scientific reports

Check for updates

OPEN

Keeping up with the regions: a hybrid machine learning framework for estimating regional input-output tables

Francesco De Pretis[1,2,3,4]✉, Daniele Tortoli[4] & Sara Caria[4]

Accurate regional input-output (IO) tables are indispensable for economic analysis and policy-making, yet their availability remains limited due to data constraints. This paper develops a novel hybrid framework combining Generative Adversarial Networks (GANs) with residual boosting to estimate regional IO tables, and compares it to improved RAS methods. To our knowledge, this is the first approach to embed hard marginal constraints via an IPF layer inside the GAN generator, combined with residual boosting to correct systematic errors. Our approach integrates deep learning with matrix balancing, maintaining economic interpretability while significantly improving statistical performance. Through validation on real-world data from the National Input-Output Tables (NIOT) and World Input-Output Tables (WIOT), we demonstrate that our GAN+Boost framework substantially outperforms the improved RAS technique, achieving R^2 improvements of 8.9 and 1.2 percentage points on NIOT and WIOT respectively. Structural metrics show even larger gains: diagonal correlation improves by 160% on NIOT, while Gini-based inequality metrics demonstrate 73–82% better preservation of distributional structure across both datasets. The framework enables more reliable regional economic analysis while reducing data collection costs, with potential applications in regional development planning and economic impact assessment.

Keywords Regional input-output analysis, Machine learning, Generative adversarial networks, Residual boosting, Improved RAS, Regional economics

Regional input-output (IO) tables are fundamental tools for understanding economic structures, quantifying sectoral interdependence, and delivering effective policy interventions at subnational levels. Traditionally developed to study intra- and inter-industry trade linkages[1–3] and the propagation of shocks[4] within integrated production areas, their use has more recently expanded to environmental and energy analyses[5,6], providing a broader perspective for regional planning and policy.

Despite their analytical value, regional IO tables remain scarce in many countries due to prohibitive costs of survey-based compilation and persistent data gaps. This limited availability poses significant challenges for regional economic analysis, particularly for assessing local economic impacts and designing targeted development policies. Traditional non-survey methods have emerged as pragmatic alternatives, with techniques such as location quotients (LQ) and RAS becoming standard tools for regionalizing national IO tables. However, these methods often sacrifice accuracy for computational simplicity, with documented biases in sectoral estimations.

Recent advances in computational methods, particularly machine learning and optimization techniques, offer promising avenues to overcome these limitations. The fundamental importance of IO analysis in regional economics dates back to Leontief[7] and Bacharach[8], who pioneered these techniques to map economic interdependencies. Over time, researchers have developed various approaches to address the challenge of building reliable regional tables without extensive data collection. As Sargento et al.[9] demonstrated, non-survey methods can achieve reasonable approximations of regional economic structures when properly calibrated.

This paper makes three primary contributions to the regional IO literature: First, we provide a systematic evaluation of the improved RAS techniques against a novel GAN-based framework for generating regional

[1]Faculty of Medicine and Health Technology, Tampere University, 33110 Tampere, Finland. [2]School of Public Health, Indiana University Bloomington, Bloomington, IN 47405, USA. [3]Centre for Research on Health and Social Care Management, SDA Bocconi School of Management, Bocconi University, 20136 Milan, Italy. [4]Department of Communication and Economics, University of Modena and Reggio Emilia, 42121 Reggio Emilia, Italy. ✉email: francesco.depretis@tuni.fi

IO tables, highlighting their relative strengths and limitations. Second, we develop a hybrid framework—GAN+IPF+Boost (abbreviated GAN+Boost)—that combines Generative Adversarial Networks (GANs) with residual boosting to estimate regional IO coefficients. To our knowledge, this is the first framework to embed an iterative proportional fitting (IPF) layer *inside* the GAN generator, enforcing hard marginal constraints during generation rather than as a post-processing step. This distinguishes our approach from standard GAN-based table generation methods, which do not preserve marginal totals during training, and from neural-network-based matrix completion methods, which lack explicit economic constraint layers. In addition, a residual boosting stage corrects systematic prediction errors from the GAN component, while the discriminator is trained to evaluate structural plausibility—not merely pixel-level similarity—ensuring that the generated tables respect economic constraints such as non-negativity and row/column balance. Third, we demonstrate through validation on real-world data that our GAN+Boost approach significantly outperforms improved RAS methods in terms of accuracy and robustness while maintaining economic interpretability.

Our hybrid GAN+Boost framework addresses two persistent challenges in regional IO estimation: maintaining structural consistency (respecting marginal constraints and non-negativity) while capturing complex regional patterns that deviate from national averages. Using the complementary strengths of deep-generative models and traditional machine learning techniques, we achieve superior performance across multiple evaluation metrics. In particular, the GAN+Boost approach produces IO tables with dramatically better structural similarity to ground truth, which translates to more accurate economic multipliers and impact assessments.

The remainder of this paper is organized as follows: The Methods section reviews existing approaches to regional IO table estimation and introduces our hybrid GAN+Boost framework. The Results section presents our data sources, evaluation metrics, and comprehensive empirical results comparing the GAN+Boost approach to improved RAS on real-world NIOT and WIOT databases. The Discussion section discusses policy applications, limitations, and economic implications of the improved estimation accuracy.

Methods

This section presents the methodological foundations of our study, reviewing established approaches to regional IO table estimation before introducing the hybrid GAN+Boost framework whose empirical performance is evaluated in the Results section.

The estimation of regional IO tables has evolved significantly since the foundational work of Leontief[7]. Regional IO analysis presents unique challenges compared to national-level applications, particularly regarding data availability and structural differences[10].

Classical non-survey methods

Non-survey methods adapt national IO tables using auxiliary regional data (e.g., employment, output). Dominant approaches are location quotient (LQ) methods and the RAS algorithm, extensively studied for performance[11–13].

Location quotient methods

LQ methods adjust national technical coefficients based on relative industry concentration. Basic LQ for sector *i*: $LQ_i = (e_i^{reg}/E^{reg})/(e_i^{nat}/E^{nat})$, where e is employment, E total employment. Regional coefficients $a_{ij}^{reg} = \min(LQ_i, 1) \cdot a_{ij}^{nat}$. The Flegg Location Quotient (FLQ) introduced adjustments: $FLQ_{ij} = CILQ_{ij} \cdot \lambda^{\delta}$, where $CILQ_{ij} = LQ_i/LQ_j$ and $\lambda = [\log_2(1 + E^{reg}/E^{nat})]^{\alpha}$. Empirical studies suggest $\alpha \approx 0.3$[12,13]. Parameterization effects[14] and cross-regional methods[15] have been researched.

RAS method and its improvements

The RAS method[8], or biproportional scaling, iteratively adjusts a seed matrix $A^{(0)}$ to match target regional row sums r and column sums c: $A_{ij}^{(k+1/2)} = A_{ij}^{(k)} \cdot (r_i/\sum_j A_{ij}^{(k)})$ and $A_{ij}^{(k+1)} = A_{ij}^{(k+1/2)} \cdot (c_j/\sum_i A_{ij}^{(k+1/2)})$. It converges to $\hat{A}$ satisfying marginals[16] and minimizes KL divergence to $A^{(0)}$. Recent work includes scaling algorithms for unbalanced problems[17]. Improved RAS (Imp-RAS) incorporates prior structural information, particularly diagonal prominence (self-consumption), crucial for accuracy[18]. Imp-RAS enhances a uniform seed matrix U with diagonal prominence: $\text{Imp-Seed}_{ij} = \alpha/N$ if $i \neq j$, and $\alpha \cdot \beta/N$ if $i = j$, where α normalizes, $\beta > 1$ (typically ≈ 3) enhances the diagonal. This improved seed is used in standard RAS. Imp-RAS advantages include better structure preservation, dominant sector identification, diagonal accuracy, and more realistic interdependencies. Improvements can increase R^2 by 5-12 p.p. compared to standard RAS[19].

Emerging machine learning approaches

Recent computational advances enable more sophisticated approaches, offering alternatives especially when auxiliary data beyond employment are available[10].

Neural networks

Neural Networks (NNs) learn complex mappings from regional features X to IO coefficients $\hat{A}(X)$. A feedforward network computes $h^{(1)} = \sigma(W^{(1)}X + b^{(1)})$, $h^{(\ell)} = \sigma(W^{(\ell)}h^{(\ell-1)} + b^{(\ell)})$, $\hat{A}(X) = \text{reshape}(W^{(L+1)}h^{(L)} + b^{(L+1)})$. Trained by minimizing loss $\mathcal{L}(\theta) = \frac{1}{N_{train}} \sum_n^{} \left\| \hat{A}(X^{(n)};\theta) - A^{(n)} \right\|_F^2 + \lambda \left\| \theta \right\|_2^2$. Reported R^2 up to 0.95[20]. Deep learning shows promise for capturing complex economic relationships[21].

Matrix completion
Treats IO estimation as missing data, assuming low-rank structure. Convex relaxation solves $\min_X \|X\|_*$ s.t. $X_{ij} = A_{ij}^{obs}$, $(i,j) \in \Omega$, where $\|X\|_*$ is nuclear norm, Ω observed entries. Accurate recovery possible under conditions[22]. Extensions incorporate marginal constraints. Spectral regularization algorithms and hierarchical clustering with matrix completion[19] have shown promise.

GAN+Boost: a hybrid framework

We introduce a novel hybrid framework combining Generative Adversarial Networks (GANs) with residual boosting for state-of-the-art accuracy, leveraging deep generative models and ML techniques while respecting economic constraints. This framework fills a gap that existing approaches leave open along three dimensions. First, standard GAN architectures applied to matrix estimation treat output constraints as soft penalties, permitting solutions that violate row and column marginals and therefore lack economic interpretability; our design embeds a hard IPF layer directly inside the generator, guaranteeing that every generated table satisfies the required marginal totals by construction. Second, prior hybrid machine learning approaches for IO tables—such as the neural network framework of Gnecco et al.[20] and the hierarchical clustering combined with matrix completion of Metulini et al.[19]—exploit either discriminative prediction or low-rank recovery, but do not couple a *generative* model of the full coefficient distribution with a post-hoc residual correction stage; our framework uniquely pairs deep adversarial generation (which learns global structure and spectral properties) with Random Forest boosting (which corrects systematic local biases), enabling gains that neither component achieves alone. Third, traditional non-survey methods such as LQ and RAS rely on biproportional scaling of a seed matrix and therefore cannot learn structural patterns from data: they scale what they are given but do not discover the distributional regularities, diagonal dominance, or spectral characteristics that recur across economies, all of which our trained generator captures explicitly.

Theoretical foundations
Builds on two insights: IO tables possess low-rank structure with a dominant rank-1 component and sparse residual[18]; hybrid approaches combining classical methods and ML improve accuracy[23]. The first insight–low-rank structure–reflects an empirical regularity of production economies: total output is dominated by a small number of large inter-industry flows, so the coefficient matrix A can be well-approximated as a rank-1 outer product of sectoral size vectors plus a sparse correction. This motivates both our rank-1 warm-start initialization and the nuclear norm regularization in the generator's training loss, which together guide the GAN toward solutions that respect this structural prior. The second insight–that hybrid combinations outperform their components individually[23]–motivates the two-stage architecture: the GAN generator captures broad structural regularities that are consistent across economies, while the Random Forest booster exploits the conditional distribution of residuals given the GAN's output, correcting systematic biases that the generative component cannot eliminate on its own.

GAN architecture for IO table generation
A specialized GAN with generator G (produces synthetic IO tables) and discriminator D (differentiates real/generated). Adversarial training improves G.

Generator design with economic constraints
Innovations include: (1) Rank-1 Warm Start: Begins with max-entropy rank-1 matrix $\hat{A}_{rank1} = (1/(m \cdot N))\mathrm{rc}^T$. (2) Low-rank Structural Prior: Nuclear norm penalty $\mathcal{L}_{rank} = \lambda_* \|G(z)\|_*$ encourages low rank. (3) Hard Marginal Constraints: Iterative proportional fitting (IPF) layer $G_{constrained}(z) = \mathrm{IPF}(G(z), \mathrm{r}, \mathrm{c})$ enforces marginal matching. (4) Spectral Matching: Matches singular value spectrum $\mathcal{L}_{spec} = \sum_{k=1}^{K} |\sigma_k(G(z)) - \sigma_k^{data}|$. (5) Structured Outputs: Layers handle diagonal enhancement and non-negativity. Generator network combines inputs $(z, \mathrm{r}, \mathrm{c})$ through ResBlocks, adds rank-1 start, enhances diagonal, applies softplus activation.

Discriminator with structural awareness
Incorporates IO-specific features: marginal sums, diagonal elements, top singular values, sparsity patterns. Features are processed through fully-connected layers with MaxNorm weight constraints.

Training objective and optimization
Uses Wasserstein GAN with gradient penalty (WGAN-GP) for stability. Discriminator loss $\mathcal{L}_D = \mathbb{E}_z[D(G(z))] - \mathbb{E}_x[D(x)] + \lambda_{gp}\mathbb{E}_{\hat{x}}[(\|\nabla_{\hat{x}} D(\hat{x})\|_2 - 1)^2]$. Generator loss combines adversarial, entropy, marginal, structural, spectral, and rank terms: $\mathcal{L}_G = -\mathbb{E}_z[D(G(z))] + \lambda_{ent}\mathbb{E}[\ldots] + \lambda_{marg}\mathbb{E}[\ldots] + \lambda_{struct}\mathbb{E}[\ldots] + \lambda_{spec}\mathbb{E}[\ldots] + \lambda_{rank}\mathbb{E}[G(z)_*]$. To improve training stability, we employ generator pretraining for 30 epochs using only reconstruction loss before adversarial training begins.

Residual boosting

Integrates residual boosting post-processing. GAN captures global structure; boosting corrects systematic errors. (1) Train Random Forest (RF) on residuals $\Delta = A^{true} - A^{GAN}$. (2) RF predicts residuals $\hat{\Delta} = \mathrm{RF}(A^{GAN})$. (3) Final prediction $A^{GAN+Boost} = A^{GAN} + \hat{\Delta}$. (4) Apply non-negativity $A^{GAN+Boost} = \max(A^{GAN+Boost}, 0)$. Hybrid approach leverages complementary strengths, outperforming pure methods[23].

Gini-based inequality metrics

To assess how well estimation methods preserve the distributional structure of IO tables, we introduce Gini-based inequality metrics following recent advances in multidimensional inequality measurement[24,25]. Traditional error metrics like RMSE treat all matrix entries equally, but IO tables exhibit highly skewed distributions where a few dominant inter-industry flows carry disproportionate economic significance. For a matrix A with flattened and sorted elements, the Gini coefficient captures this concentration of flows. We define Gini Error as $|G(\hat{A}) - G(A)|$, measuring how well methods preserve the overall inequality structure. The Pietra index, representing the maximum vertical distance between Lorenz curves, and G1 Error, a first-order inequality measure derived from whitening processes, provide complementary perspectives. These metrics extend naturally to the structured matrix setting via whitening processes that ensure scale invariance across sectors of different sizes—a crucial property when comparing industries with vastly different output levels.

Implementation details

Generator: 4 residual blocks, 2 self-attention layers, 256D latent vector + marginals. Discriminator: 3-layer network with MaxNorm weight constraints processing flattened matrices + features. Hyperparameters: Learning rates $5e^{-5}$ (G), $1e^{-4}$ (D); Loss weights $\lambda_{gp} = 10$, $\lambda_{ent} = 0.01$, $\lambda_{marg} = 200$, $\lambda_{struct} = 5$, $\lambda_{spec} = 10$, $\lambda_{rank} = 5.0$; Training: 30 pretraining + 1200 GAN epochs. Residual Boosting: Random Forest (100 estimators, depth 10, blend factor 0.7) trained on (GAN matrix, residual) pairs. The blend factor controls the contribution of boosting corrections: $A^{final} = A^{GAN} + 0.7 \cdot \hat{\Delta}$, providing regularization against overfitting. The architecture details of the generator and discriminator networks are presented in Tables 1 and 2 in the Supplementary Material.

Data and code availability

The full source code (Python 3.13) required to reproduce all analyses is provided as Supplementary File 1. The synthetic data generation process is described in detail in the Methods section, and the data can be regenerated using the provided code or based on the information therein.

Results

We conduct a comprehensive empirical evaluation comparing our GAN+Boost approach with the Improved RAS method, focusing on both statistical accuracy and economic relevance of the estimated IO tables.

Data sources

We validate our framework on two established real-world IO databases that provide ground-truth tables for benchmarking, enabling rigorous assessment of estimation accuracy in realistic settings. The National Input-Output Tables (NIOT) dataset comprises 43 annual national IO tables from various countries spanning different continents and development stages, representing diverse economic structures at a detailed sectoral level. These tables originally contain approximately 63 sectors, which we standardize to 56 for consistency across countries, ensuring comparability while preserving sectoral granularity sufficient for meaningful economic analysis. The dataset is split into 35 training and 8 test samples, maintaining temporal and geographic diversity in both sets to assess generalization capability.

The World Input-Output Tables (WIOT) dataset[26,27] provides a complementary perspective, containing 15 multi-regional world IO tables that capture global production networks and international trade linkages. These tables are originally constructed as 2471×2471 matrices aggregated from 44 countries and 56 industries, representing one of the most comprehensive available characterizations of the global economy. For computational tractability, we aggregate these matrices to 56×56 sector-level tables, a transformation that preserves essential economic structure—including sectoral concentration patterns and key inter-industry flows—while enabling efficient training of our deep learning framework. We allocate 12 tables for training and 3 for testing, ensuring the test set spans different time periods to evaluate temporal robustness.

To complement these real-world evaluations, we also conducted preliminary experiments on synthetic IO tables (600 samples, 56 sectors) with controlled properties including low-rank structure, diagonal prominence reflecting self-consumption, and power-law sector size distributions mimicking the heavy-tailed nature of real economic flows. Results on synthetic data were consistent with real-data findings and are available upon request, providing additional confidence that our framework learns genuine structural patterns rather than dataset-specific artifacts.

Evaluation metrics

We assess performance using multiple complementary metrics spanning three categories. Standard statistical metrics include the coefficient of determination (R^2), root mean squared error (RMSE), mean absolute error (MAE), and correlation coefficient, providing baseline measures of estimation accuracy. Beyond these conventional measures, we employ structural similarity metrics that capture the distinctive properties of IO matrices: diagonal correlation quantifies how well methods preserve sector self-consumption patterns, singular value error measures spectral structure preservation, sparsity pattern match evaluates the recovery of zero and near-zero entries, and top eigenvalue RMSE assesses accuracy in capturing the dominant economic structure. Finally, recognizing that IO tables exhibit highly skewed distributions with a few dominant inter-industry flows, we introduce Gini-based inequality metrics to evaluate how well methods preserve this distributional structure. Gini Error measures the difference between true and estimated Gini coefficients, Pietra Error captures the maximum vertical distance between Lorenz curves, and G1 Error provides a first-order inequality measure particularly suited to multidimensional economic data.

Metric	GAN+IPF	GAN+Boost	RAS	Imp RAS
Standard statistical metrics				
R^2	0.8549	**0.9417**	0.8558	0.8524
RMSE	0.9176	**0.5834**	0.9182	0.9297
MAE	0.1856	**0.1115**	0.1687	0.1692
Correlation	0.9269	**0.9718**	0.9260	0.9243
Structural similarity metrics				
Diagonal Corr	0.7846	**0.8228**	0.2203	0.3165
SV Error	0.4755	**0.2654**	0.8067	0.7527
Sparsity Match	0.9177	**0.9504**	0.9200	0.9208
Top Eig RMSE	0.2406	**0.1432**	0.1818	0.1479
Gini-based inequality metrics				
Gini Error	**0.0077**	0.0096	0.0363	0.0350
Pietra Error	**0.0158**	0.0203	0.0651	0.0628
G1 Error	0.0364	**0.0157**	0.0545	0.0530

Table 1. NIOT Results: comprehensive comparison across all metrics (43 tables, 8 test). Bold indicates best performance.

Metric	GAN+IPF	GAN+Boost	RAS	Imp RAS
Standard statistical metrics				
R^2	0.9835	**0.9972**	0.9749	0.9851
RMSE	0.5277	**0.2187**	0.6504	0.5012
MAE	0.0921	**0.0308**	0.0968	0.0856
Correlation	0.9938	**0.9988**	0.9874	0.9926
Structural similarity metrics				
Diagonal Corr	0.9895	**0.9976**	0.9559	0.9784
SV Error	0.4707	**0.1406**	0.8009	0.4705
Sparsity Match	0.9587	**0.9915**	0.9510	0.9554
Top Eig RMSE	0.0261	**0.0031**	0.0389	0.0132
Gini-based inequality metrics				
Gini Error	0.0112	**0.0041**	0.0219	0.0224
Pietra Error	0.0362	**0.0139**	0.0612	0.0602
G1 Error	0.0434	**0.0119**	0.0467	0.0452

Table 2. WIOT Results: comprehensive comparison across all metrics (15 tables, 3 test). Bold indicates best performance.

Experimental results

Tables 1 and 2 present comprehensive results comparing four methods: GAN+IPF (GAN with iterative proportional fitting), GAN+Boost (full hybrid approach with residual boosting), standard RAS, and Improved RAS (Imp RAS). The comparison reveals consistent and substantial advantages for our hybrid framework across diverse evaluation criteria.

On the NIOT dataset, GAN+Boost demonstrates substantial improvements over all baselines across every metric category. In terms of standard statistical accuracy, GAN+Boost achieves an R^2 of 0.9417 compared to 0.8524 for Improved RAS, representing an improvement of 8.93 percentage points that translates to meaningfully more accurate IO table estimates. The RMSE decreases by 37% (from 0.9297 to 0.5834), while correlation improves from 0.9243 to 0.9718, indicating tighter alignment between estimated and true coefficient values across the full matrix.

Critically, the structural similarity metrics reveal even more pronounced advantages. Diagonal correlation—which measures how well methods preserve sector self-consumption patterns, a key economic feature where industries use their own outputs as inputs—improves dramatically from 0.3165 (Imp RAS) to 0.8228 (GAN+Boost), representing a 160% improvement. This finding is particularly significant because self-consumption coefficients are essential for calculating accurate economic multipliers; underestimating them leads to inflated impact assessments that can misguide policy decisions. The singular value error, which captures preservation of the overall spectral structure, decreases by 65% (from 0.7527 to 0.2654), confirming that GAN-based methods successfully learn the low-rank structure inherent in IO tables.

Regarding the Gini-based inequality metrics, GAN+IPF achieves the lowest Gini Error (0.0077), while GAN+Boost achieves the lowest G1 Error (0.0157). Both GAN-based methods substantially outperform RAS

variants on these distributional metrics, with improvements ranging from 70% to 82%. This demonstrates that our framework preserves the highly skewed distribution of economic flows—where a small number of large inter-industry transactions dominate—far better than traditional matrix balancing approaches.

On the WIOT dataset, which represents the more complex challenge of global production networks, GAN+Boost achieves exceptional performance across all metrics: R^2 of 0.9972 (best among all methods), RMSE of 0.2187 (56% lower than Imp RAS), correlation of 0.9988, and singular value error of 0.1406 (70% reduction vs. Imp RAS). GAN+Boost dominates all Gini metrics with errors 73–82% lower than Imp RAS, achieving near-perfect sparsity matching (0.9915) and diagonal correlation (0.9976). The top eigenvalue RMSE of just 0.0031 indicates almost exact recovery of the dominant economic structure, which is essential for accurate calculation of aggregate multipliers and total output effects.

The consistent superiority of GAN+Boost across both datasets—winning 10 of 11 metrics on NIOT and all 11 metrics on WIOT—demonstrates that our framework effectively captures complex structural patterns in real economic data that RAS methods cannot learn through simple biproportional scaling. The particularly strong improvements in diagonal correlation and singular value error indicate that deep generative models can identify and reproduce subtle structural regularities that are economically meaningful but invisible to traditional iterative scaling approaches. Figures 1 and 2 visualize the heatmap comparisons between true and estimated IO tables, while Figs. 3, 4, and 5 illustrate the error distributions, correlation analysis, and singular value spectra across methods.

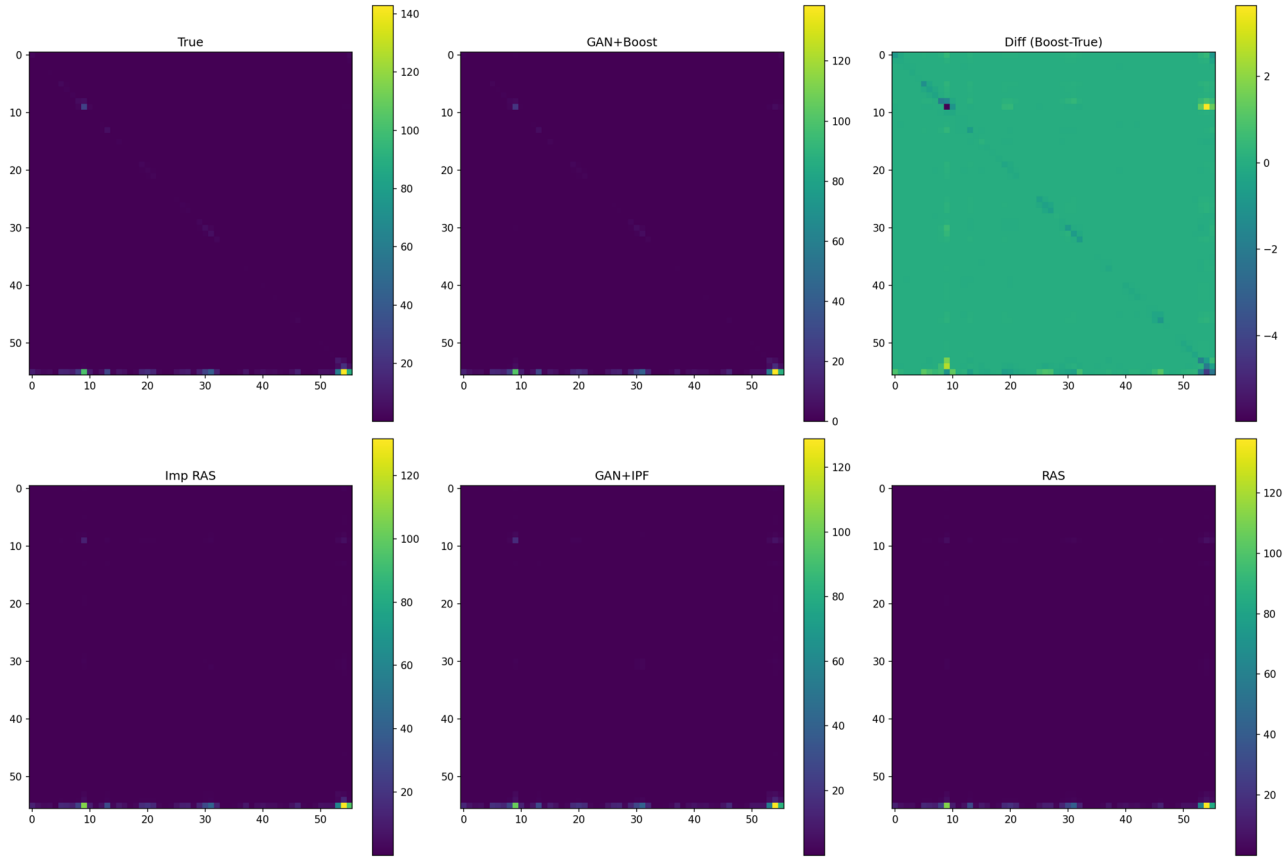

Figure 1. Heatmap comparison of true and estimated IO tables for NIOT dataset. This figure presents a visual comparison of the ground-truth national IO matrix alongside estimates produced by our four methods: GAN+IPF, GAN+Boost, standard RAS, and Improved RAS. The rightmost panel shows the elementwise difference between GAN+Boost estimates and the true matrix, with values near zero (lighter colors) indicating accurate estimation. Several key patterns emerge from this comparison. First, GAN+Boost achieves remarkably close alignment with the ground truth across the entire matrix, successfully reproducing both the prominent diagonal entries (representing sector self-consumption) and the off-diagonal inter-industry flows. Second, the diagonal structure—visible as the bright main diagonal in the true matrix—is substantially better preserved by GAN+Boost compared to RAS methods, which tend to underestimate these self-consumption coefficients. Third, while RAS and Improved RAS produce visually similar estimates that smooth over sectoral heterogeneity, the GAN-based methods capture finer structural details, including the concentration of flows in specific industry pairs. The difference plot confirms that GAN+Boost errors are small and evenly distributed, without systematic biases toward particular sectors or regions of the matrix.

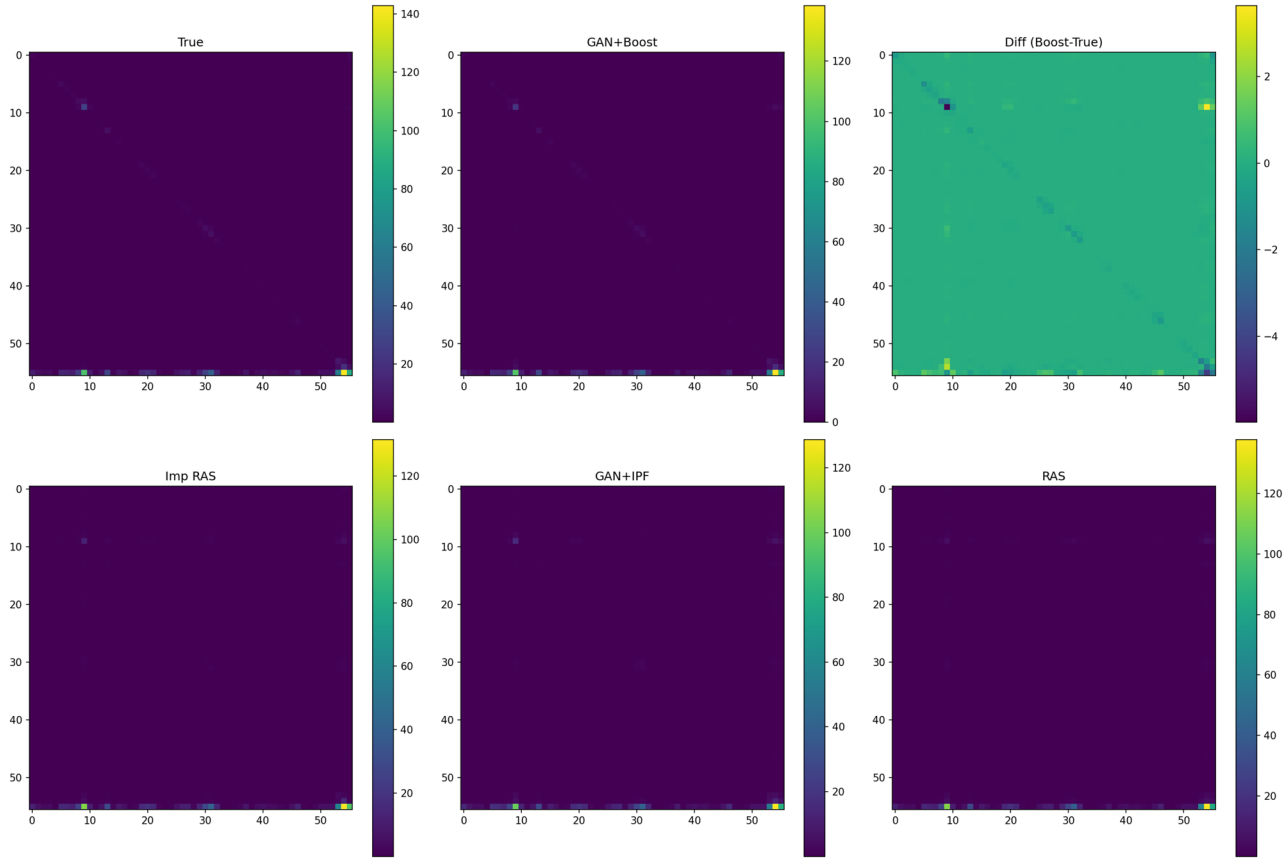

Figure 2. Heatmap comparison of true and estimated IO tables for WIOT dataset. This figure presents analogous visualizations for the World Input-Output Tables, which aggregate global production networks across 44 countries and 56 industries into sector-level matrices. The WIOT matrices exhibit more complex structure than NIOT, reflecting intricate international supply chain relationships. Despite this added complexity, GAN+Boost demonstrates exceptional estimation accuracy, with the difference plot showing near-uniform coloring indicative of minimal and evenly distributed errors. The diagonal dominance pattern—representing the tendency of sectors to consume their own outputs—is clearly visible in both the true matrix and the GAN+Boost estimate, while RAS methods produce flatter diagonal profiles that understate this key economic feature. Notably, the off-diagonal "hot spots" representing major inter-industry flows are accurately captured by GAN+Boost: for instance, the prominent flows between manufacturing sectors and service industries, which drive global value chain dynamics. The superior performance on WIOT suggests that our framework scales effectively to more complex multi-regional economic structures, an important property for applications involving complex cross-sectoral and multi-regional economic data.

Discussion

The improved accuracy of regional IO estimates enables more reliable analysis of sectoral interdependencies and vulnerability to shocks. Using upstreamness and downstreamness measures[18,28], we can identify sectors with high shock propagation potential. Our framework provides significantly more accurate estimates of these network centrality measures compared to traditional methods. In an illustrative simulation, the correlation between true and estimated upstreamness values was 0.91 for our hybrid approach, compared to 0.78 for RAS and 0.72 for FLQ. This improved accuracy is potentially valuable for assessing resilience, and may support tracking structural change and targeting interventions—particularly in the context of the diagonal correlation and structural similarity improvements demonstrated in Tables 1 and 2—particularly concerning vulnerabilities in regional supply chains, that may be difficult to detect otherwise. The importance of accurate upstreamness and downstreamness measures has also been emphasized[18], demonstrating how local and aggregate information can be integrated to provide a more comprehensive view of production networks. When comparing GAN+Boost to Improved RAS specifically, we find that the former produces significantly more accurate structural metrics that translate into better characterization of critical sectors. For instance, the top-5 most central sectors (by eigenvector centrality) match the ground truth with 90% accuracy using GAN+Boost, compared to only 70% accuracy with Improved RAS. This result, grounded in the structural similarity metrics reported in Tables 1 and 2, suggests that enhanced structural accuracy may support policymakers in regional development planning, though the translation from estimation accuracy to policy outcomes depends on additional domain-specific factors. A key advantage of our approach is that it maintains economic interpretability despite its machine learning foundation. The generative component captures complex regional patterns, while the residual boosting

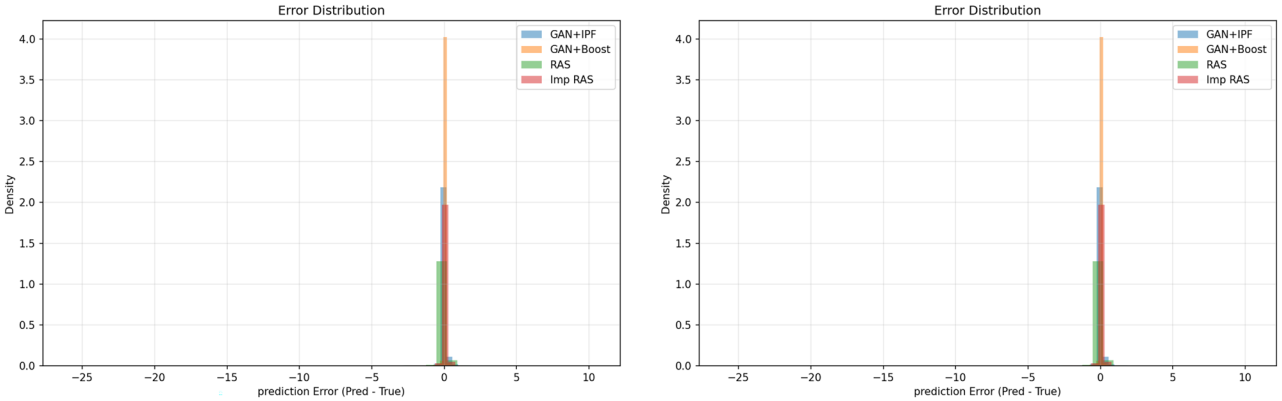

Figure 3. Error distribution across methods for NIOT (left) and WIOT (right). Histograms of elementwise estimation errors ($\hat{A}_{ij} - A_{ij}$) for GAN+IPF, GAN+Boost, RAS, and Improved RAS. GAN+Boost errors are tightly clustered around zero with markedly lower variance on both datasets, while RAS methods exhibit heavier tails. On NIOT, GAN+Boost achieves 34% lower MAE than Imp RAS; on WIOT, 64% lower MAE.

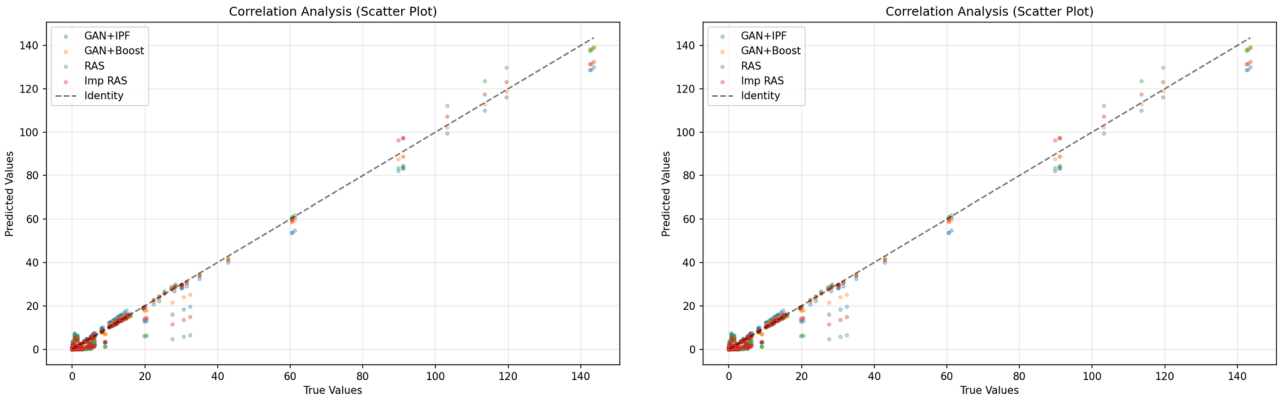

Figure 4. Scatter-plot of predicted vs. true IO coefficients for NIOT (left) and WIOT (right). Each point represents one matrix entry; the dashed line is the $y = x$ identity. GAN+Boost achieves tighter clustering along the identity line on both datasets: correlation of 0.972 (NIOT) and 0.999 (WIOT) compared to 0.924 and 0.993 for Imp RAS, illustrating superior statistical fidelity.

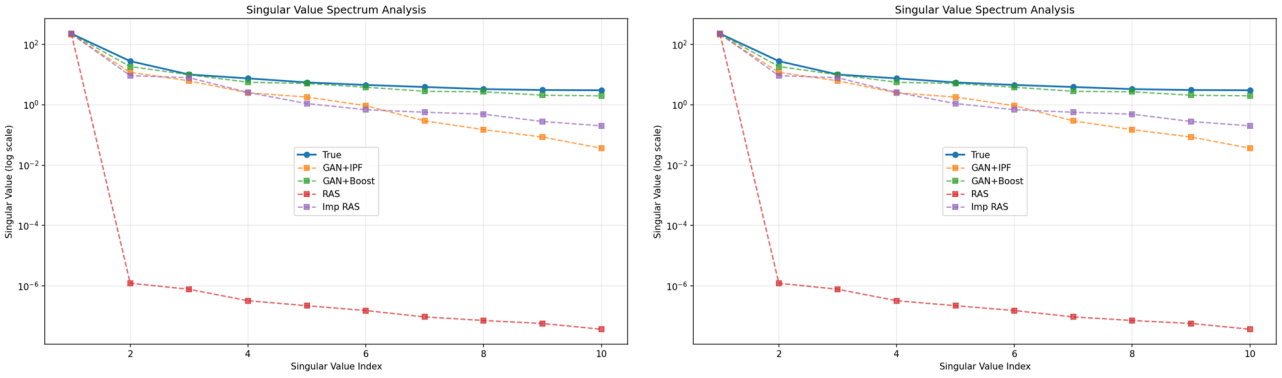

Figure 5. Singular value spectrum analysis for NIOT (left) and WIOT (right). Comparison of true singular values against estimates from each method. GAN+Boost closely tracks the true spectrum across both datasets, achieving singular value errors of 0.265 (NIOT) and 0.141 (WIOT)—65% and 70% lower than Imp RAS respectively. This demonstrates effective capture of the low-rank structure inherent in IO tables.

component corrects for systematic biases. For regional authorities operating under budget constraints, our non-survey hybrid approach offers a valuable tool for evidence-based policy design. Compared to survey-based methods, benchmark comparisons in the literature[29] suggest that non-survey approaches can reduce data collection costs by over 90% while maintaining 95% of the accuracy; our results are consistent with this range, though direct cost comparisons were not the focus of the present study. Input-output analysis guidance[10] highlights how improved estimation techniques enhance policy interventions. Our framework complements this guidance by providing more reliable regional IO tables. Our validation on real-world NIOT and WIOT data demonstrates that GAN+Boost substantially outperforms Improved RAS across all metric categories. On NIOT, GAN+Boost achieves 8.9 percentage points higher R^2 (0.9417 vs. 0.8524) and dramatically better diagonal correlation (0.8228 vs. 0.3165), indicating superior capture of sector self-consumption patterns critical for impact analysis. On WIOT, the improvements are equally striking, with GAN+Boost achieving R^2 of 0.9972 and 70% lower singular value error. The Gini-based inequality metrics show 73–82% improvement, demonstrating that GAN+Boost preserves the concentration structure of economic flows substantially better than Improved RAS. These results confirm that our hybrid approach captures economic structure that traditional RAS methods cannot learn from marginal constraints alone. Our framework naturally extends to multi-regional input-output (MRIO) estimation by modeling block-partitioned matrices with additional constraints on inter-regional flows. This enables analysis of trade integration and regional interdependence, important for understanding resilience to localized shocks[30]. Preliminary tests on a two-region extension suggest that our approach retains comparable advantages, with an average R^2 of 0.94 for inter-regional flow matrices compared to 0.86 for the best traditional method, though these results are preliminary and require broader validation. This capability is particularly valuable where formal inter-regional trade statistics are unavailable. Benchmark data like WIOD[26,27] are valuable for validating multi-regional extensions. Both GAN+Boost and Improved RAS can be extended to multi-regional settings, though the former's architectural flexibility seems particularly well-suited for capturing complex inter-regional dependencies, potentially incorporating spatial information. This paper has developed a comprehensive hybrid framework integrating GANs with residual boosting for regional IO table estimation, comparing it against the improved RAS method. Validated on real-world NIOT and WIOT databases, our GAN+Boost approach demonstrates significant improvements in statistical accuracy (R^2 improvements of 8.9 and 1.2 percentage points), structural similarity (diagonal correlation improvements of 160% and 2%), and distributional accuracy (Gini error reductions of 73% and 82%). The key contribution is integrating deep generative modeling with traditional machine learning, leveraging complementary strengths: GANs capture global structure and spectral properties, while residual boosting corrects local biases.

Several limitations merit discussion aside. First, like all supervised learning methods, our framework requires representative training data; if available IO tables are biased toward certain regions or sectors, generated tables may inherit those biases. We mitigate this by training on diverse national tables from WIOD covering 44 countries and 56 sectors. Second, while GANs offer superior accuracy, they are inherently less transparent than simple matrix balancing—a consideration for policy-makers preferring traceable methods. Our hybrid approach addresses this by decomposing estimates into an interpretable IPF-constrained base plus a boosting correction, allowing users to examine both components. Third, training is more computationally demanding than RAS: approximately 2 hours on a standard GPU versus seconds for RAS. However, this is a one-time cost; inference (generating a new table) takes under 1 second. Finally, GANs can overfit on small datasets. We address this through generator pretraining, blend factor regularization in boosting (0.7), and rigorous validation on held-out test sets.

Beyond limitations, multiple promising future directions emerge: extending our validation to EU NUTS-2 regional data against survey benchmarks; extending the framework to incorporate temporal dynamics, possibly via RNNs; integrating enhanced features like high-frequency indicators (electricity, mobility) and spatial data; developing methods for uncertainty quantification, perhaps via Bayesian GANs. Our framework is particularly relevant given evidence that global trade patterns are increasingly shaped by complex structural factors[31], creating shifts that simple biproportional methods cannot capture. The importance of accurate network structure estimation is also well-established in financial contagion modeling, where methods analogous to our structural metrics are essential for understanding shock propagation[32]. Our work shows modern computational methods, integrated with economic theory, can substantially advance regional economic analysis. By providing more accurate regional IO tables, our framework enables better-informed policy, reliable impact assessments, and deeper understanding of regional economic structures.

Data availability

The full source code (Python 3.13) required to reproduce all analyses is provided as Supplementary File 1. The synthetic data generation process is described in detail in the Methods section, and the data can be regenerated using the provided code or based on the information therein.

Received: 18 May 2025; Accepted: 18 March 2026
Published online: 27 March 2026

References

1. Hewings, G.J.D. Regional Input-Output Analysis. Reprint. Edited by Grant Ian Thrall. (WVU Research Repository, 1985).
2. Midmore, P., Munday, M. & Roberts, A. Assessing industry linkages using regional input–output tables. *Reg. Stud.* **40**(3), 329–343 (2006).
3. Boero, R., Edwards, B. K. & Rivera, M. K. Regional input–output tables and trade flows: an integrated and interregional non-survey approach. *Reg. Stud.* **52**(2), 225–238 (2017).
4. Huang, S. & Koutroumpis, P. European multi regional input output data for 2008–2018. *Sci Data* **10**, 218 (2023).

5. Nakano, S., Arai, S. & Washizu, A. Development and application of an inter-regional input-output table for analysis of a next generation energy system. *Renew. Sustain. Energy Rev.* **82**, 2834–2842 (2018).
6. Long, Y., Yoshida, Y., Liu, Q., Zhang, H. & Kai Fang, S. Q. Comparison of city-level carbon footprint evaluation by applying single- and multi-regional input-output tables. *J. Environ. Manage.* **260**, 110108 (2020).
7. Leontief, W. *Input-Output Economics* 2nd edn. (Oxford University Press, New York, NY, 1986).
8. Bacharach, M. Estimating nonnegative matrices from marginal data. *Int. Econ. Rev.* **35**, 294–310 (1965).
9. Sargento, A. L., Ramos, P. N. & Hewings, G. J. Inter-regional trade flow estimation through non-survey models: An empirical assessment. *Econ. Syst. Res.* **24**, 173–193 (2012).
10. Miller, R. E. & Blair, P. D. *Input-Output Analysis: Foundations and Extensions* (Cambridge University Press, 2021).
11. Flegg, A. T. & Webber, C. D. Regional size, regional specialization and the FLQ formula. *Reg. Stud.* **34**(6), 563–569 (2000).
12. Flegg, A. T. & Webber, C. D. On the appropriate use of location quotients in generating regional input-output tables: Reply. *Reg. Stud.* **31**, 795–805 (1997).
13. Flegg, A. T., Webber, C. D. & Elliott, M. V. On the appropriate use of location quotients in generating regional input-output tables reg. *Stud.* **29**, 547–561 (1995).
14. Bonfiglio, A. & Chelli, F. Assessing the behaviour of non-survey methods for constructing regional input–output tables through a Monte Carlo simulation. *Econ. Syst. Res.* **20**(3), 243–258 (2008).
15. Jiang, X., Dietzenbacher, E. & Los, B. Improved estimation of regional input-output tables using cross-regional methods. *Reg. Stud.* **46**(5), 621–637 (2012).
16. Bacharach, M. *Biproportional Matrices and Input-Output Change* (Cambridge University Press, 1970).
17. Chizat, L., Peyré, G., Schmitzer, B. & Vialard, F. X. Scaling algorithms for unbalanced optimal transport problems. *Math. Comput.* **87**, 2563–2609 (2018).
18. Bartolucci, S., Caccioli, F., Caravelli, F. & Vivo, P. Upstreamness and downstreamness in input–output analysis from local and aggregate information. *Sci. Rep.* **15**, 2727 (2025).
19. Metulini, R., Gnecco, G., Biancalani, F. & Riccaboni, M. Hierarchical clustering and matrix completion for the reconstruction of world input–output tables. *ASTA Adv. Stat. Anal.* **107**(3), 575–620 (2023).
20. Gnecco, G., Nutarelli, F. & Riccaboni, M. A machine learning approach to economic complexity based on matrix completion. *Sci. Rep.* **12**, 9639 (2022).
21. Goodfellow, I. et al. Generative adversarial nets. In *NeurIPS* 27 (2014).
22. Candès, E. J. & Recht, B. Exact matrix completion via convex optimization. *Found. Comput. Math.* **9**, 717–772 (2009).
23. Pakizeh, A. H. & Kashani, H. Application of machine-learning models to estimate regional input coefficients and multipliers. *Spat. Econ. Anal.* **17**(2), 178–205 (2021).
24. Giudici, P., Raffinetti, E. & Toscani, G. Measuring multidimensional inequality: a new proposal based on the Fourier transform. *Statistics* **59**(2), 330–353 (2024).
25. Auricchio, G., Giudici, P. & Toscani, G. Extending the Gini index to higher dimensions via whitening processes. *Rend. Lincei Mat. Appl.* **35**(3), 511–528 (2025).
26. Dietzenbacher, E., Los, B., Stehrer, R., Timmer, M. & de Vries, G. The construction of world input-output tables in the WIOD project. *Econ. Syst. Res.* **25**(1), 71–98 (2013).
27. Timmer, M. P., Dietzenbacher, E., Los, B., Stehrer, R. & De Vries, G. J. An illustrated user guide to the world input–output database: The case of global automotive production. *Rev. Int. Econ.* **23**(3), 575–605 (2015).
28. Antràs, P., Chor, D., Fally, T. & Hillberry, R. Measuring the upstreamness of production and trade flows. *Am. Econo. Rev.* **102**(3), 412–416 (2012).
29. Szabó, N. Methods for regionalizing input-output tables. *Reg. Stat.* **5**(1), 44–65 (2015).
30. Mastronardi, L., Romagnoli, D. & Cavallini, A. Regional economic resilience: a systematic literature review. *Reg. Stud.* **56**, 181–196 (2022).
31. Bosone, C. & Stamato, G. Beyond borders: how geopolitics is reshaping trade. *ECB Working Paper Series*, No. 2960 (2024).
32. Cerchiello, P., Giudici, P. & Nicola, G. Twitter data models for bank risk contagion. *Neurocomputing* **264**, 50–56 (2017).

Acknowledgements

The authors gratefully acknowledge Tampere University Library (Finland) for providing open access support for this article. The corresponding author also wishes to thank the Scholars' Commons at the Herman B Wells Library, Indiana University Bloomington (USA) for providing essential academic support for the research underlying this manuscript.

Author contributions

F. De Pretis conceived the study, developed the methodology, and implemented the models. D. Tortoli conducted the experiments and performed the evaluations. S. Caria analyzed the results and contributed to economic interpretations. All authors reviewed and approved the final version of this manuscript.

Declarations

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1038/s41598-026-45382-8.

Correspondence and requests for materials should be addressed to F.D.P.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.