform Netherlands, 2026).

Vegetation can moderate local thermal conditions through shading and evapotranspiration, although the relationship depends on spatial scale, urban form and atmospheric conditions (Rafiee et al., 2016; Theeuwes et al., 2017). A

review of Dutch built-environment research emphasizes the relevance of scale to heatwave vulnerability (Ahmed et al., 2023). Empirical work in Rotterdam documents variation in urban heat and thermal comfort (van Hove et al., 2015), while recent datasets provide detailed information for Dutch urban heat assessment (Habib et al., 2025). Municipality-level analysis offers a complementary perspective, but neighbouring municipalities may share landscapes, weather conditions and settlement patterns. Their observations therefore need not be statistically independent.

Spatial dependence can affect estimation and inference when it is not represented appropriately (Anselin, 1988). A spatial pattern in the outcome alone does not determine which model is suitable. This study therefore examines residual dependence, relative model fit and residual diagnostics after estimation. A spatial-lag specification includes a weighted average of neighbouring outcome values. A spatial-error specification represents dependence in the disturbance term. Neither specification, by itself, identifies the physical processes responsible for the observed pattern.

The research question is whether municipal heat exposure retains spatial structure after adjustment for canopy cover, groundwater depth, flood exposure and heavy-rain exposure, and whether their conditional associations differ across model specifications. The study applies a common set of indicators to a national municipal dataset, compares OLS with spatial-lag and spatial-error models, and assesses sensitivity to an alternative weights matrix. Its contribution is an empirical comparison of these specifications for the selected indicators. It does not introduce a new spatial method or estimate causal effects.

2 Data and Methods

2.1 Data and variables

The indicators were obtained from the Klimaateffectatlas Buurtdashboard and its associated metadata workbook (Klimaateffectatlas, 2026). The analysis script aggregated neighbourhood values to municipalities using neighbourhood population weights. For each indicator, the weighted mean used available values with valid positive population weights. These municipal summaries are therefore population-weighted averages of neighbourhood environmental indicators; they are not direct estimates of the percentage of municipal land or residents exposed. The assembled dataset contained 342 municipalities. Five lacked the groundwater-depth indicator, leaving 337 complete observations (98.54%). No duplicate municipal rows were detected. The main regression models used this complete-case sample; the maximum-value exclusion check used 336 observations.

Table 1 Environmental variables and analytical roles

Variable	Unit	Role	Interpretation
Heat exposure 41 plus	%	Outcome	Population-weighted mean of neighbourhood area percentages in PET categories of 41 °C and above
Canopy cover	%	Predictor	Population-weighted mean of neighbourhood canopy-cover percentages

Variable	Unit	Role	Interpretation
Drought depth	m	Predictor	Population-weighted mean of modelled extreme low groundwater depth below ground level
Flood exposure over 50 cm	%	Predictor	Population-weighted mean of neighbourhood area percentages with flood depth above 50 cm in the source scenario
Heavy-rain exposure 20 plus	%	Predictor	Population-weighted mean of neighbourhood area percentages with water depth of at least 20 cm in the source heavy-rain scenario

Note. “Drought depth” is retained as a source-related label in the tables and figures. It refers to modelled extreme low groundwater depth below ground level under current-climate conditions, rather than observed drought damage. The heat indicator combines the source categories with physiological equivalent temperature (PET) of 41 °C and above. Flood and heavy-rain indicators describe scenario-based water-depth thresholds. The source layers represent different environmental conditions and should not be interpreted as measurements of a single observed event.

Descriptive statistics were examined before modelling. The main models retained variables in their original units: coefficients for percentage indicators describe percentage-point differences, and the groundwater-depth coefficient describes the difference associated with one additional metre. The strongly right-skewed groundwater-depth distribution motivated additional post hoc checks using a $\log(1 + \text{depth})$ transformation and exclusion of the maximum observation. All coefficients are interpreted as conditional associations.

2.2 Non-spatial regression and diagnostics

The initial OLS model related municipal heat exposure to the four environmental predictors:

$$\text{Heat}_i = \beta_0 + \beta_1 \text{Canopy}_i + \beta_2 \text{Drought}_i + \beta_3 \text{Flood}_i + \beta_4 \text{HeavyRain}_i + \varepsilon_i \quad (1)$$

Model adequacy was examined using residual plots, the Shapiro–Wilk test, the Breusch–Pagan test, Ramsey’s RESET, variance-inflation factors and Cook’s distance (Breusch & Pagan, 1979; Cook, 1977; Ramsey, 1969; Shapiro & Wilk, 1965). Coefficient inference in the non-spatial benchmark used HC3 heteroskedasticity-consistent standard errors to reduce sensitivity to heteroskedasticity and influential leverage (MacKinnon & White, 1985). The Shapiro–Wilk result was considered separately as evidence of departure from residual normality and as a reason for cautious interpretation of finite-sample inference. Observations with Cook’s distance above $4/N$ were screened in a reduced-sample sensitivity model.

2.3 Spatial weights and residual dependence

Municipal polygons were linked to the analytical dataset by municipality code. The primary weights matrix used queen contiguity: municipalities sharing a boundary point were treated as neighbours. Row standardization made each municipality's weights sum to one. The 337 municipalities had a mean of 5.22 neighbours (range 1–17). Every

municipality had at least one neighbour, but the graph contained two connected components. One consisted of Hulst, Terneuzen and Sluis; the other contained the remaining 334 municipalities. Residual spatial autocorrelation was assessed using Moran's I (Moran, 1950):

$$I = (n/S_0)[\sum_i \sum_j w_{ij}(e_i - \bar{e})(e_j - \bar{e})]/[\sum_i (e_i - \bar{e})^2] \quad (2)$$

Positive Moran's I indicates that neighbouring municipalities tend to have similar residuals under the specified weights matrix. The statistic was evaluated for OLS and spatial-model residuals. A non-significant result is treated as a failure to detect dependence at the chosen significance level, rather than proof that the residuals are spatially independent.

2.4 Spatial models and sensitivity

Two maximum-likelihood models were fitted. The spatial-lag model included a spatially lagged outcome, $y = \rho W y + X\beta + \varepsilon$. The spatial-error model instead represented the disturbance process as spatially dependent:

$$y = X\beta + u, \quad u = \lambda W u + \varepsilon \quad (3)$$

Here, W denotes the row-standardized weights matrix, ρ the spatial-lag parameter and λ the spatial-error parameter. The models were compared using log likelihood, Akaike's information criterion (AIC) and residual Moran tests. Lower AIC indicates better relative fit when models use the same observations and outcome; it does not establish causal validity or predictive performance. The spatial-error model was also estimated with row-standardized five-nearest-neighbour weights to assess sensitivity to the definition of connectivity. This check addresses the weights choice, rather than all aspects of model robustness. The original analysis was conducted in R (R Core Team, 2026). Additional verification and post hoc sensitivity models were estimated independently in Python using PySAL spreg 1.9.1 (ML_Error, sparse LU likelihood evaluation). Municipal geometries were projected to EPSG:28992; the verification adjacency reconstruction used a 0.01 m distance tolerance to match the R queen-contiguity graph. Baseline queen and five-nearest-neighbour spatial-error coefficients and AIC values reproduced the original results to the reported precision. AIC for verification models counted seven estimated parameters, including the disturbance variance.

Two post hoc groundwater-depth checks were conducted. First, depth in metres was replaced by $\ln(1 + \text{depth})$, retaining all 337 observations and the existing weights; this transformation was assessed under both queen and five-nearest-neighbour weights. Second, Landgraaf (GM0882), the municipality with the maximum depth of 59.80 m, was excluded and queen weights were reconstructed for the remaining 336 municipalities. Exclusion was a diagnostic check, not a judgement that this observation was erroneous. Other covariates and the outcome were unchanged. Coefficient p values used asymptotic two-sided normal tests. Raw and transformed coefficients have different scales. AIC comparisons were restricted to models fitted to the same observations.

3 Results

3.1 Descriptive patterns and OLS benchmark

The mean heat-exposure indicator was 65.62% (SD = 6.11), and mean canopy cover was 13.31% (SD = 4.76). Flood exposure had the greatest absolute dispersion among the percentage indicators (mean = 40.06%, SD = 35.91%).

Drought depth was right-skewed: its mean was 3.95 m, its median 2.17 m and its maximum 59.80 m. The large difference between the mean and maximum warrants caution when interpreting a linear groundwater-depth coefficient.

The OLS model explained 45.89% of variation in the heat-exposure indicator (adjusted $R^2 = 0.452$). Its conventional overall F statistic was $F(4, 332) = 70.38$, $p < 0.001$. HC3 standard errors were used for coefficient inference. Canopy cover ($B = -0.943$, $p < 0.001$) and flood exposure ($B = -0.047$, $p < 0.001$) had negative conditional associations; heavy-rain exposure had a positive association ($B = 0.705$, $p = 0.001$). The drought-depth coefficient was not statistically distinguishable from zero ($B = 0.054$, $p = 0.246$). Excluding the 12 observations above the Cook's-distance screening threshold did not change this OLS significance pattern. This reduced-sample check does not establish the robustness of the spatial-model coefficients.

Table 2 OLS benchmark for municipal heat exposure

Predictor	B	HC3 SE	t	p	95% CI
Intercept	77.613	0.981	79.10	< 0.001	[75.683, 79.543]
Canopy cover	-0.943	0.047	-20.24	< 0.001	[-1.035, -0.852]
Drought depth	0.054	0.046	1.16	0.246	[-0.037, 0.144]
Flood exposure	-0.047	0.007	-6.40	< 0.001	[-0.061, -0.032]
Heavy-rain exposure	0.705	0.213	3.31	0.001	[0.286, 1.125]

Note. $N = 337$. $R^2 = 0.459$; adjusted $R^2 = 0.452$; AIC = 1970.742. B values are unstandardized.

3.2 Spatial dependence and model comparison

OLS residuals exhibited positive spatial autocorrelation (Moran's $I = 0.356$, $z = 10.241$, $p < 0.001$). Thus, the four predictors did not account for all spatial structure in the outcome. Both spatial specifications had lower AIC than OLS. The spatial-error model had the lowest AIC (1881.874), 14.38 points below the spatial-lag model and 88.87 points below OLS. It was therefore selected for detailed interpretation among the specifications examined.

Table 3 Comparison of non-spatial and spatial models

Model	Log likelihood	AIC	Spatial parameter	Residual Moran's I	p
OLS	-979.371	1970.742	—	0.356	< 0.001
Spatial lag	-941.127	1896.253	$\rho = 0.454$	0.027	0.387
Spatial error	-933.937	1881.874	$\lambda = 0.616$	-0.062	0.090

Note. All models use the same 337 observations. Spatial models use row-standardized queen-contiguity weights.




Fig. 1 AIC comparison for the non-spatial and spatial models. Lower values indicate better relative fit

3.3 Spatial-error estimates

In the selected spatial-error model, λ was 0.616 (SE = 0.052, $p < 0.001$), indicating positive dependence in the modelled disturbance process. With the other indicators held constant, a one-percentage-point higher canopy-cover indicator was associated with a 0.818-percentage-point lower heat-exposure indicator (95% CI [-0.917, -0.718], $p < 0.001$). Drought depth had a positive coefficient ($B = 0.160$, $p = 0.002$). The confidence intervals for flood and heavy-rain exposure included zero. These estimates describe associations under the specified model and weights, rather than effects of interventions.

Table 4 Spatial-error model for municipal heat exposure

Term	B	SE	z	p	95% CI
Intercept	76.549	1.143	66.99	< 0.001	[74.309, 78.788]
Canopy cover	-0.818	0.051	-16.04	< 0.001	[-0.917, -0.718]
Drought depth	0.160	0.052	3.08	0.002	[0.058, 0.261]
Flood exposure	-0.011	0.009	-1.27	0.205	[-0.029, 0.006]
Heavy-rain exposure	-0.140	0.241	-0.58	0.563	[-0.613, 0.333]
Spatial error λ	0.616	0.052	11.78	< 0.001	[0.513, 0.718]

Note. N = 337. AIC = 1881.874. Coefficients are conditional associations, not causal effects.




Fig. 2 Spatial-error coefficient estimates and 95% confidence intervals. The intercept and λ are omitted for scale

Relative to OLS, the canopy coefficient decreased in magnitude but retained its negative sign. Drought depth was significant in the spatial-error model, whereas flood and heavy-rain exposure were not. These differences demonstrate that the inferential pattern depends on the treatment of spatial dependence; they do not identify the underlying mechanism. Residual Moran's I was -0.062 ($p = 0.090$). Dependence was not detected at the 5% level, but the result should not be interpreted as conclusive evidence of independent residuals.

3.4 Spatial-weights sensitivity

With five-nearest-neighbour weights, the spatial-error model retained the same coefficient signs and significance pattern. It estimated $\lambda = 0.632$, a negative canopy coefficient ($B = -0.849$, $p < 0.001$) and a positive drought-depth coefficient ($B = 0.129$, $p = 0.016$). Flood and heavy-rain coefficients remained non-significant. Residual Moran's I was approximately -0.062 ($p = 0.071$). The queen-contiguity specification had lower AIC (1881.874 versus 1888.348). These results indicate stability across the two weights definitions tested, for the raw-depth specification. The transformation checks below further qualify the groundwater-depth finding.

Table 5 Sensitivity to the spatial-weights definition

Weights	AIC	λ	Residual Moran's I	p
Queen contiguity	1881.874	0.616	-0.062	0.090
Five nearest neighbours	1888.348	0.632	-0.062	0.071

Note. Both matrices were row standardized.

3.5 Groundwater depth sensitivity

The independently reproduced baseline queen model yielded $B = 0.160$ for raw groundwater depth ($p = 0.002$; AIC = 1881.874). Removing Landgraaf retained a positive coefficient ($B = 0.167$, $p = 0.003$; $N = 336$). With all 337 municipalities retained, replacing raw depth with $\ln(1 + \text{depth})$ produced $B = 1.025$ ($p = 0.185$) under queen weights and $B = 0.806$ ($p = 0.311$) under five-nearest-neighbour weights. Canopy cover remained negative and significant in every additional model ($p < 0.001$). These findings show stability to the single maximum-value exclusion, but sensitivity to the depth specification.

Table 6 Groundwater depth sensitivity in spatial error models

Specification	N	Depth B	Depth p	Canopy B	AIC
Queen raw depth	337	0.160	0.002	−0.818	1881.874
Queen maximum excluded	336	0.167	0.003	−0.820	1877.883
Queen log depth	337	1.025	0.185	−0.830	1889.409
Five nearest neighbours raw depth	337	0.129	0.016	−0.849	1888.348
Five nearest neighbours log depth	337	0.806	0.311	−0.856	1893.040

Note. Log depth denotes $\ln(1 + \text{depth in metres})$; its coefficient is not directly comparable with the raw-depth coefficient. Canopy $p < 0.001$ in all models. The excluded maximum was Landgraaf (59.80 m). AIC for the 336-observation model must not be compared with AIC for the 337-observation models. These additional checks did not reassess all original diagnostics.

4 Discussion

4.1 Principal findings

The analysis identifies spatial dependence in municipal heat-exposure residuals and shows that accounting for it changes the interpretation of the selected environmental indicators. OLS explained approximately 46% of outcome variation, yet its residuals remained spatially clustered. The spatial-error specification improved relative fit and produced a different coefficient-significance pattern. This is evidence for considering spatial structure in this dataset, rather than evidence that a particular regional mechanism has been identified.

Canopy cover showed the most consistent association across the specifications examined. Its coefficient was negative in OLS, both baseline spatial-error models and all additional groundwater-depth sensitivity models. This direction is compatible with shading and evapotranspiration and with Dutch research on vegetation and thermal conditions (Rafiee et al., 2016; Theeuwes et al., 2017). However, the municipal aggregate does not distinguish canopy location, species, maturity or water availability. Nor does the analysis establish a causal cooling benefit from additional tree planting. Canopy can be considered in environmental screening, with these measurement and interpretation limits made explicit.

The groundwater-depth association was sensitive to functional form. Excluding Landgraaf did not remove the positive raw-depth association, so the maximum observation alone does not explain that finding. However, the coefficient for $\ln(1 + \text{depth})$ was not significant under either weights matrix. The transformed queen model had higher AIC than the raw-depth model on the same 337 observations. This favours the raw specification in relative fit among these models, but does not establish that its linear association is substantively correct. Statistical suppression or shared regional structure could contribute to the difference from OLS; these mechanisms were not identified. The indicator measures modelled groundwater depth rather than observed drought severity, and its association should be reported as specification dependent. Soil, hydrogeological, land-use and meteorological covariates would help assess its interpretation.

Flood and heavy-rain coefficients were significant in OLS but not in the spatial-error specification. This change cautions against interpreting the OLS estimates as stable cross-hazard relationships. The spatial model accounts for a particular form of residual dependence, but it does not fully control for every spatially patterned confounder. Non-significant coefficients also do not establish that flooding or heavy rain is unrelated to heat in other settings, at other scales or under other specifications.

4.2 Implications for spatial policy

Municipal screening should place local indicators in their regional context. The observed residual structure makes it reasonable to examine neighbouring conditions alongside municipal summaries, rather than assume that administrative units are independent. Regional monitoring and coordinated planning are possible applications, but their effectiveness was not tested here. Decisions about greening or other interventions require finer-scale exposure information, local feasibility assessment and evidence of intervention effects.

The results also caution against treating coefficient significance as a fixed property of an indicator or using this model alone to rank adaptation priorities. Reporting the weights definition, model comparisons, residual tests and sensitivity checks makes the basis for interpretation more transparent. This study evaluates explanatory associations; it does not validate a prediction or decision-support tool.

4.3 Limitations and research priorities

Several limitations constrain interpretation. The cross-sectional observational design cannot establish temporal ordering or causal effects. Municipality-level aggregation can conceal local inequalities and is susceptible to ecological and modifiable-areal-unit effects. Population-weighted neighbourhood averages depend on the availability and validity of neighbourhood values, and they are not direct population-exposure estimates. The source layers have different thresholds, modelling assumptions and reference conditions, so their combination does not describe a single concurrent hazard event. Complete-case analysis excluded five municipalities and may introduce selection bias despite the high retained proportion. The spatial-error model represents residual structure without identifying its source. Its linear form and distributional assumptions also require caution for a bounded percentage outcome. Post hoc groundwater-depth checks revealed sensitivity to transformation, despite stability after exclusion of the maximum value. These checks do not exhaust possible non-linear forms, influential observations or omitted covariates. They were prompted by the observed distribution and should be interpreted as exploratory. Agreement between independent implementations verifies numerical estimation on the supplied dataset; it does not validate source measurements, aggregation choices or model assumptions.

Further work should examine alternative non-linear groundwater-depth specifications, broader influence diagnostics and residual distributional assumptions. Environmental indicators could then be linked to urban form, soil and groundwater conditions, meteorology and demographic sensitivity. Neighbourhood or multilevel analysis would help distinguish local variation from municipality-level patterns. Longitudinal or intervention-based designs would be needed to investigate temporal change and causal effects. Validation under alternative boundaries and connectivity definitions would further clarify the scope of the findings.

5 Conclusion

Municipal heat-exposure residuals exhibited spatial autocorrelation after adjustment for four environmental indicators. Among the models examined, the spatial-error specification had the lowest AIC and a non-significant residual Moran test at the 5% level. Canopy cover retained a negative conditional association under both spatial-weights definitions and the additional groundwater-depth checks. The positive raw-depth association survived exclusion of the maximum value, but was not significant after log transformation under either weights definition. It should therefore be regarded as specification dependent. Flood and heavy-rain coefficients were not significant in the spatial-error models. These findings illustrate the value of spatial diagnostics and transparent model comparison, while leaving causal mechanisms and adaptation effectiveness unresolved.

Statements and Declarations

Funding: The author received no specific funding for this work.

Competing interests: The author declares no competing interests.

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

Consent to participate and publish: Not applicable.

Data and code availability: The source indicators and metadata are available through the Dutch Climate Impact Atlas Buurtdashboard (<https://www.klimaat-effectatlas.nl/nl/buurtdashboard>). Analysis scripts, documentation and selected verification outputs are available at <https://github.com/joshuaaguele/dutch-environmental-analysis>. The repository contains source-data acquisition guidance rather than a complete redistribution of all input data. Reuse of source datasets is subject to the providers' licensing conditions. No archived release DOI is claimed in this version.

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

Use of generative AI: ChatGPT (OpenAI) assisted with manuscript organization, English-language revision, methodological interpretation and code used for independent numerical verification and post hoc sensitivity analysis. The original R spatial-error results were reproduced with a separate Python implementation, and additional groundwater-depth checks are reported explicitly in this revision. This computational agreement does not constitute independent peer review. The author is responsible for checking the manuscript, references, code and results before public posting and for all claims in the final version.

References

- Ahmed, I., van Esch, M. M. E., & van der Hoeven, F. D. (2023). Heatwave vulnerability across different spatial scales: Insights from the Dutch built environment. *Urban Climate*, 51, 101614. <https://doi.org/10.1016/j.uclim.2023.101614>
- Anselin, L. (1988). *Spatial econometrics: Methods and models*. Kluwer Academic Publishers.

- Breusch, T. S., & Pagan, A. R. (1979). A simple test for heteroscedasticity and random coefficient variation. *Econometrica*, 47(5), 1287–1294. <https://doi.org/10.2307/1911963>
- Climate Adaptation Platform Netherlands. (2026). National Climate Adaptation Strategy. <https://klimaatadaptatienederland.nl/en/policy-programmes/national-climate-adaptation-strategy/>
- Cook, R. D. (1977). Detection of influential observation in linear regression. *Technometrics*, 19(1), 15–18. <https://doi.org/10.2307/1268249>
- 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>
- Habib, M. M., van Esch, M., van Ham, M., & Timmermans, W. J. (2025). High-resolution datasets for urban heat vulnerability assessment in urbanized areas of the Netherlands. *Data in Brief*, 60, 111525. <https://doi.org/10.1016/j.dib.2025.111525>
- KlimaatEffectatlas. (2026). Buurtdashboard and associated metadata workbook. <https://www.klimaatEffectatlas.nl/nl/buurtdashboard>
- MacKinnon, J. G., & White, H. (1985). Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of Econometrics*, 29(3), 305–325. [https://doi.org/10.1016/0304-4076\(85\)90158-7](https://doi.org/10.1016/0304-4076(85)90158-7)
- Moran, P. A. P. (1950). Notes on continuous stochastic phenomena. *Biometrika*, 37(1–2), 17–23. <https://doi.org/10.1093/biomet/37.1-2.17>
- R Core Team. (2026). R: A language and environment for statistical computing. R Foundation for Statistical Computing.
- Rafiee, A., Dias, E. S., & Koomen, E. (2016). Local impact of tree volume on nocturnal urban heat island: A case study in Amsterdam. *Urban Forestry & Urban Greening*, 16, 50–61. <https://doi.org/10.1016/j.ufug.2016.01.008>
- Ramsey, J. B. (1969). Tests for specification errors in classical linear least-squares regression analysis. *Journal of the Royal Statistical Society Series B*, 31(2), 350–371.
- Shapiro, S. S., & Wilk, M. B. (1965). An analysis of variance test for normality. *Biometrika*, 52(3–4), 591–611. <https://doi.org/10.1093/biomet/52.3-4.591>
- Theeuwes, N. E., Steeneveld, G. J., Ronda, R. J., & Holtslag, A. A. M. (2017). A diagnostic equation for the daily maximum urban heat island effect for cities in northwestern Europe. *International Journal of Climatology*, 37(1), 443–454. <https://doi.org/10.1002/joc.4717>
- van Dorland, R., Beersma, J., Bessembinder, J., Bloemendaal, N., van den Brink, H., Brotons Blanes, M., Drijfhout, S., Haarsma, R., Keizer, I., Krikken, F., Le Bars, D., Lenderink, G., van Meijgaard, E., Meirink, J. F., Reerink, T., Selten, F., Severijns, C., Siegmund, P., Sterl, A., ... van der Wiel, K. (2023). KNMI National Climate Scenarios 2023 for the Netherlands (WR-23-02). Royal Netherlands Meteorological Institute. <https://www.knmi.nl/kennis-en-datacentrum/publicatie/knmi-national-climate-scenarios-2023-for-the-netherlands>

van Hove, L. W. A., Jacobs, C. M. J., Heusinkveld, B. G., Elbers, J. A., van Driel, B. L., & Holtslag, A. A. M. (2015). Temporal and spatial variability of urban heat island and thermal comfort within the Rotterdam agglomeration. *Building and Environment*, 83, 91–103. <https://doi.org/10.1016/j.buildenv.2014.08.029>