The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume L-4/W1-2026
Free and Open Source Software for Geospatial (FOSS4G) 2026 – Academic Track "Bridging Geospatial Technology and Humanity",
30 August–5 September 2026, Hiroshima, Japan

Multivariate Spatio-Temporal Modeling for Regional GIS Data: A Statistical Framework for Analyzing Multidimensional Spatial Interactions

Saeko Ohta[1], Shojiro Tanaka[2], Ryuei Nishii[3]

[1] Faculty of Human Health Sciences, Meio University, Nago, Okinawa, Japan – s.ota@meio-u.ac.jp
[2] Faculty of Media Business, Hiroshima University of Economics, Hiroshima, Japan – sh-tanaka@hue.ac.jp
[3] School of Information and Data Sciences, Nagasaki University, Nagasaki, Japan – nishii.ryuei@nagasaki-u.ac.jp

Keywords: Choropleth visualization, interactive web mapping, model comparison, open-source GIS, spatial econometrics, spatio-temporal model.

Abstract

Open geospatial datasets increasingly contain multiple regional indicators that evolve jointly across space and time, but applied workflows often fit separate regressions for each indicator and communicate results through static choropleth maps. We present an open-source, reproducible workflow for comparing multivariate spatio-temporal regression models that allow for cross-equation dependence against separate single-equation spatial regressions using interactive map-based diagnostics. The statistical core is a multivariate generalized nesting spatio-temporal (MGNST) framework that represents spatial-lag dependence, spatial-error dependence, temporal autoregression, and cross-equation dependence; independent spatial error models (SEMs) and ordinary regressions are obtained as nested restrictions. Eleven specifications are estimated by maximum likelihood and evaluated using AIC and BIC. A Shiny–leaflet web application loads precomputed model outputs and allows users to inspect observed, fitted, residual, and residual-difference choropleths for municipalities in the Kansai region of Japan. In a synthetic benchmark generated from the full MGNST model, AIC selects the true multivariate specification for all lattice sizes and ranks it above all separate regressions. In the Kansai case study, the independent SEM is marginally preferred by AIC and BIC, while the multivariate SEM estimates statistically significant cross-equation spatial-error dependence. The application makes this near-tie interpretable beyond numerical values by showing where residuals from the two models differ across municipalities. The released data, spatial weights, source code, model outputs, and application provide a reusable FOSS4G workflow for moving from choropleth visualization to spatial model diagnosis.

1. Introduction

Open geospatial data have changed regional analysis. Government statistics, administrative boundaries, population registers, economic surveys, environmental indicators, and infrastructure data can now be joined, mapped, and analyzed with free and open-source software. Yet the analytical habit has not kept pace. Regional data rarely consist of a single isolated outcome: income, population change, ageing, density, public services, and economic activity often vary together, show spatial dependence among neighboring municipalities, and persist over time. The common practice, however, is still to fit a *separate single-equation regression* to each variable and to communicate the result through a static choropleth map. Such a workflow is useful for exploration, but it cannot show which spatial patterns are jointly explained across variables, which are persistent over time, and which remain as spatially structured residuals.

Spatial regression models provide a statistical language for this problem. The spatial autoregressive (SAR) model represents direct spillovers in the response, whereas the spatial error model (SEM) represents spatially correlated unobserved factors. The generalized nesting spatial (GNS) model combines both structures and has become a flexible model family in spatial econometrics (Anselin, 1988, LeSage and Pace, 2009, Elhorst, 2014). Panel and spatio-temporal extensions additionally incorporate temporal dependence (Lee and Yu, 2010). Applied equation by equation, however, these models still analyze each response on its own. For modern open regional data this is restrictive: several outcomes may share latent spatial causes, and one response may help explain another through cross-equation dependence (Cressie and Wikle, 2011, Banerjee et al., 2014, Bradley et al., 2018, Mastrantonio et al., 2019).

A multivariate spatio-temporal model can couple the responses, but this raises a practical question that a coefficient table answers poorly: for a given dataset, does multivariate coupling actually change the spatial fit, and *where*? Practitioners accustomed to map-based reasoning require a means of *seeing* this answer rather than inferring it from parameter estimates alone. This paper therefore places an interactive, open-source map application at the center of the workflow. The application loads pre-computed outputs from a multivariate generalized nesting spatio-temporal (MGNST) model and from its single-equation special cases, and allows the analyst to compare them visually—observed, fitted, residual, and residual-difference choropleths—for any pair of specifications.

The contribution is fourfold. First, we provide a unified multivariate spatio-temporal regression framework (MGNST) that represents spatial lag, spatial error, temporal autoregressive, and cross-equation dependence, and that contains separate SAR, SEM, GNS, and ordinary regressions as nested restrictions. Second, we estimate eleven nested specifications by maximum likelihood and compare them primarily with the Akaike information criterion (AIC), while reporting the Bayesian information criterion (BIC) as a more parsimonious sensitivity criterion. Third—and centrally—we release an interactive open-source web map application that turns model comparison into a navigable visual diagnosis. Fourth, we package the whole pipeline—data, spatial weights, estimation scripts, model outputs, and the deployed application—as a reproducible repository (Ohta et al., 2026).

The paper does not present a proprietary GIS procedure or a closed empirical workflow. All stages use open data and free/open-source statistical and geospatial software. Section 2 describes the workflow. Section 3 summarizes the MGNST model, the nested family, and the information-criterion comparison. Section 4 reports the synthetic benchmark, in which AIC recovers the true multivariate model when cross-equation structure is present. Section 5 describes the interactive web application. Section 6 applies the workflow to municipal data for the Kansai region of Japan and uses the application for visual comparison. Section 7 discusses reproducibility and reuse.

2. Open-Source Geospatial Workflow

The workflow starts from open regional data and ends with an interactive map application for visual model comparison (Figure 1). In the empirical case study, the analysis unit is the municipality. Processed panel data, a row-standardized spatial weight matrix, model-estimation scripts, and model outputs are stored together so that the statistical results are not separated from the geospatial preprocessing steps. The repository keeps the model library, execution scripts, and the application as separate, clearly labeled components.

The estimation layer is R-centered. The model-estimation scripts use open-source packages including `spdep`, `spatialreg`, `numDeriv`, and `systemfit`; the geospatial preparation uses `sf`, `dplyr`, `tidyr`, `stringr`, and `rmapshaper`. This keeps the analysis close to established spatial statistics and geospatial tooling (Pebesma and Bivand, 2023, Bivand and Piras, 2015). The municipal boundary polygons are obtained from public Japanese administrative boundary data, and the row-standardized adjacency matrix is supplied as a processed CSV so that the models can be re-estimated without reconstructing the geometry.

The workflow has six reproducible stages. First, open socioeconomic indicators and boundary data are obtained from public sources (Ministry of Land, Infrastructure, Transport and Tourism, 2021). Second, municipalities are harmonized and joined to the panel data. Third, a row-standardized spatial weight matrix is prepared from adjacency relationships. Fourth, each candidate specification—both the multivariate models and the separate single-equation regressions—is estimated by maximum likelihood with the same responses, covariates, time lag, and weight matrix. Fifth, AIC provides the primary ranking of the nested specifications, BIC is reported as a sensitivity criterion, and observed values, fitted values, and residuals are exported as pre-computed tables. Sixth, these outputs are published as an interactive map application, so that the observed, fitted, residual, and residual-difference choropleths of any specification can be inspected and compared without a local R installation. The same directory structure also contains synthetic data and true parameter values, making the repository a benchmark for testing independent implementations.

The application serves as the visual hub of the workflow rather than as a supplementary component. It is built with open-source web-mapping tools, deployed to a public cloud endpoint, and released with the full workflow in a public repository (Ohta et al., 2026); it exposes the observed, fitted, residual, and residual-difference choropleths for each candidate model. A reader can place a multivariate specification alongside a separate single-equation regression and view the residual maps and their difference side by side. This transforms a static figure sequence into a navigable diagnostic surface, while keeping the underlying data, weights, and estimation scripts identical to the reproducible package. A snapshot of the deployed application is shown in Figure 2.

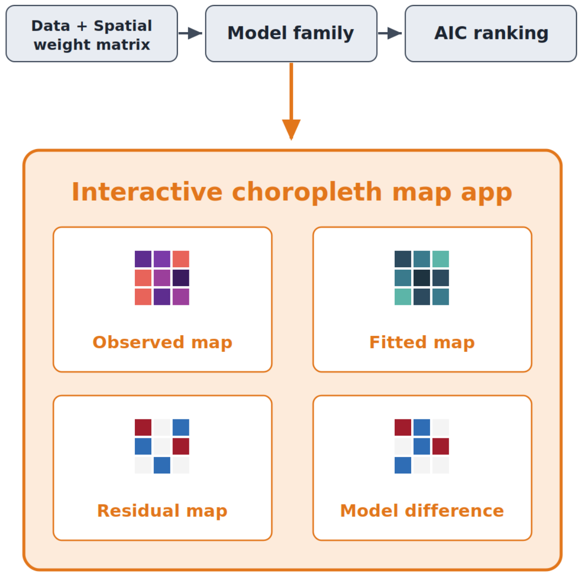

Figure 1. Open-source geospatial workflow from data preparation and model estimation to pre-computed interactive diagnostic maps.

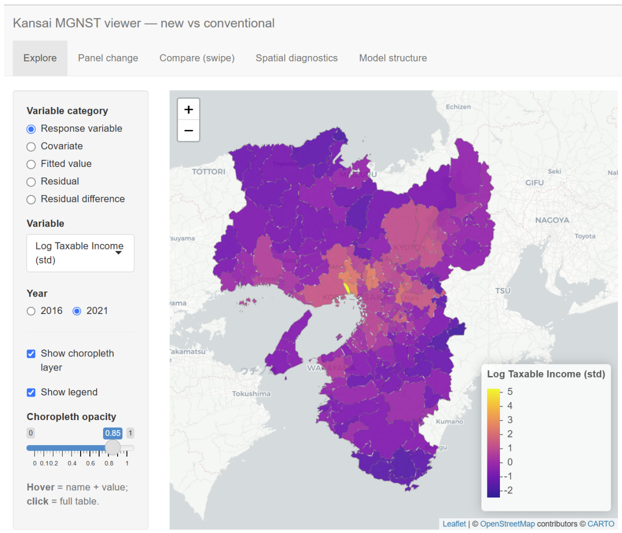

Figure 2. Main Explore view of the Shiny–leaflet application, showing variable, model-output, year, legend, and opacity controls.

In this workflow the statistical model does not terminate at a coefficient table. Maps are employed at three points: before modeling, to reveal regional patterns in the observed variables; during diagnosis, to compare fitted surfaces and residual structure across specifications; and after model comparison, to communicate which dependence structures the data support. The interactive application carries this final stage from static publication to interactive exploration, completing the transition from choropleth mapping to spatial diagnosis.

3. Statistical Model and Model Comparison

3.1 MGNST Formulation

Let $\boldsymbol{y}_{1,t}, \ldots, \boldsymbol{y}_{K,t}$ be K response vectors observed over the same n regions at time t, and define the stacked vector $\boldsymbol{y}_t = (\boldsymbol{y}_{1,t}^\top, \ldots, \boldsymbol{y}_{K,t}^\top)^\top$ of length Kn. Let W denote an $n \times n$ row-standardized spatial weight matrix. The MGNST model is

$$\boldsymbol{y}_t = (R \otimes W)\boldsymbol{y}_t + X_t\boldsymbol{\beta} + (A \otimes I_n)\boldsymbol{y}_{t-1} + \boldsymbol{u}_t, \quad (1)$$

with spatially correlated errors

$$\boldsymbol{u}_t = (\Lambda \otimes W)\boldsymbol{u}_t + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim N(\boldsymbol{0}_{Kn}, \Sigma \otimes I_n). \quad (2)$$

Here $\boldsymbol{\beta}$ is the regression-coefficient vector, $\boldsymbol{u}_t$ is the spatially correlated disturbance vector, $\boldsymbol{\varepsilon}_t$ is the innovation vector, I_n is the $n \times n$ identity matrix, I_{Kn} is the identity matrix of order Kn, and $\boldsymbol{0}_{Kn}$ is a zero vector of length Kn. The matrices R, Λ, and A are $K \times K$ matrices representing spatial-lag dependence, spatial-error dependence, and one-period temporal autoregressive dependence, respectively. The symbol $\otimes$ denotes the Kronecker product, so that, for example, $R \otimes W$ expands the $K \times K$ cross-equation matrix against the $n \times n$ spatial weight matrix to act on the full stacked vector of length Kn. The covariance matrix Σ is $K \times K$ and captures contemporaneous cross-equation dependence. The design matrix X_t is block diagonal, allowing common and response-specific covariates to enter different equations with different coefficients. The likelihood is evaluated over the parameter space in which $I_{Kn} - R \otimes W$ and $I_{Kn} - \Lambda \otimes W$ are nonsingular.

Equations (1)–(2) contain familiar models as special cases. If $\Lambda = O_K$, where O_K denotes the $K \times K$ zero matrix, the model becomes a multivariate SAR-type specification; if $R = O_K$ it becomes a multivariate SEM-type specification. Crucially, when the spatial and covariance matrices are restricted to be *diagonal* (or absent), the cross-equation links are severed and each response follows its own spatial process: the model then reduces to the separate single-equation spatial regressions that an analyst would obtain by fitting one equation at a time. This nesting is the basis of the comparison developed below.

3.2 Nested Model Family and Information-Criterion Comparison

The implementation labels each candidate by a four-character identifier describing restrictions on (R, Λ, A, Σ), where `1`, `d`, and `0` indicate full, diagonal, and absent matrices. The eleven specifications split into two groups (Figure 3). The *multivariate methodology* comprises the specifications with at least one full off-diagonal matrix, so that information is shared across equations: the fully connected MGNST model `1111`, the multivariate spatial-error model `0111`, the multivariate spatial-lag model `1011`, and the cross-equation temporal/covariance variants `1101`, `0101`, `1001`, and the vector-autoregressive `0011`. The *single-equation regressions* comprise the specifications in which every coupling matrix is diagonal or absent, reproducing what separate, one-equation-at-a-time spatial regressions would yield: the independent GNS `dddd`, independent SAR `d0dd`, independent SEM `0ddd`, and the ordinary independent regression `000d`.

This notation turns model comparison into a transparent question: do the data support cross-equation coupling on top of the separate single-equation structure? It also bridges statistical modeling and map-based interpretation. If a multivariate specification is preferred, its residual maps should show weaker spatial structure than the corresponding single-equation regression; if the single-equation regression already captures the structure, the residual maps—and their difference—should be nearly indistinguishable.

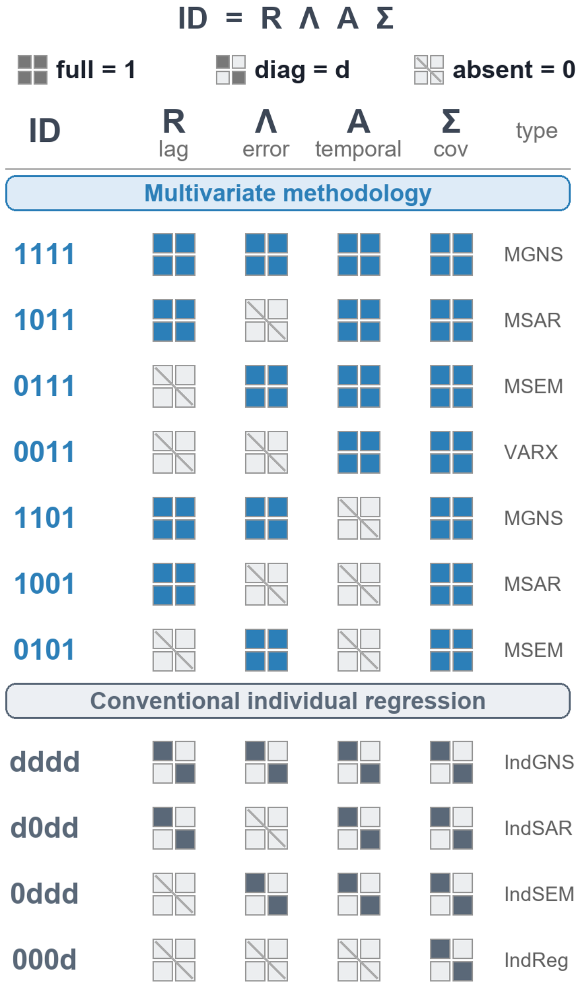

Figure 3. Nested model family. In each four-character ID, `1`, `d`, and `0` denote full, diagonal, and absent matrices for R, Λ, A, and Σ, respectively.

Every specification is fitted by maximum likelihood. The fitted specifications are then compared by the Akaike and Bayesian information criteria,

$$\mathrm{AIC} = -2\ell(\hat{\theta}) + 2p, \qquad \mathrm{BIC} = -2\ell(\hat{\theta}) + \log(Kn)\,p, \quad (3)$$

where $\ell(\hat{\theta})$ is the maximized log-likelihood and p is the number of free parameters of the specification. AIC is used as the primary comparison criterion because the benchmark is designed to test recovery of the full data-generating structure, whereas BIC is reported as a parsimony-oriented sensitivity criterion. With two observation years, the likelihood is evaluated conditionally on the first year; accordingly, the temporal term is interpreted as a one-period persistence component rather than as evidence from a long time series. The BIC penalty uses Kn, the number of stacked regional responses in the modeled transition.

4. Synthetic Benchmark

The synthetic benchmark verifies that, when the data contain genuine cross-equation structure, AIC recovers the true multivariate model and assigns lower values (better values) than the separate equation regressions, which omit that structure. Synthetic data are generated on square lattices with $n = 100$, 400, and 900 regions and 2 time points. The true data-generating process is the full `1111` MGNST model with $K = 2$ responses, shared covariates, response-specific covariates, spatial lag dependence, spatial error dependence, temporal autoregressive effects, and cross-equation covariance. The repository includes the generated response/covariate data, spatial weight matrices, and true parameter values for each lattice size, so that the benchmark can be reproduced or re-used to test independent implementations.

For each n, all eleven specifications are estimated and compared by AIC and BIC (Table 1). In the AIC panel, the full multivariate model `1111` is selected at every lattice size, with AIC 108.85 ($n = 100$), 448.06 ($n = 400$), and 810.66 ($n = 900$). The gap to the best separate single-equation regression is large and grows with n: at $n = 900$ the best single-equation model, the independent GNS `dddd`, has AIC 1070.05 (a difference of about 260), the independent SEM `0ddd` has 1427.71, and the ordinary independent regression `000d` has 5229.24. Thus, when several responses genuinely share spatial, temporal, and cross-equation structure, fitting them one equation at a time leaves substantial dependence unmodeled, and AIC penalizes this severely. The closest AIC competitor is the multivariate spatial-lag model `1011` (AIC 827.50 at $n = 900$); at every lattice size the two best-ranked AIC specifications, `1111` and `1011`, are multivariate and outrank every separate single-equation regression. BIC, with its stronger penalty, instead selects the more parsimonious multivariate spatial-lag model `1011` at all three sizes. Both criteria therefore select a multivariate specification rather than a separate single-equation regression in the synthetic benchmark.

(a) AIC

Model	$n = 100$	$n = 400$	$n = 900$
Multivariate methodology			
`1111`	**108.85**	**448.06**	**810.66**
`1011`	109.29	459.77	827.50
`0111`	156.20	581.70	1177.34
`0011`	198.53	868.96	1632.89
`1101`	471.32	1850.83	4223.10
`1001`	475.07	1860.84	4244.97
`0101`	477.34	1857.27	4259.44
Individual spatial regressions			
`dddd`	133.45	540.67	1070.05
`d0dd`	130.55	559.42	1090.23
`0ddd`	175.81	677.12	1427.71
`000d`	579.45	2286.12	5229.24

(b) BIC

Model	$n = 100$	$n = 400$	$n = 900$
Multivariate methodology			
`1111`	184.71	555.80	937.06
`1011`	**171.95**	**548.78**	**931.92**
`0111`	218.87	670.71	1281.75
`0011`	248.00	939.23	1715.32
`1101`	533.99	1939.84	4327.52
`1001`	524.55	1931.11	4327.40
`0101`	526.82	1927.54	4341.88
Individual spatial regressions			
`dddd`	186.22	615.63	1157.98
`d0dd`	176.73	625.00	1167.17
`0ddd`	221.99	742.70	1504.65
`000d`	612.43	2332.96	5284.19

Table 1. AIC and BIC for synthetic data ($n = 100, 400, 900$; true model `1111`). Minima in each panel are shown in bold.

5. Open-Source Interactive Web Application

The interactive map application (Figure 4) is the central deliverable of the workflow. Its purpose is to make the comparison between the multivariate cross-equation methodology and separate single-equation regressions a visual, exploratory act rather than a reading of coefficient tables. It is built with the open-source `leaflet` and `shiny` stack. The source code, pre-computed inputs, model outputs, and application files are archived in the versioned workflow release (Ohta et al., 2026), and the same application is also deployed as a live Shiny instance for browser-based inspection.[1]

The application is organized around three diagnostic questions: what spatial pattern is observed, what part of that pattern is fitted by each model, and where residuals change when moving from a separate single-equation model to a multivariate specification. Because all map layers are generated from the same pre-computed maximum-likelihood outputs reported in the tables, visual exploration remains directly tied to the reproducible statistical workflow.

Architecture. The application follows a strict "read-only" principle. All eleven specifications are fitted in advance by the batch estimation scripts, and their observed values, fitted values, and residuals are written to pre-computed files that the application loads at start-up. This keeps the application lightweight and cleanly separates the reproducible estimation pipeline (raw data $\rightarrow$ maximum-likelihood estimation $\rightarrow$ derived outputs) from the visual layer. The simplified municipal polygons and the derived map tables are shipped with the repository so that the application also runs locally after a single clone.

Layers and interaction. The map shows the 198 Kansai municipalities as the active layer, with neighboring prefectures left to the base tiles for geographic context. A category selector exposes five layer types: the observed responses, the covariates, the fitted values of any of the eleven models, the residuals of any of the eleven models, and the *residual difference* between any two chosen models. Responses and covariates are available for both observation years (2016 and 2021), while fitted values, residuals, and residual differences are mapped for the most recent year. Observed responses use sequential color scales (one per response), while residuals and residual differences use a symmetric, zero-centered diverging scale so that positive and negative departures are immediately distinguishable. Hovering reveals a municipality's name and the selected value; clicking opens a panel with the full set of values for that municipality. A legend toggle, a layer-opacity slider, and a busy-state overlay complete the interface.

Visual comparison of multivariate and individual equation models. The residual-difference layer is central to the comparison. By selecting a multivariate specification as model A and a separate single-equation regression as model B, the analyst obtains the map of $e_{A,i} - e_{B,i}$, the per-municipality difference in residuals for municipality i. Where the multivariate coupling changes the fitted surface, this map shows coherent

[1] Live application: bit.ly/4e8Tmnt (14 June 2026).

spatial structure; where the single-equation regression already suffices, it is close to flat. The same control therefore answers, visually and locally, the question that a single information criterion cannot: not only *whether* the multivariate methodology changes the fit, but *where* on the map it does so. Section 6 uses exactly this control to compare the multivariate spatial-error model `0111` with the independent spatial-error model `0ddd`.

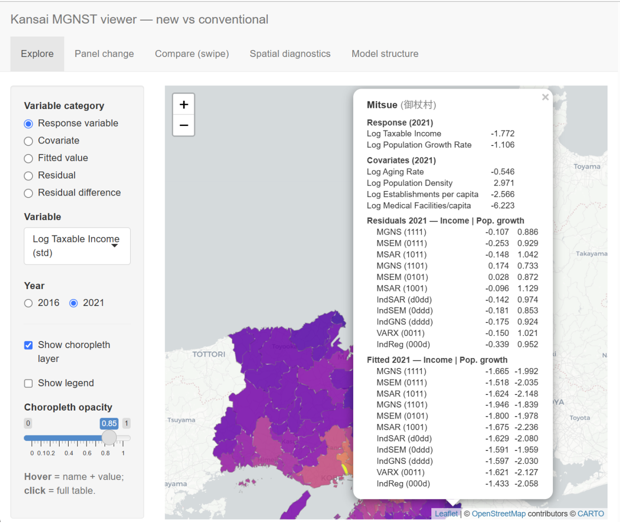

Figure 4. Municipality-level inspection in the interactive application: clicking a municipality opens observed, fitted, and residual values for the selected specifications.

6. Empirical Application: Kansai Municipal Data

6.1 Data Sources and Variables

The empirical workflow is applied to municipal-level socioeconomic data for the Kansai region of Japan (Figure 5). The dataset contains $n = 198$ municipalities observed at two time points, 2016 and 2021. The response variables are standardized log taxable income per capita ($\boldsymbol{y}_1$) and standardized log population growth ($\boldsymbol{y}_2$). Here population growth is represented in the processed data as a log population ratio over the corresponding observation interval, rather than as the logarithm of a signed percentage change. The common explanatory variables are log ageing rate and log population density. The response-specific covariates are log establishments per capita for the income equation and log medical facilities per capita for the population-growth equation. All variables are standardized after the logarithmic transformations used in the repository.

The data are built from public sources including Japanese government statistics and administrative boundary data. Taxable income is obtained from municipal taxation surveys, population and ageing variables from the Basic Resident Register statistics, establishment counts from the Economic Census, medical-facility counts from medical institution surveys, and municipal areas/boundaries from public geospatial data sources (Statistics Bureau of Japan, 2021, Ministry of Land, Infrastructure, Transport and Tourism, 2021). The spatial weight matrix is constructed from municipal adjacency relationships and row-standardized before estimation. The repository contains processed inputs and documents the original source URLs so that the data can be re-obtained if needed.

The observed maps in Figure 6 motivate joint modeling. Taxable income concentrates around the Kyoto–Osaka–Kobe

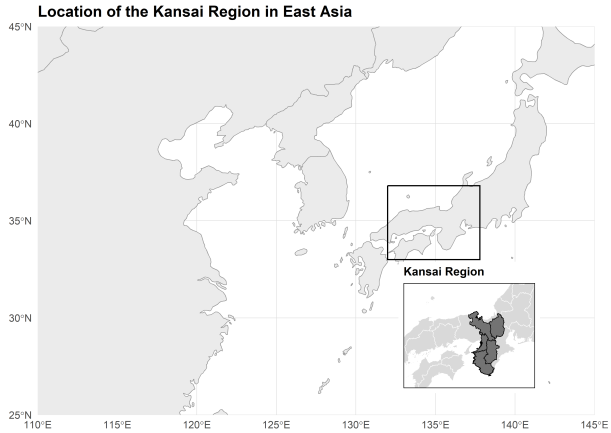

Figure 5. Study area: the Kansai region of Japan.

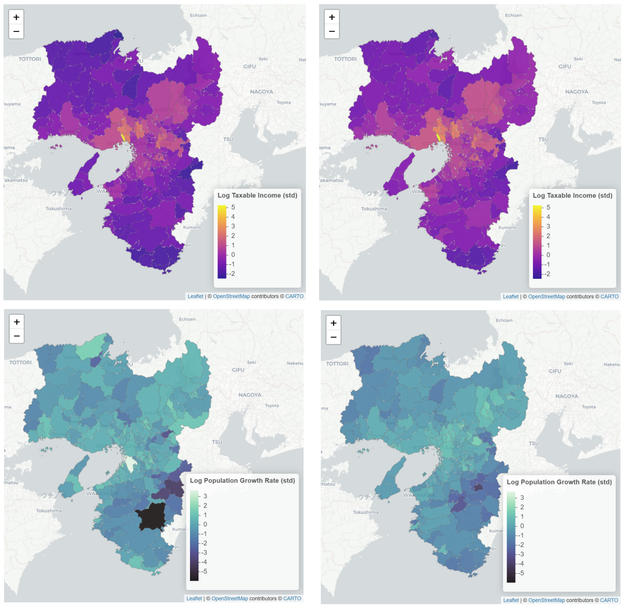

Figure 6. Observed standardized responses by municipality: log taxable income ($\boldsymbol{y}_1$, top) and log population growth ($\boldsymbol{y}_2$, bottom); 2016 (left), 2021 (right).

metropolitan axis, while population growth shows a different pattern of suburban growth and peripheral decline. The two outcomes are related but not spatially identical—the type of case in which separate single-equation regressions might miss shared latent regional structure, and in which a multivariate methodology should be examined rather than assumed.

6.2 Model Comparison

The eleven specifications are estimated by maximum likelihood $\ell(\hat{\theta})$ in equation (3) using the same responses, covariates, one-period time lag, and row-standardized municipal adjacency matrix, and are ranked primarily by AIC, with BIC reported as a sensitivity criterion (Table 2). The four best-fitting specifications by AIC all combine spatial-error and temporal dependence: the independent SEM `0ddd` (AIC 69.71), the multivariate SEM `0111` (AIC 70.38), the independent GNS `dddd` (AIC 71.98), and the full MGNST `1111` (AIC 74.12); the latter two additionally include a spatial-lag term. The pure spatial-lag specifications (`d0dd`, `1011`) rank lower, so spatially structured unobserved factors and one-period persistence—rather than di-

https://doi.org/10.5194/isprs-archives-L-4-W1-2026-197-2026

(a) AIC

Model	p	AIC
Multivariate methodology		
`0111`	19	70.38
`1111`	23	74.12
`1011`	19	80.50
`0011`	15	80.62
`0101`	15	554.93
`1101`	19	555.28
`1001`	15	567.88
Individual spatial regressions		
`0ddd`	14	**69.71**
`dddd`	16	71.98
`d0dd`	14	79.80
`000d`	10	645.71

(b) BIC

Model	p	BIC
Multivariate methodology		
`0111`	19	146.03
`1111`	23	165.69
`1011`	19	156.15
`0011`	15	140.35
`0101`	15	614.65
`1101`	19	630.92
`1001`	15	627.60
Individual spatial regressions		
`0ddd`	14	**125.45**
`dddd`	16	135.68
`d0dd`	14	135.54
`000d`	10	685.53

Table 2. AIC and BIC for the Kansai data; p = number of free parameters. Minima in bold.

rect spatial spillover—dominate the income and population dynamics.

For these two responses the joint multivariate and separate single-equation structures fit comparably: the best single-equation regression `0ddd` and the best multivariate specification `0111` differ by less than one AIC unit, and AIC does not favor the five additional parameters of `0111`. Both AIC and BIC select `0ddd` for these data, with BIC giving a stronger parsimony penalty. This near-tie is itself informative, and it is precisely the situation in which a single criterion value is insufficient: it indicates that the two models fit comparably, but not whether they describe the same spatial residual surface. The multivariate model `0111` still estimates cross-equation spatial-error dependence—an off-diagonal element of Λ that is statistically significant—and a cross-equation covariance, both of which the independent model `0ddd` is *structurally unable to represent*; that this added structure is statistically detectable yet not preferred by AIC reflects the cost of the extra parameters. The visual comparison in Section 6.3 lets the analyst inspect where the two specifications differ across municipalities.

The leading specifications also give coherent and stable substantive effects. The ageing rate is strongly negatively associated with both income and population growth. Establishments per capita have a positive association with taxable income, whereas medical facilities per capita do not show a stable effect on population growth after controlling for the other variables and the dependence structure. Population density is weakly associated with population growth. Income shows strong one-period persistence, while population growth shows much less. These effects are estimated within the same spatio-temporal workflow used to generate the diagnostic maps.

6.3 From Choropleth Maps to Residual Diagnosis

The application makes the comparison concrete. The observed maps reveal regional patterns but cannot separate explained structure from remaining structure. Fitted-value maps show the component explained by covariates, the one-period lag, and the estimated dependence; residual maps show what remains. Figure 7 places the residual maps of the multivariate spatial-error model `0111` next to those of the independent spatial-error model `0ddd`. Both are visually much weaker than the observed response maps, suggesting that the spatial-error-plus-temporal structure shared by the two models captures much of the dominant pattern.

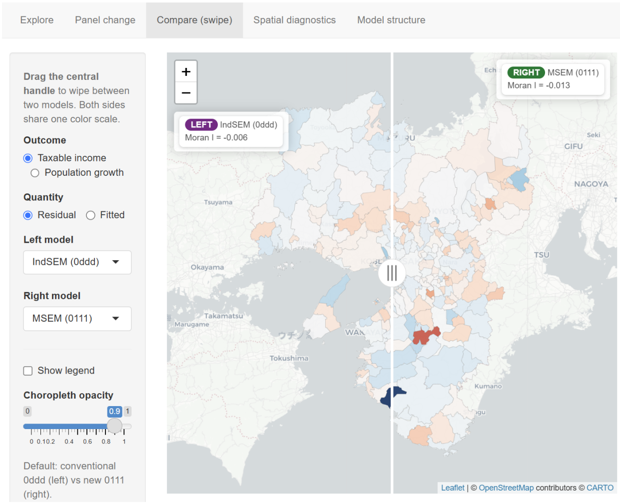

Figure 7. The app lets users compare residual maps from the independent SEM model `0ddd` (left) and the multivariate SEM model `0111` (right) by dragging the vertical divider left or right.

The application also provides a residual-difference layer (Figure 8), the $e_A - e_B$ map with $A =$ `0111` and $B =$ `0ddd`. Although the two models are an AIC near-tie, their residuals are not identical, and the map shows where the two specifications differ—information that an AIC value or a single residual map does not convey. Positive values indicate municipalities where the multivariate SEM has a larger residual than the independent SEM for the selected response, whereas negative values indicate the reverse. The analyst can thus see not only whether the multivariate model changes the residual surface, but *where* on the map that change occurs.

This map-based diagnosis constitutes the core of the transition from choropleth mapping to spatial diagnosis. The methodology is not introduced merely to improve a numerical criterion; it provides a means of decomposing regional maps into observed variation, modeled variation, and residual spatial structure, and the application enables practitioners to compare that decomposition across competing specifications. The decomposition is directly reusable in other workflows where practitioners already use choropleth maps but lack a formal, visual diagnostic step tied to model comparison.

7. Discussion: Reproducibility and Reuse

Three properties of the workflow follow from the design above. It makes a multivariate spatial econometric model reproducible end to end: the repository contains the model library, execution scripts, processed input data, spatial weights, output tables, and

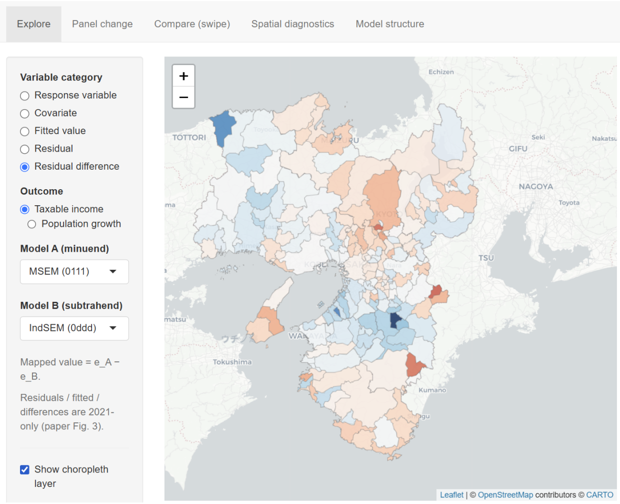

Figure 8. Residual difference for 2021 taxable income, $e_{0111,i} - e_{0ddd,i}$. Positive values indicate municipalities where the multivariate SEM has a larger residual than the independent SEM.

the application, so a reader can re-estimate every result. It turns model comparison into a visual diagnostic act, because the analyst compares the multivariate methodology against separate single-equation regressions on the map rather than only in a table. And it links statistical output back to maps, so the results are inspectable by geospatial practitioners rather than confined to a coefficient table.

The repository (Ohta et al., 2026) is organized to support these properties in practice. It provides a clear software license, a data-license statement distinguishing processed public statistics from original source data, exact commands for regenerating every table, and the application together with its pre-computed inputs. It also documents the original source URLs and contains the full set of execution scripts, so that every reported result can be reproduced and audited; standard errors, coefficient tables, and complete numerical outputs are available there.

The approach is not limited to the Kansai case study. Any dataset with multiple regional responses, a spatial weight matrix, and temporal observations can be analyzed with the same family and inspected with the same application. Possible applications include health indicators, housing markets, transport accessibility, environmental monitoring, disaster recovery, and local demographic change. In each case the diagnostic sequence is the same: observed maps, fitted maps, residual maps, residual-difference maps, and AIC/BIC rankings—and the same visual comparison reveals whether, and where, joint multivariate parameter estimation changes the spatial diagnosis.

8. Conclusions

This paper presented a reproducible open-source workflow whose central deliverable is an interactive map application for visually comparing a multivariate spatio-temporal regression methodology with separate single-equation regression models. The statistical core is the MGNST framework, which jointly represents spatial lag, spatial error, one-period temporal persistence, and cross-equation dependence and which contains independent SAR, SEM, GNS, and ordinary regressions as nested restrictions. Eleven specifications are estimated by maximum likelihood and compared primarily with AIC, with BIC reported as a parsimony-oriented sensitivity criterion.

The synthetic benchmark confirms that, when the data contain genuine coupled structure, AIC recovers the true multivariate model and ranks it above every separate single-equation regression, with the gap growing as the sample grows; BIC instead favors a more parsimonious multivariate specification. The Kansai municipal case study then demonstrates the value of the visual comparison in a more ambiguous empirical setting. There AIC and BIC select the independent spatial-error model, but the multivariate spatial-error model is a close AIC competitor and estimates statistically detectable cross-equation spatial-error dependence. The residual-difference map localizes where the two fitted residual surfaces diverge. Maps of observed values, fitted values, residuals, and residual differences therefore turn information-criterion results into a visual geospatial diagnostic workflow that is transparent, interactive, and reusable.

Future work includes improving computational efficiency for larger n and K, extending the implementation to non-Gaussian responses, adding formal residual spatial-autocorrelation tests, p-values, and local diagnostics to the application, and packaging the estimator as a more accessible open-source library.

Acknowledgements

The authors thank the developers of the open-source geospatial and statistical software used in this work, and the public agencies that provide the open statistical and geospatial data used in the case study. This study was supported by JSPS KAKENHI Grant-in-Aid for Scientific Research (Project No. 24K14857).

References

Anselin, L., 1988. *Spatial Econometrics: Methods and Models*. Kluwer Academic Publishers, Dordrecht.

Banerjee, S., Carlin, B. P., Gelfand, A. E., 2014. *Hierarchical Modeling and Analysis for Spatial Data*. 2nd edn, Chapman & Hall/CRC, Boca Raton.

Bivand, R., Piras, G., 2015. Comparing implementations of estimation methods for spatial econometrics. *J. Stat. Softw.*, 63(18), 1–36.

Bradley, J. R., Holan, S. H., Wikle, C. K., 2018. Computationally efficient multivariate spatio-temporal models for high-dimensional count-valued data. *Bayesian Anal.*, 13(1), 253–310.

Cressie, N., Wikle, C. K., 2011. *Statistics for Spatio-Temporal Data*. Wiley, Hoboken.

Elhorst, J. P., 2014. *Spatial Econometrics: From Cross-Sectional Data to Spatial Panels*. Springer, Heidelberg.

Lee, L.-F., Yu, J., 2010. Estimation of spatial autoregressive panel data models with fixed effects. *J. Econom.*, 154(2), 165–185.

LeSage, J. P., Pace, R. K., 2009. *Introduction to Spatial Econometrics*. CRC Press, Boca Raton.

Mastrantonio, G., Jona Lasinio, G., Pollice, A., Capotorti, G., Teodonio, L., Genova, G., Blasi, C., 2019. A hierarchical multivariate spatio-temporal model for clustered climate data with annual cycles. *Ann. Appl. Stat.*, 13(2), 797–823.

Ministry of Land, Infrastructure, Transport and Tourism, 2021. National Land Numerical Information download service. National Land Numerical Information (Kokudo Suuchi Jouhou). nlftp.mlit.go.jp/ksj/gml/datalist/KsjTmplt-N03-2021.html (1 June 2026).

Ohta, S., Tanaka, S., Nishii, R., 2026. Reproducible workflow and interactive diagnostic maps for multivariate spatio-temporal regression of kansai municipal data. DOI: 10.5281/zenodo.20693415. GitHub repository: github.com/s-ohta-s/mst-modeling-kansai.git (1 June 2026).

Pebesma, E., Bivand, R., 2023. *Spatial Data Science: With Applications in R*. Chapman & Hall/CRC, Boca Raton.

Statistics Bureau of Japan, 2021. e-Stat: Portal site of official statistics of Japan. Portal Site of Official Statistics of Japan (e-Stat). www.e-stat.go.jp (25 March 2026).