Can We Measure from the Bottom Up? Constructing an Index of Gas Station Infrastructure to Identify Regional Economic Development

Filkoski, Vasil and Tevdovski, Dragan

Ss. Cyril and Methodius University in Skopje, Faculty of Economics - Skopje

9 September 2025

CAN WE MEASURE FROM THE BOTTOM UP? CONSTRUCTING AN INDEX OF GAS STATION INFRASTRUCTURE TO IDENTIFY REGIONAL ECONOMIC DEVELOPMENT

Vasil Filkoski

Ss. Cyril and Methodius University in Skopje, Faculty of Economics - Skopje

Dragan Tevdovski

Ss. Cyril and Methodius University in Skopje, Faculty of Economics - Skopje

dragan@eccf.ukim.edu.mk

Abstract

We develop a methodology that leverages open-source geospatial data on fuel station infrastructure and related services to construct the Gas Station Index (GSI), a novel indicator that augments official and alternative measures of regional economic development. Gas stations serve as consumer-facing infrastructure nodes, and their density and quality reflect local demand, purchasing power, and mobility. Using data on 19,033 stations across 62 regions in nine European countries, the GSI explains 64% of the cross-regional variation in GDP per capita - a notable result for a single-variable indicator. Beyond its statistical fit, the GSI uncovers meaningful economic patterns. It reflects diminishing returns to infrastructure, consistent with core economic theory; it maps spatial inequality both visually and statistically, highlighting clusters of prosperity in capitals, port cities, transit corridors, and tourist destinations; and it classifies regional development typologies through bivariate LISA analysis. The unexplained variation underscores the structural differences between infrastructure-based indicator and GDP per capita, driven by sectoral specialization, mobility patterns, and informal economic activity. The GSI should therefore be viewed not as a substitute for national accounts, but as a complementary indicator *particularly relevant at the subnational level. Compared to existing indicators, it offers distinct advantages: GDP per capita is delayed and masks heterogeneity, while night-time lights suffer from saturation and rural undercoverage. By contrast, the GSI provides a ground-level, behaviorally grounded, and real-time measure of economic development. By capturing both infrastructure and consumption dynamics, it complements—and in certain respects surpasses—conventional indicators in tracing regional growth trajectories and spatial inequality.*

Keywords: regional income, regional inequality, economic development measurement, infrastructure, geospatial data, nowcasting.

JEL Codes: C43, C55, E01, O18, O47, R12.

1. Introduction

The accurate measurement of economic development remains one of the central challenges in economics. GDP remains the most widely used measure of economic activity, yet it is often delayed, unevenly reported, and inadequate to subnational economic heterogeneity. For decades, scholars and policymakers have turned to alternative indicators to fill this gap, with a tradition that spans both developing and advanced economies. Electricity consumption, for example, has often been employed as a substitute measure of economic activity. For instance, an IMF study on Jamaica used electricity consumption as a proxy for GDP and found that official statistics understated true output growth by an average of 2.7% per year during the 1990s, largely due to the expansion of the informal sector (IMF, 2006). Similarly, until 2005 the Federal Reserve Board based its monthly index of industrial production in part on a survey of utilities that measured electricity delivered to different classes of industrial customers. Young (2009) developed proxies for the level and growth rate of consumption in 56 developing countries using microeconomic data from the Demographic and Health Surveys. More recently, researchers have employed satellite data to construct proxies for economic activity observable from outer space. Sutton et al. (2007) and Henderson et al. (2012) were among the first to demonstrate that night-time light intensity is a reliable proxy for GDP growth at the national level. Together, this literature underscores both the necessity and the potential of creative proxies when conventional measures fall short.

This paper advances that tradition by proposing a novel indicator of regional economic development. The central idea is straightforward: if aggregate economic activity can be observed from outer space through night-time lights, then it may also be revealed "on the ground" through infrastructure that is both essential and ubiquitous. Yet unlike composite indicators that include multiple dimensions measuring roads, schools, hospitals, buildings, and other layers that together blur the picture - we argue that a single, well-chosen component can serve as a clear lens through which to approximate regional economic development. We suggest that gas station infrastructure provides precisely such a lens. Gas stations are directly linked to mobility, commerce, and household consumption, while their distribution reflects both purchasing power and regional development.

The logic for the Gas Station Index can also be anchored in economic theory. In models of urban and regional development, households and firms cluster where access to services and markets is best, and infrastructure nodes emerge as natural outcomes of this process (Fujita, Krugman & Venables 1999; Glaeser 2008). Gas stations function as such nodes: they require significant private investment, but remain viable only where local demand and purchasing power are high enough to sustain them and the surrounding services they attract. The growth literature further shows that infrastructure is both a driver and a signal of productivity (Calderón & Servén 2004). In parallel, evidence on household behavior suggests that fuel use and travel are income elastic: as incomes rise, people drive more and spend more on discretionary services, many of which co-locate around gas stations (Goodwin, Dargay & Hanly 2004). Finally, structural transformation theory highlights how rising prosperity shifts consumption toward services (Herrendorf, Rogerson & Valentinyi 2014), a pattern captured in the cafés, shops, and amenities embedded in modern station ecosystems. Taken together, these insights clarify why gas station infrastructure is more than a measure of fuel supply: it reflects demand, investment, and service clustering that together mirror regional development.

Figure I

A tale of gas station infrastructure in two contrasting regions – Budapest (Hungary) and Northeastern Region (North Macedonia)

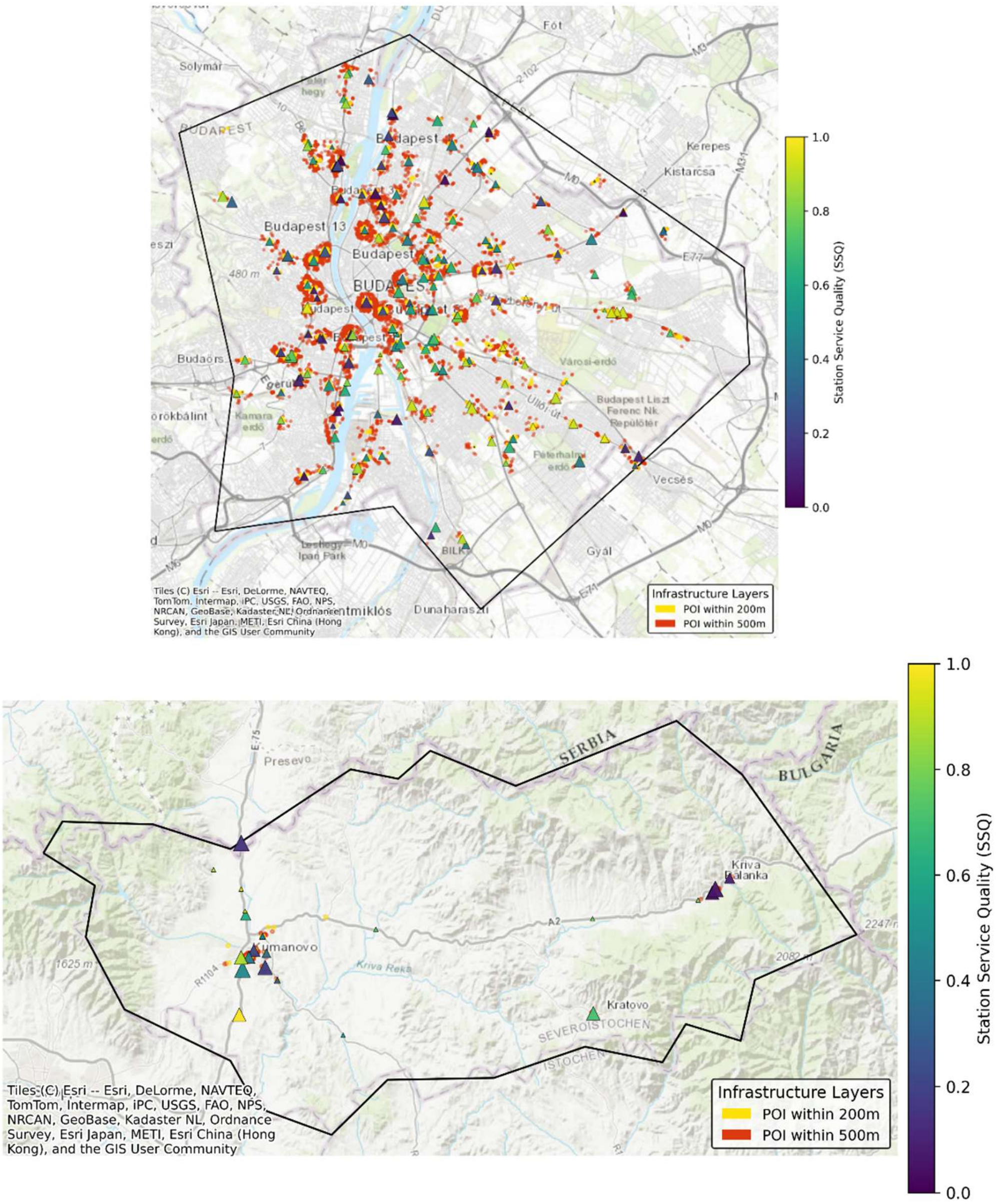

Note. Budapest (GDP per capita €44,837, 2023) shows dense station infrastructure; Northeastern Macedonia (GDP per capita €4,346, 2023) shows sparse infrastructure.

Figure I illustrate this proposition with a comparison of two regions in post-transitional economies from Europe. The Budapest region in Hungary, among the wealthiest in the area (GDP per capita equal to 44,837 euros in 2023), shows a dense and complex network of gas stations, mirroring its high GDP per capita and elevated living standards. In sharp contrast, the Northeastern Region in North Macedonia—one of the poorest in the area (GDP per capita equal to 4,346 euros in 2023), - exhibits a strikingly sparse gas station infrastructure. This stark difference suggests that gas station infrastructure may serve as a simple yet powerful indicator for regional economic development, providing policymakers and researchers with an alternative measure where conventional data are insufficient.

The logic for using gas stations infrastructure as a identification tool is threefold. First, the density and distribution of gas stations capture household fuel consumption, a discretionary good closely linked to income, commuting, and access to services; higher vehicle usage typically correlates with higher living standards. Second, because fuel sales yield thin margins, stations are strategically located in economically vibrant urban centers, market towns, and transit corridors. Zhu et al. (2024) show that site selection depends on population density, road connectivity, and fuel demand, underscoring their alignment with areas of economic vitality. Third, sustaining a viable network requires a consumer base with sufficient disposable income, reflected not only in vehicle ownership and travel but also in the profitability of on-site services—such as cafés and convenience shops—that command price premiums. Taken together, gas station infrastructure reflects private capital investment, consumer demand, and urban-commercial development, making it a meaningful composite indicator of regional GDP per capita.

To operationalize this idea, we introduce the Gas Station Index (GSI), a novel data-driven indicator constructed from open, crowd-sourced geospatial data. The GSI aggregates five dimensions: (i) service quality of each station, (ii) amenity density within 200 meters, (iii) amenity variety within 200 meters, (iv) amenity variety within 500 meters, and (v) accessibility (stations per 10,000 inhabitants and per 100 km^2). The index is scalable to any country, region, or time period, providing a transparent, low-cost tool for nowcasting regional prosperity.

We apply the GSI to 62 regions across nine European countries—Austria, Slovenia, Croatia, Hungary, Romania, Bulgaria, Serbia, North Macedonia, and Greece—covering 19,033 gas stations. The results merit particular attention: a single-variable index predicts 64% of the variance in regional GDP per capita, outperforming many traditional multi-variable indices. Beyond statistical fit, the GSI reveals meaningful economic patterns. It captures diminishing returns in infrastructure, consistent with core economic theory; it visually and statistically maps spatial inequality, highlighting clusters of prosperity in capitals, port cities, transit corridors, and tourist zones; and it identifies regional development typologies through bivariate LISA analysis.

This paper makes three main contributions. It first introduces a novel indicator for regional economic development that integrates infrastructure data, service quality, and amenity scoring into a unified framework. It then shows that the Gas Station Index (GSI) offers clear advantages over existing identification tools: unlike GDP per capita, which is delayed and masks heterogeneity, or night-time lights, which suffer from saturation and undercoverage in rural or energy-efficient areas, the GSI provides a ground-level, behaviorally grounded, and real-time measure of economic development. Third, it shows that the GSI identifies development stages, and spatial patterns: trade corridors, development hotspots, and stagnating peripheries, that are often misvalued in GDP and nighttime light measurements. By capturing both infrastructure and consumption dynamics at ground level, the GSI complements and in some aspects surpasses conventional development indicators.

The rest of the paper is organized as follows. Section 2 reviews the relevant literature on alternative indicators for economic dynamics, situating our contribution within this body of work. Section 3 presents the construction of the Gas Station Infrastructure Index (GSI), outlining its conceptual foundations and methodological design. Section 4 describes the data and illustrates the application of the GSI to selected European regions. Section 5 reports the empirical results, assessing the explanatory power of the GSI in relation to regional GDP per capita. Section 6 discusses the implications of these findings, highlighting both the strengths and limitations of the GSI as a proxy measure. Section 7 concludes by summarizing the main contributions and suggesting directions for future research.

2. Literature Review

While no prior studies have explored gas stations as a metric for economic development, a growing body of literature has focused on identifying alternative indicators for measurement of economic dynamics. These indicators include satellite night lights, bank and ATM density, and coworking space presence.

Although Sutton et al. (2007) was the first study in the economics literature to introduce satellite night lights as an indicator of economic activity, the approach gained broader recognition only after the influential publication by Henderson et al. (2012). They demonstrate that the amount of light observable from space is a reliable proxy for GDP growth at the national level. Their analysis shows that light intensity effectively captures both long-term growth trends and short-term GDP fluctuations. Since then, over 450 papers in the economics literature (based on IDEAS/RePEc) have used night lights as a indicator for economic activity (e.g., Pinkovskiy, and Sala-i-Martin, 2016; Henderson et al., 2018). However, Gibson et al. (2020) caution that night lights data may have limitations, particularly in temporal comparisons and in capturing activity in small or rural areas.

Another strand of the literature highlights the role of financial infrastructure—specifically, bank branch and ATM density—as a indicator for economic activity. Kireyeva et al. (2021) find a positive association between bank/ATM density and economic growth across emerging markets, reinforcing the relevance of financial infrastructure as a proxy for economic performance. However, in advanced economies, traditional bank branch density—measured as the number of branches relative to total deposits—has significantly declined over the past decade, due to the rise of digital banking. Benmelech, Yang, and Zator (2023) also highlight the importance of bank density in the context of financial distress, particularly in explaining vulnerability to bank runs.

The rise of self-employment, freelancing, and remote work has led to increased interest in coworking spaces as an indicator of entrepreneurial activity and the growth of the service economy. Mariotti et al. (2021) argue that coworking spaces are contemporary drivers of regional economic growth. Empirical studies by Gauger et al. (2021) and Clifton et al. (2022) find significant correlations between coworking space density and local startup intensity.

Taken together, this literature demonstrates a growing acceptance of alternative, infrastructure-based proxies in economic research. In line with this methodological trend, the Gas Station Infrastructure (GSI) Index offers a grounded yet novel indicator of regional economic well-being—one that captures visible, accessible, and trackable aspects of economic life through the distribution and evolution of fuel station infrastructure.

3. Constructing the Gas Station Infrastructure Index (GSI)

The idea of Gas Station Infrastructure Index (GSI) is to serve as a bottom-up indicator for regional economic development. It is a novel, multi-dimensional indicator designed to measure economic development, on a regional and national level. It is constructed using micro-geospatial data derived from gas station infrastructure and its surrounding development.

The GSI is based on the scores of five pillars: Station Service Quality, Amenity Density within 200 meters, Amenity Variety within 200 meters, Amenity Variety within 500 meters, and Gas Station Accessibility (Table I).

TABLE I

Pillars of the Gas Station Infrastructure Index

Pillars	Station Service Quality (SSQ)	Amenity Density 200m (AD200)	Amenity Variety 200m (AV200)	Amenity Variety 200-500m (AV500)	Gas Station Accessibility (GSA)
Description	- Score based on intrinsic station attributes: fuel variety, payments, brand tier, retail services, metadata. - Reflects internal investment and service quality.	- Weighted Places of Interest (POIs) density within 200 m using economic multipliers. - Captures local commercial intensity.	- Shannon entropy of POI types within 200 m radius - Measures functional diversity and land-use complexity.	- Same as AV200 but measured between 200 m and 500 m radius. - Captures transit corridors or economic sprawl.	- Station density per km² and per 10k residents. - Reflects accessibility and spatial service reach.
Economic Interpretation	- Station-level investment, expected demand, and service differentiation.	- Density of localized economic clustering and consumer commercial activity.	- Urban complexity, service variety, and potential consumer income.	-Macrolocation integration into broader economic networks.	- High GSA reflects strong infrastructure, demand, and mobility access; low GSA may indicate rural neglect or transit poverty.

Note. This table summarizes the five pillars of the Gas Station Index: Station Service Quality (SSQ), Amenity Density within 200 meters (AD200), Amenity Variety within 200 meters (AV200), Amenity Variety within 200–500 meters (AV500), and Gas Station Accessibility (GSA). Each pillar reflects distinct aspects of investment, demand, spatial density, and service variety.

3.1. Station Service Quality (SSQ)

This pillar evaluates the station's level of services and economic footprint based on the following elements: brand tier, fuel availability, opening hours, payment options, retail services, and metadata completeness.

- Brand Tier: The gas station's brand is classified into global, regional, national and local based on presence across country. The classification is confirmed through both the dataset and the market share (Global = 20, Regional = 15, National = 10, Local = 5).
- Fuel Availability: The presence of fuel: tags is evaluated. The scoring is +8 points for having at least one fuel:* = yes tag and +2 for each additional, capped on 30.
- Opening Hours: The presence and type of opening_hours tags is evaluated. 20 points are awarded if the gas station has an opening_hours = 24/7 tag, 10 if there is another tag, and 0 if the tag is missing
- Payment Options: A similar logic to the one of Fuel Availability is applied. The station gets +5 score for first payment:* = yes method, +2 for each additional.
- Retail Service: The station gets +10 score if a retail element (like shop=convenience or amenity=fast_food) is present on-site or within 20m using BallTree radius checks.
- Metadata Completeness: The station gets +10 score if the station includes at least 3 of: capacity, website, ref, wheelchair, toilets, internet_access, or any contact:* key.

The station service quality score (0–100) reflects gas station standards, with higher values indicating better quality. To avoid undervaluation in data-sparse regions, two fallback mechanisms are applied: missing data are partly imputed by brand tier, and for stations with very low scores (<10) but above-average amenity density/diversity, a baseline score of 30 is assigned, assuming minimal service is likely in active areas.

3.2. Amenity Density 200 m (AD200)

This pillar reflects the economic activity intensity in a 200 meter radius. All Points of Interest (POIs) within this buffer are collected via spatial query (with prefiltered bounding box + exact distance using haversine). The POIs selected for the GSI are: restaurant, supermarket, cafe, bar, car_repair, hotel, pharmacy, fast_food, convenience, bank, charging_station, atm, bus_station, parking. Each POI is assigned a quality weight based on its key and value (e.g., restaurant = 1.0, fast_food = 0.6)

$$AD200_i = \frac{\sum_{j \in \text{POIs within } 200m} w_j}{\pi \cdot (0.2)^2}$$

Distances and counts are computed using a two-stage spatial indexing approach. An R-tree (rtree.Index) enables efficient bounding-box queries to preselect candidate POIs, followed by a BallTree (sklearn.neighbors) with Haversine distance to identify exact neighbors within 200 m and 500 m. This procedure allows consistent estimation of POI density and entropy-based diversity measures while maintaining computational feasibility. Intermediate POI lists, station objects, and scoring components are cached at the regional level to facilitate reuse and iterative analysis.

3.3. Amenity Variety 200 m (AV200) and Amenity Variety 200-500 m (AV500)

These two pillars reflect the diversity and complexity of economic services surrounding each gas station. While AD200 captures the intensity of POI activity, AV200 and AV500 aim to capture how many distinct types of amenities are nearby—an important proxy for urbanization and service richness. To quantify variety, a weighted Shannon entropy score is calculated based on all Points of Interest (POIs) within the corresponding buffer radius (200 m and 200-500 m, respectively). For each unique POI type (defined as a (key, value) pair such as amenity=restaurant or shop=supermarket), the proportion of its weighted count relative to total weight is computed, and the entropy formula is applied:

$$AV_r = -\sum_{k=1}^{K} p_k \cdot \ln(p_k) \qquad\qquad p_k = \frac{w_k \cdot c_k}{\sum_{j=1}^{K} w_j \cdot c_j}$$

The higher the entropy, the greater the amenity diversity around the station. AV200 focuses on immediate economic heterogeneity, while AV500 captures broader neighborhood complexity, and also better represents rural and transit gas stations.

3.4. Gas Station Accessibility (GSA)

This is the only pillar not derived directly at the station level. It is precomputed at the **regional level** and injected into each station during the pipeline. GSA is a composite score combining number of stations per 10,000 inhabitants and number of stations per 100 km². Both values are log-transformed to avoid extreme region size bias and weighted 50:50 before being scaled globally.

$$GSA_r = 1/2\left(\ln\left(\frac{\text{Stations}_r}{\text{Population}_r/10{,}000}\right) + \ln\left(\frac{\text{Stations}_r}{\text{Area}_r/100}\right)\right)$$

The calculation of the 4 pillars on station level gives the raw scores of the pillars before any scaling or normalization is applied. Afterwards, for each administrative region (NUTS2 or NUTS3, depending on the country), the station level pillar scores are averaged to produce a single raw regional value per pillar:

$$P^{\text{region}} = \frac{1}{n}\sum_{i=1}^{n} P_i$$

This aggregation happens before the global scaling or index construction.

Once raw region-level pillar values are obtained for all regions, each pillar is normalized across the full dataset using a two-phase process: i) clipping of the pillars on a station level, per region, and ii) global across regions min-max normalization of each pillar.

i) For every region, each station's SSQ, AD200, AV200, and AV500 are winsorized at each region's 5th and 95th percentile, reducing the impact of outliers. This prevents data quality-related false positives and false negatives from biasing the results. Afterwards, the region aggregation is calculated from these trimmed scores.

$$x_i^{\text{wins}} = \begin{cases} Q_{0.05}(P_r), & \text{if } x_i < Q_{0.05}(P_r) \\ Q_{0.95}(P_r), & \text{if } x_i > Q_{0.95}(P_r) \\ x_i, & \text{otherwise} \end{cases}$$

This ensures limited influence of extreme outliers on a regional level.

ii) After aggregating regional scores, for each pillar a global min-max normalization is applied. All regional pillar values are now in the [0, 1] range and directly comparable across countries.

$$P_r^* = \frac{P_r - \min_j(P_j)}{\max_j(P_j) - \min_j(P_j)}$$

After the min-max normalization the GSI is constructed with equal weighs of the pillars, effectively performing a mean aggregation of the 5 pillars:

$$\text{GSI} = \frac{1}{5}\,(\text{SSQ} + \text{AD200} + \text{AV200} + \text{AV500} + \text{GSA})$$

The following section applies this methodology to 62 NUTS2 and NUTS3 regions across 9 European countries to test the GSI's predictive validity.

To assess the usefulness of equal weighting of the five pillars, 10,000 randomly generated weight combinations were tested to explore whether different weightings of certain pillars could improve model fit. The gains over equal weighting were marginal. The best-performing custom weights only slightly outperformed equal weighting in terms of R^2, and given the trade-off between complexity, overfitting risk, methodological replicability, and interpretability, equal weights across all five pillars were retained to ensure simplicity and replicability of the index.

4. Data: Applying the Gas Station Index to Selected Regions

4.1. Data Collection and Processing

During data collection, we evaluated several approaches, including the Overpass API, Google Places API, and locally running OpenStreetMap data. The Overpass API, which enables programmatic queries of OpenStreetMap, was rejected due to technical instability in large-scale regional extraction, including frequent timeouts, rate limiting, and inconsistent performance that rendered automated scripts unreliable. The Google Places API, while offering potentially higher coverage in some areas, was also excluded because of high usage costs that undermine replicability and open-access use. Instead, this paper relies on pre-downloaded *.osm.pbf* files from Geofabrik.de, which provide freely available, regularly updated OpenStreetMap extracts by country. This choice ensures scalability, reproducibility, and zero cost.

The core data for the GSI are derived from these *.osm.pbf* files. Using *osmium*, the dataset is parsed to extract all gas stations tagged as *amenity=fuel*, together with attributes such as location, brand, name, presence of a shop, car wash, toilet, payment terminal, and other service features. In parallel, a curated set of additional POIs (e.g., *amenity=restaurant*, *shop=convenience*, *amenity=pharmacy*, *tourism=hotel*) is extracted from the same files. These contextual POIs are used to construct the index's supporting pillars.

The Gas Station Index (GSI) is applied to 62 regions across 9 European countries[1]: Austria, Slovenia, Croatia, Hungary, Romania, Bulgaria, Serbia, North Macedonia, and Greece - providing variability both across countries with differing levels of economic development and within countries that exhibit disparities between more and less developed regions. In total, 19,033 gas stations were extracted, scored, and aggregated using the methodology outlined in Section 3. The selected countries and their respective regions

[1] The countries from the region that are missing: Bosnia and Herzegovina, Montenegro, Kosovo, and Albania; were excluded for the following reasons. Bosnia and Herzegovina and Kosovo lack NUTS-standardized statistical regions, Montenegro has only one NUTS2 region, and Albania although fulfilling the NUTS condition, hasn't released regional GDP data since 2021 deeming it unfitting for the empirical statistical analysis and machine learning modeling.

fulfill the requirements of both Method of Most Differences and Method of Least Differences. Most differences in the sense that the selected countries vary in level of economic development, EU membership status (long-standing EU members, recent EU members, non-EU members), and economic structure of the national and regional economies. The least differences comparability is expressed through providing comparability between regions in the same country and across countries of similar levels of development (example: Steiermark in Austria and Eastern Slovenia).

TABLE II

Scores of the GSI pillars for an example gas station in Vienna (AT13)

Pillar	Data	Raw Value	Winsorized Value (within Vienna region)
SSQ	The gas station got an SSQ score of 65/100. 20 points from being a global brand, it has 4 payment methods 5+2*3 = 11 points, 4 fuel types, 8+2*3 = 14 points, it has rich metadata tags, an additional 10 points, and it has a kiosk tag, an additional 10 points.	65/100	0.926
AD200	The gas station has an AD200 score of 364.46 POI objects in a 200 m radius around it. This signifies an exceptionally economically active area.	364.46	1.000
AV200	The gas station got an AV200 score of 3.51, signifying high Shannon entropy and variety of the surrounding POI objects in a 200 m radius	3.51	1.000
AV500	The gas station got an AV500 score of 4.10, signifying high Shannon entropy and variety of the surrounding POI objects in a 200-500 m radius	4.10	1.000
GSA		/	/

Note. Pillar scores reflect raw and winsorized values for a single station in the Vienna region. SSQ is based on station attributes; AD200 and AV200 measure density and variety of amenities within 200 m; AV500 measures variety in the 200–500 m buffer.

To illustrate the Gas Station Infrastructure Index (GSI) methodology, we apply it to a gas station, which is part AT13 - Vienna region in Austria.[2] This station achieves one of the highest GSI scores in the dataset and exemplifies the mechanics behind the index's construction. The gas station is located on the main street in a high land value area near Reumannplatz, close to the Vienna Central Train Station. Among extra services, it has a market, fast food, cafe, and car wash. It also has one of the central boulevards, multiple bus stops and a metro stop in its vicinity. Overall, its location represents an economic urban hotspot in Vienna (Figure II). The scores of each pillar of this gas station are provided in Table II.

Figure II

A gas station in AT13 - Vienna region

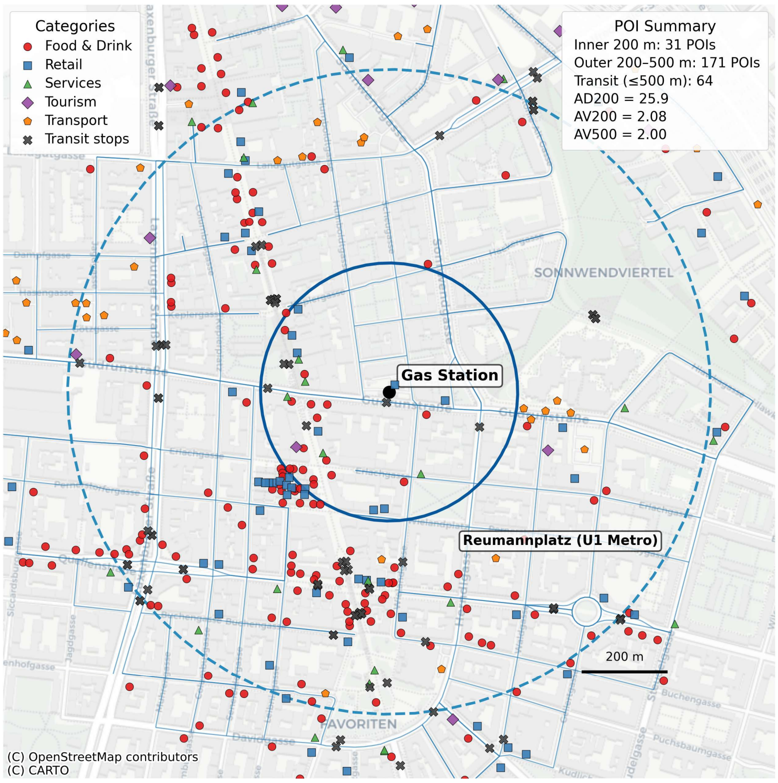

Note. Example of a high-scoring station in Vienna, with multiple on-site services and high surrounding amenity density and variety.

[2] The metadata of this station are: brand = Eni, region_id = AT13 (Wien), coordinates: lat= 48.178058, lon = 16.378547, id = 329280055 (its OpenStreetMap ID number).

Table III

Regional pillar scores for Vienna (AT13)

Region	SSQ	AD200	AV200	AV500	GSA	GSI
AT-13, Wien	1.00	0.385	1.00	0.8523	0.566	0.761

Note. Mean winsorized pillar scores aggregated across all stations in Vienna. Each pillar is normalized globally to [0, 1].

Similarly, the remaining 207 stations in the AT13 - Vienna region are scored and winsorized. Afterwards, the mean score of the pillars is calculated inside the region, creating the regional pillar scores. The regional pillar scores, across all regions, undergo a min-max normalization, giving the following scores for AT13 - Vienna, that are presented in Table III.

4.2. Overview of the Regions

This map in Figure III visualizes the spatial distribution of the GSI across NUTS regions of the study's selected countries. The different GSI scores outline level of prosperity patterns. There are urban powerhouses like Vienna, Budapest, Bucharest, Belgrade, Skopje, and Athens emerge as GSI hotspots. Peripheral regions with lower GSI scoring are found in rural Bulgaria, northern Greece, and the infrastructurally lagging regions of North Macedonia. Slovenia and Austria show a uniformly strong regional performance, pointing to consistent regional development. The findings from the GSI scores, especially the overperforming and underperforming outliers (in relation to regional GDP) will be thoroughly discussed further down in the text in Section 5.

Figure III

Gas Station Index (GSI) scores across selected European regions

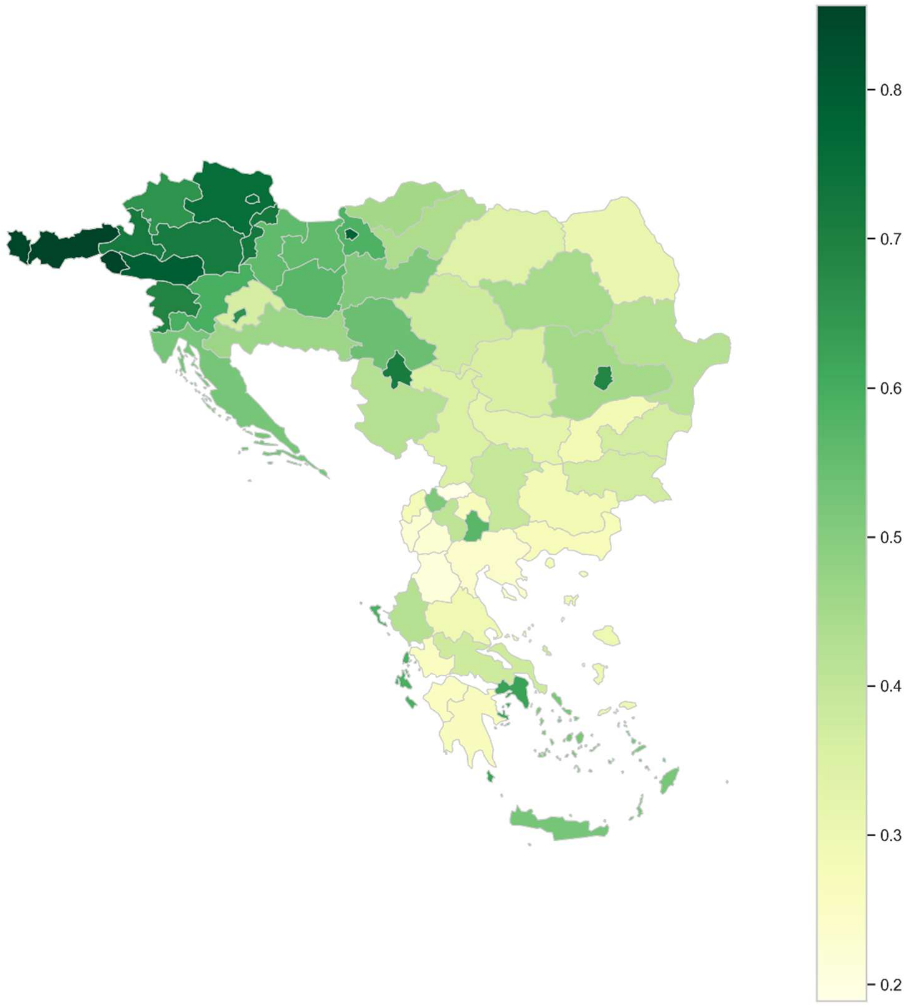

Note. Map displays regional variation in GSI across nine European countries, highlighting spatial patterns of prosperity and underdevelopment.

Having constructed and aggregated the GSI scores across 62 regions, we now turn to an empirical assessment of its predictive power. In the next section, we apply correlation analysis, linear and non-linear regression models, spatial diagnostics (e.g. Moran's I, LISA), and SHAP-based interpretation to test how well the GSI explains regional economic output.

5. Empirical Results: Assessing the explanatory power of the GSI in relation to regional GDP per capita

This section presents the empirical results of the analysis, with the aim of evaluating the explanatory power of the Gas Station Infrastructure Index (GSI) and its relationship to regional GDP per capita. The results are organized into four subsections. First, we assess the explanatory power and model fit of GSI through a comparative application of linear regression, random forest, and XGBoost models, complemented by correlation analysis. Second, we examine the residual patterns of the OLS and random forest models to identify systematic deviations and nonlinearities in the relationship between GSI and GDP per capita. Third, we explore the spatial dimension of the data by applying measures of spatial autocorrelation and local indicators of spatial association (LISA), which reveal regional clustering and highlight spatial heterogeneity. Finally, we conduct robustness checks to validate the stability and consistency of our findings across alternative specifications. Together, these analyses provide a comprehensive empirical basis for assessing the extent to which GSI can serve as a reliable predictor of regional economic outcomes.

5.1. Explanatory Power and Model Fit of the GSI

In this subsection, we examine the capacity of the GSI to explain variations in GDP per capita across the observed regions. To this end, we apply correlation analysis, ordinary least squares (OLS) linear regression, as well as two non-linear machine learning approaches: Random Forest and XGBoost.

Correlation Analysis

We begin the analysis with a graphical representation of the correlation between the Gas Station Index (GSI) and the regional GDP per capita. Figure IV shows the scatter plot of these variables, where each point represents an analyzed region. It reveals a clear positive association: as GSI increases, regional GDP per capita tends to increase, and vice versa.

Figure IV

Scatter plot of GSI and GDP per capita with regional development tiers

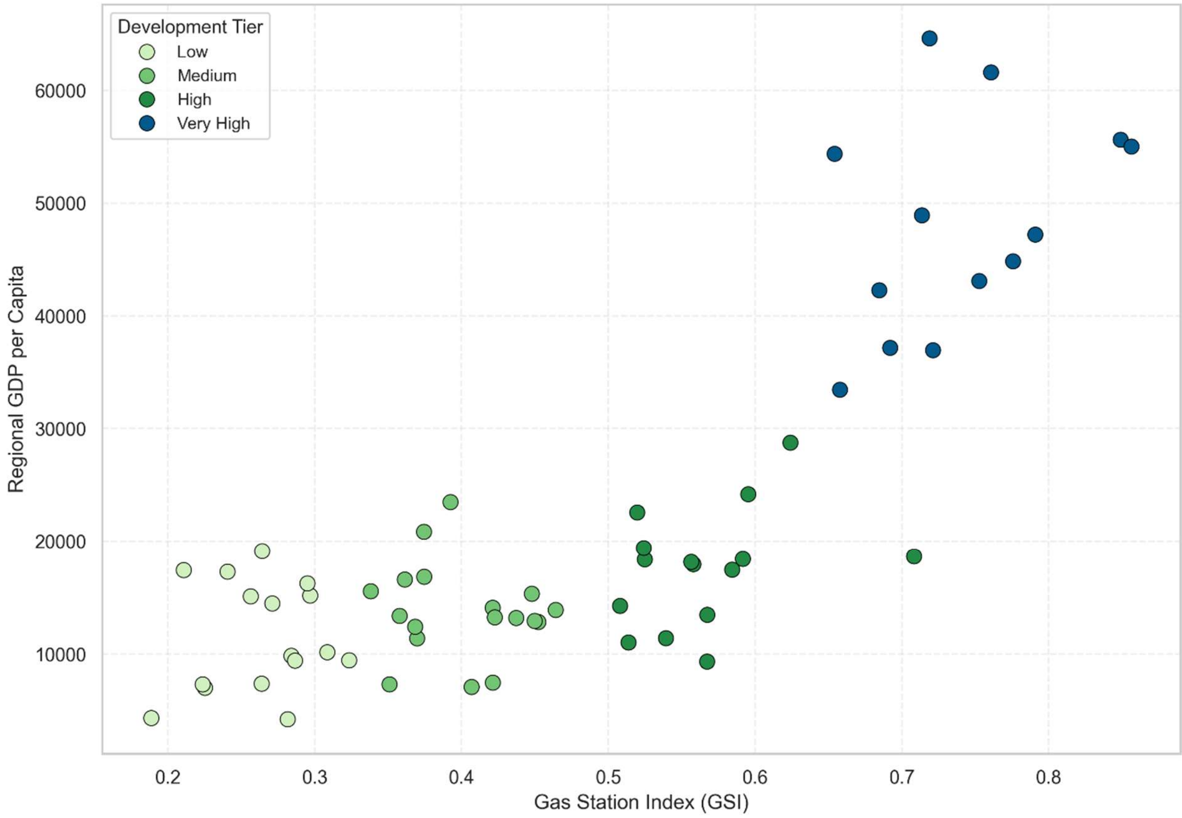

Note. Scatter plot shows strong positive correlation between GSI and GDP per capita across 62 regions. K-means clustering identifies four development tiers (Low, Medium, High, Very High).

The correlation analysis confirms a strong positive association between the Gas Station Index (GSI) and regional GDP per capita. The Pearson correlation coefficient is 0.80, indicating a strong linear relationship, while the Spearman rank correlation is 0.72, suggesting that the association also holds in terms of rank-order across regions.

The proximity of the two values indicates that the relationship is not only strong but also monotonic and relatively stable across the output distribution, with limited influence from extreme outliers or nonlinear distortions. These results provide further support that the GSI captures core variation in regional GDP per capita. Both in magnitude and in relative position across the regional income spectrum.

Moreover, the regions are clustered into four development tiers - Low, Medium, High, Very High using k-means training. Through the tier coloring, the motion that infrastructure and service quality (GSI) scales with regional GDP across the entire distribution is reinforced. The GSI distinguishes not just between high and low income regions, but also captures signals of gradients of development in the middle of the distribution. This represents more than a binary signal.

This association provides intuitive and visually compelling validation of the Gas Station Index (GSI). Regions with higher GSI scores typically exhibit better infrastructure, a wider variety of retail and payment options, and more comprehensive service offerings - characteristics commonly linked to greater economic activity and consumer purchasing power. The strong positive correlation between the GSI and regional GDP per capita suggests that regions with more dense and higher fuel station service quality also demonstrate stronger economic performance. This correspondence supports the view that consumer-oriented infrastructure and service provision can serve as reliable indicators of broader economic prosperity of the region, positioning the GSI as a valuable tool for regional economic analysis.

Ordinary Least Squares (OLS) Regression

We estimated an ordinary least squares (OLS) regression to examine the relationship between the Gas Station Index (GSI) and regional GDP per capita. The results in Table IV show that the GSI is a highly statistically significant predictor ($p < 0.001$), with a coefficient of approximately €67,800. This implies that a 0.1 increase in the GSI is associated with an increase of around €6,780 in GDP per capita.

Table IV

OLS regression results

	Coefficients	Lower confidence bound (0.025)	Upper confidence bound (0.975)
Constant	-1,102** (0.002)	-1,770	-4,358
GSI	6,780*** (0.000)	5,470	8,090
R-squared: 0.642		AIC: 1310; BIC:1314	
Adj. R-squared: 0.636		Durbin-Watson: 2.118	
F-statistic: 107.6*** (0.000)		Jarque-Bera (JB): 2.807 (0.246)	
No. of observations: 62		Skew: 0.470; Kurtosis: 3.450	

Note. The dependent variable is regional GDP per capita. Standard errors are in parentheses. ** $p < 0.01$, *** $p < 0.001$.

The model explains about 64.22% of the variance in regional GDP per capita ($R^2 = 0.642$), which is substantial for a single-variable model in regional economics. The F-statistic (107,6, $p < 0.001$) confirms the model's strong explanatory power, while the narrow confidence intervals for the GSI coefficient (€5,470 to €8,090) reinforces its robustness.

Residual diagnostics reveal no major concerns: the Jarque-Bera test ($p = 0.246$) indicates that residuals are overall normally distributed; skewness (0.47) and kurtosis (3.45) are within acceptable bounds. The Durbin-Watson statistic (2.12) suggests no problematic autocorrelation, appropriate in a cross-sectional context where spatial autocorrelation is separately evaluated via Moran's I.

Overall, this OLS model validates a strong and linear relationship between gas station infrastructure quality and regional economic prosperity, forming a statistically sound baseline against which more complex models like Random Forest or XGBoost can be compared.

Non-Linear Model Performance: Random Forest and XGBoost

To capture potential non-linearities and interaction effects in the association between GSI and regional GDP per capita, the same data was modeled using two machine learning algorithms: Random Forest and XGBoost. Table V demonstrates that their in-sample explanatory power exceeds that of the OLS linear regression (R^2 of 0.642). Random Forest, a non-parametric ensemble algorithm that averages predictions from multiple decision trees, achieved a very high in-sample R^2 of 0.951. It effectively captured complex non-linearities that the linear model could not, such as threshold effects and regional heterogeneity in the impact of infrastructure quality. Similarly, XGBoost, an advanced gradient boosting algorithm, performed even better on in-sample data $R^2 = 0.968$, (MAE = 2,156, RMSE = 2,722), outperforming both Random Forest and Linear Regression. These discrepancies in the in-sample explanatory power between the two non-linear machine learning algorithms and OLS linear regression imply that the relationship between GSI and GDP per capita is not strictly linear. This non-linearity arises because, as regions become more developed, further improvements in GSI are likely to exhibit diminishing marginal returns—a pattern that flexible non-linear models such as Random Forest and XGBoost are capable of capturing from the data, but which a linear specification is structurally unable to represent. Figure V, which displays the actual values alongside the model predictions, provides graphical confirmation of this association.

To assess the generalizability of the GSI across unseen regional contexts, a 5-fold cross-validation procedure was implemented. Table V presents cross-validated R^2 of the models. Although machine learning algorithms achieve superior in-sample fit, OLS performs comparably better in cross-validation, highlighting the strength of the GSI's linear association with GDP per capita. The fact that all models retain over 57% explanatory power across folds indicates that the GSI is not only theoretically sound, but statistically stable across regions, despite variations in national context, infrastructure, or development stage.

Table V

Comparison of model performance across OLS, Random Forest, and XGBoost

Model	In-Sample R^2	Cross-Validated R^2
OLS	0.642	0.603
Random Forest	0.951	0.588
XGBoost	0.968	0.572

Note. Table reports in-sample and cross-validated R^2 for each model. OLS shows stable performance across validation folds; Random Forest and XGBoost achieve higher in-sample fit but lower cross-validated stability.

Figure V

Linear Regression, Random Forest, and XGBoost Regression Fits

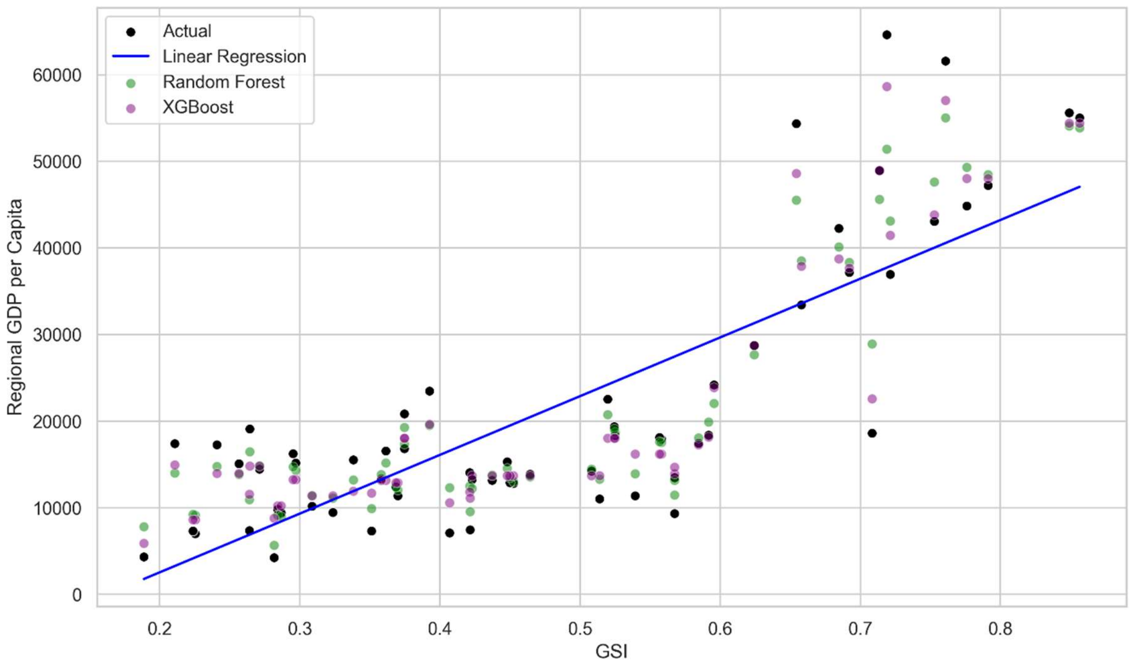

Note. Comparison of OLS, Random Forest, and XGBoost predictions of regional GDP per capita against actual values. Non-linear models capture threshold effects more effectively.

5.2. Diminishing Returns and Welfare Plateaus

While the Gas Station Index (GSI) is a strong explanatory variable for regional GDP per capita, its effect is not constant across all development levels, and the relationship appears not fully linear. As regions become wealthier, the impact of GSI begins to flatten. This section explores that non-linear relationship using both residual analysis and SHAP interpretation from the Random Forest model.

Examining the Residuals of OLS and Random Forest

Figure VI shows that the OLS regression residuals are not randomly scattered but instead follow a systematic U-shaped pattern. This curvature suggests that the relationship between GSI and GDP per capita is not strictly linear. In regions with low GDP per capita (below €15,000), the residuals are generally small, indicating that the model fits relatively well. However, in regions with middle GDP per capita values (between €15,000 and €30,000), the model systematically overestimates GDP per capita. These areas are typically more developed regions in a country, including capital cities, transportation corridors, and tourism hubs. In such cases, gas station infrastructure appears to be running ahead of income levels. One explanation lies in EU-funded investments and the rapid penetration of international brands, which raise infrastructure standards more quickly than regional production capacity and GDP per capita growth, with income levels catching up only after a lag. Similarly, coastal regions with significant tourism potential often feature dense networks of amenities, services, and infrastructure. Here, regional GDP per capita may under-record visitor spending or leak value to other regions, while GSI captures the more immediate infrastructure response to tourism demand. A further explanation relates to the role of the informal economy. In regions where informal activity is widespread, GSI may reflect higher levels of economic prosperity than are captured by official GDP statistics, leading the model to predict higher GDP per capita levels than those formally recorded.

Figure VI

OLS regression residuals

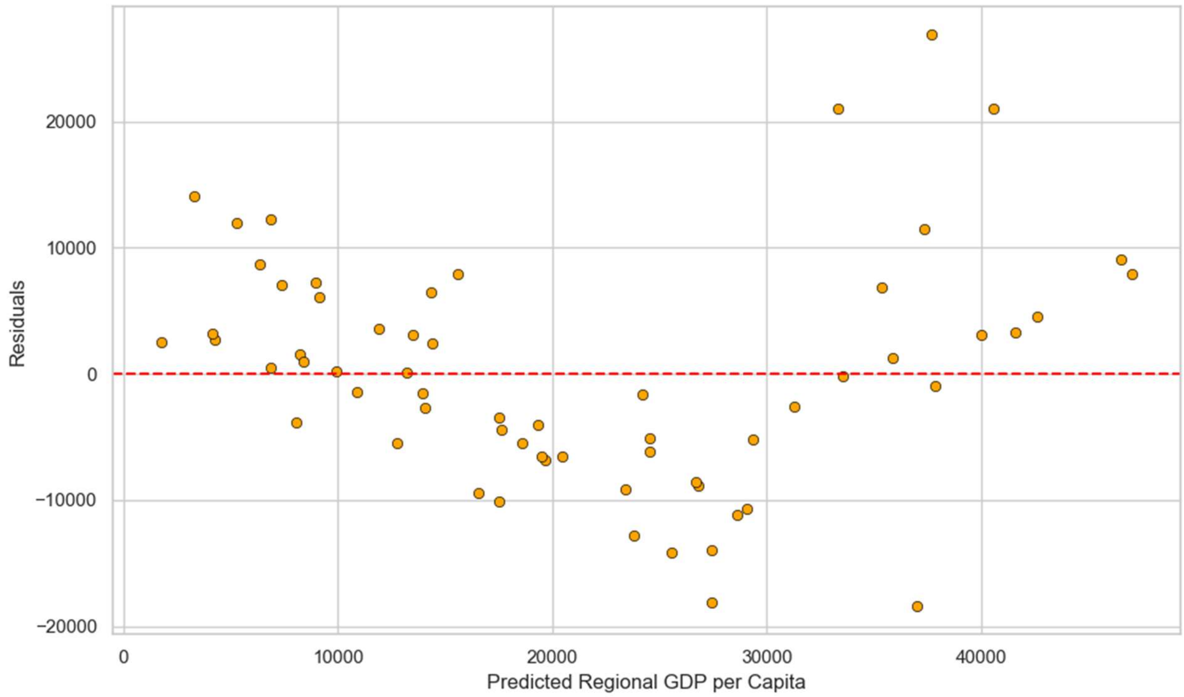

Note. Residual plot shows systematic U-shaped pattern, indicating model overestimation in middle-income regions and underestimation in high-income regions.

In regions with high GDP per capita (above €30,000), the residuals rise again. These areas, characterized by both high GSI scores and high regional GDP per capita, are systematically underestimated by the model. This suggests that once a solid level of infrastructure is in place, other factors of economic prosperity - such as the expansion of the service economy - become more decisive. At the same time, the principle of diminishing marginal returns to infrastructure becomes evident. Once gas station infrastructure has reached saturation, further improvements in infrastructure yield progressively smaller gains in GDP per capita.

Figure VII

Random Forest regression residuals

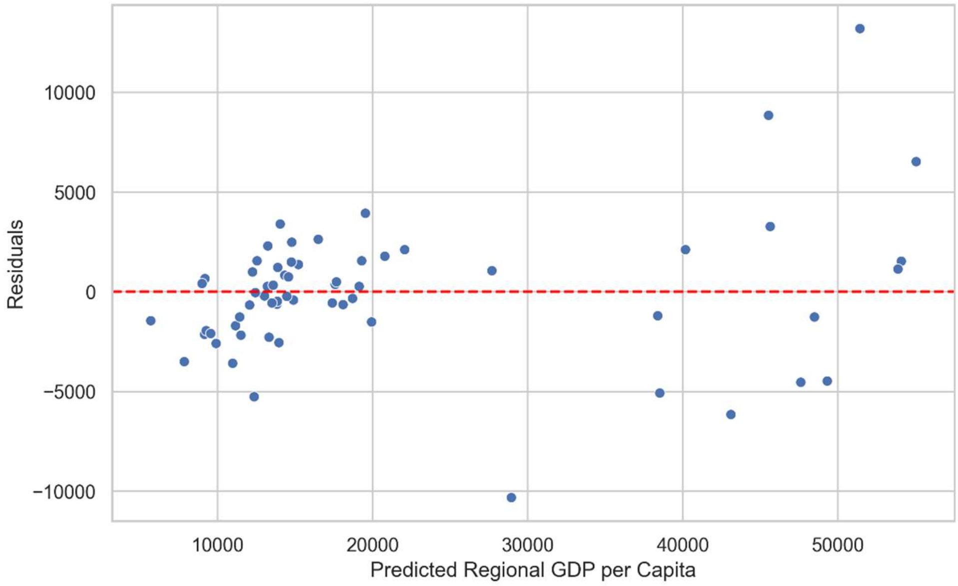

Note. Residuals are more balanced than in OLS, capturing non-linearities in the GSI–GDP relationship. Variance widens in high-income regions where GSI effects diminish.

Figure VIII

SHAP values for the Random Forest model

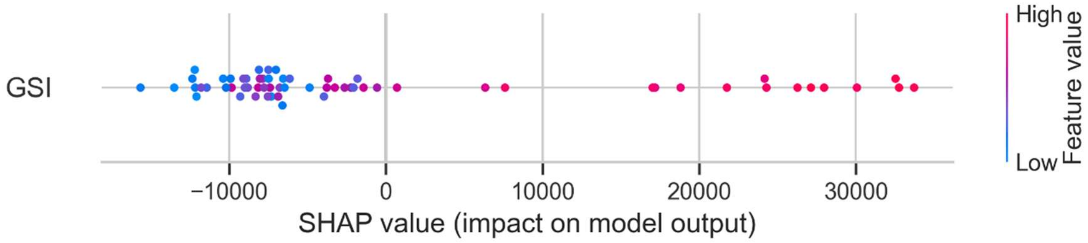

Note. SHAP summary plot shows that low-GSI regions receive strong positive contributions to predicted GDP per capita, while marginal returns diminish in high-GSI regions.

Figure VII presents that Random Forest residuals appear far more balanced compared to OLS. The systematic U-shaped curvature disappears, indicating that the non-linear model more accurately reflects the diminishing-returns dynamics in the GSI–GDP relationship. In low- and mid-income regions, prediction errors are small and symmetrically distributed, consistent with the idea that infrastructure expansion remains a binding constraint and translates efficiently into output gains. At the top of the distribution, however, residuals widen once again, even if they no longer follow a systematic bias. This greater variance reflects the fact that, beyond a certain threshold, infrastructure saturates and additional GDP per capita is increasingly driven by factors outside the scope of GSI explanatory power - such as innovation, finance, global integration, and quality-of-life improvements. Thus, the Random Forest results reinforce the same insight as the OLS model: while infrastructure is decisive for growth at lower and middle stages, its marginal contribution diminishes in advanced economies, setting the stage for alternative interpretations of development plateaus. This pattern is also evident in the SHAP summary plot (Figure VIII): regions with low GSI values (blue) get large negative SHAP scores — meaning GSI helps the model "pull up" their predicted GDP per capita. In contrast, high-GSI regions (pink) have SHAP scores that flatten. Their GSI still helps, but not by much. It adds little to what the model already knows. In other words, GSI has strong effects in low and mid-income areas, and diminishing effects at the top.

5.3. Spatial Autocorrelation and Regional Clustering

We apply in this subsection a set of spatial autocorrelation methods to analyze the relationship between GSI and GDP per capita across regions. Global Moran's I is first used to test whether regions with similar values of the GSI and GDP per capita are geographically clustered, thereby assessing the presence of global spatial dependence. Next, Univariate Local Indicators of Spatial Association (LISA) is then employed to identify the specific locations of significant clusters of GSI, capturing local variation in spatial patterns. Finally, Bivariate LISA extends the analysis by examining spatial co-clustering between GSI and GDP per capita, providing evidence on localized interactions between gas stations infrastructure and regional economic outcomes.

Moran's I (Global Spatial Autocorrelation)

The Global Moran's I statistic for GRP per capita and GSI reveals a value of 0.659 with a simulated p-value of 0.001, indicating strong and statistically significant positive spatial autocorrelation. This means that regions with similar GSI and GDP per capita values tend to be geographically clustered, rather than randomly distributed. In other words, high-GSI/high-GDP per capita regions tend to be located near each other, and the same holds for low-performing areas—suggesting that spatial spillovers, regional convergence/divergence dynamics, or shared infrastructure networks may be at play. Otherwise, the outliers can reveal significant implications worth analyzing.

Univariate LISA

The univariate LISA is applied to identify the spatial distribution of clusters based on GSI. The resulting cluster map (Figure IX) reveals four distinct cluster types. Regions with high GSI scores surrounded by similarly high GSI neighbors (High–High cluster), shown in red, form a compact concentration in the north. This includes Austria, Slovenia, northwestern and coastal Croatia, and western Hungary—areas where strong gas stations infrastructure reinforce each other across borders. Regions with low GSI scores values surrounded by high GSI neighbors (Low–High cluster), displayed in light blue, appear mainly on the periphery of these High–High concentrations, representing weaker pockets within otherwise well-performing gas stations infrastructure environments.

Figure IX

Univariate LISA clusters of GSI across regions

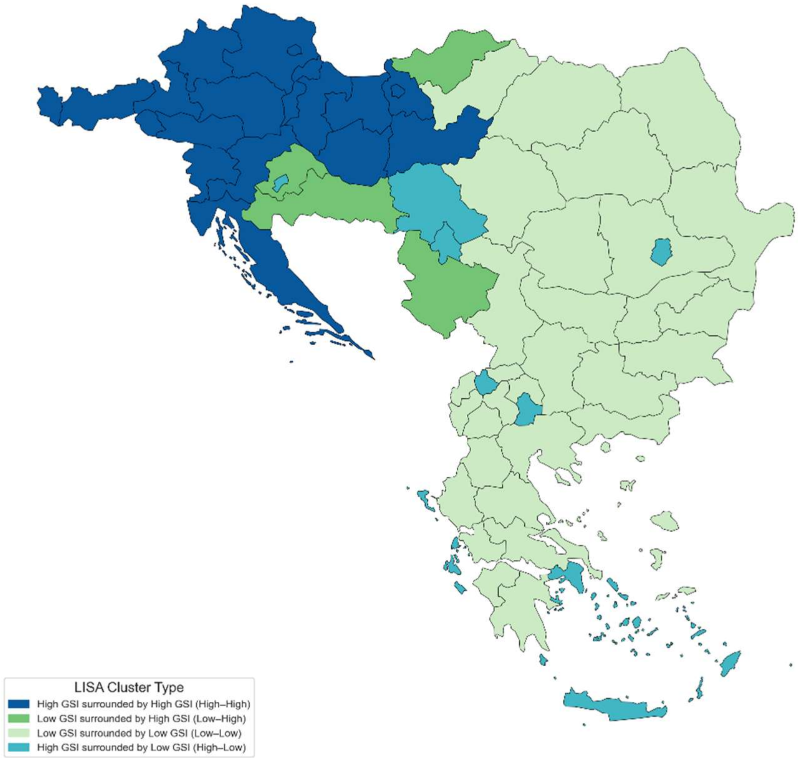

Note. Map identifies High–High, Low–Low, High–Low, and Low–High clusters, showing strong spatial dependence in GSI distribution.

A very different picture emerges in the southeast, where regions with low GSI values surrounded by low GSI neighbors (Low–Low cluster) dominate. These areas, indicated in dark blue, extend across most of Romania and Bulgaria, as well as inland and northern Greece, southeastern Serbia, and large parts of North Macedonia. Here, the spatial association reveals a pattern of cumulative disadvantage, where underdeveloped regions are clustered together, reflecting persistent structural gaps in gas station infrastructure. Finally, regions with high GSI values surrounded by low GSI neighbors (High–Low cluster), shown in orange, emerge as isolated pockets of stronger infrastructure embedded within weaker environments. These are largely concentrated around capital cities, along major transport corridors, and in tourist destinations—such as Bucharest, Belgrade, Vojvodina, North Macedonia's main highway corridor, Athens, and the Greek islands. These outliers function as regional hubs or gateways, where higher-quality gas station infrastructure provides strategic access and connectivity despite the surrounding context of relative underdevelopment.

Bivariate LISA

The bivariate LISA extends this analysis by examining the local spatial association between GSI and GDP per capita. This approach assesses whether the GSI score in a given region is correlated with the GDP per capita values in neighboring regions. The results, presented in Figure X, show that GSI has strong explanatory power for regional GDP per capita clustering. Once again, regions with high GSI scores surrounded by neighbors with high GDP per capita (High-High cluster) are concentrated in the north. The pattern of low GSI scores surrounded by neighbors with low GDP per capita (Low–Low cluster) is even more telling. Compared to the univariate analysis, their extent has expanded, now also encompassing regions that were previously classified as Low–High in Figure IX. This finding underscores that low GSI scores are not only reflective of gas stations infrastructural weakness but are also closely aligned with low GDP per capita environments, making them a powerful predictor of regional underdevelopment. At the same time, High–Low cluster, where advanced gas station infrastructure is embedded in economically weaker surroundings, become more prominent. These typically represent capital cities, transit corridors, or tourism-driven centers, where gas stations infrastructure development has outpaced the broader economic environment. Conversely, Low–High cluster, where regions with weak infrastructure are surrounded by wealthy neighbors, are substantially reduced compared to the univariate results, suggesting that the explanatory capacity of GSI for predicting low-income contexts is particularly robust.

Figure X

Bivariate LISA of GSI and GDP per capita

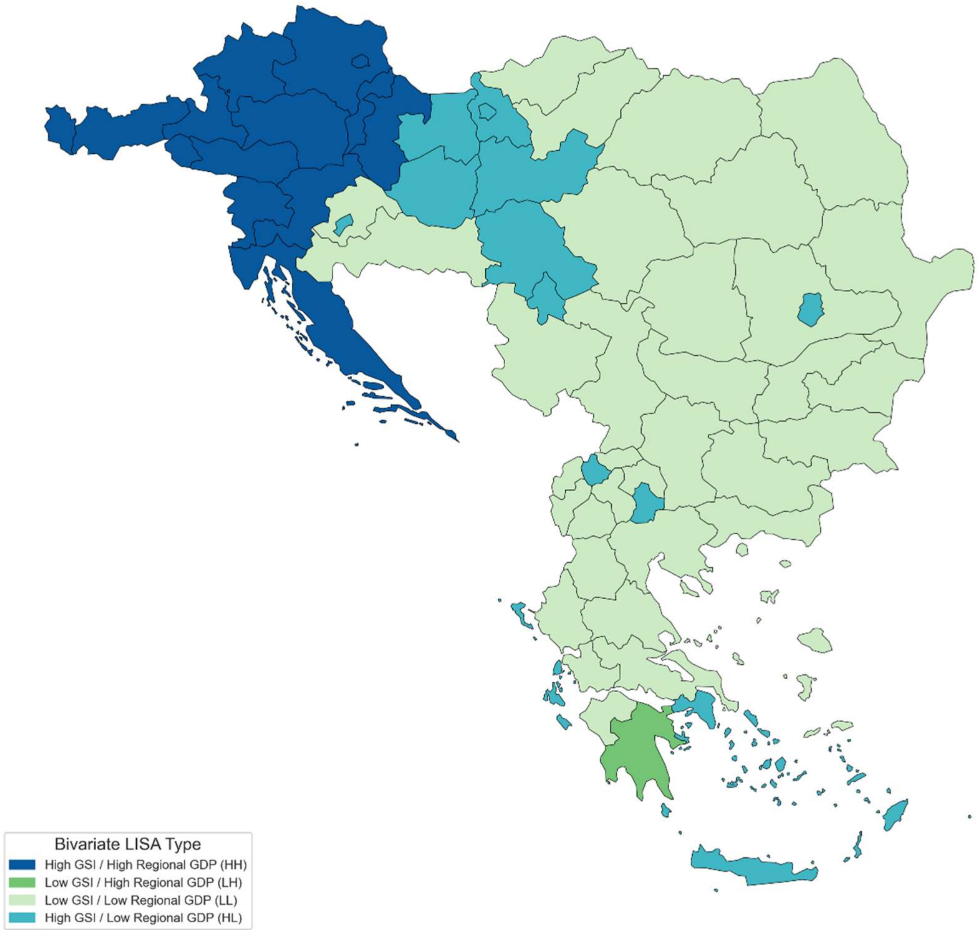

Note. Map shows local co-clustering between GSI and GDP per capita, with High–High clusters in Central Europe and Low–Low clusters in Southeast Europe.

5.4. Regional Development Patterns

Figure XI provides a direct comparison between the GSI and GDP per capita across the study regions. The regional GDP map (Figure XI, right side) shows a familiar but relatively flat topography: beyond the stark difference between Austria/Slovenia and the rest of the region, few patterns clearly stand out. GDP averages smooth over the diversity of regional industries, purchasing power, and economic trajectories. In contrast, the bivariate GSI-GDP map (Figure XI, left side) immediately reveals a richer geography of development: transit corridors, coastal hubs, and capital regions emerge sharply, while peripheral provinces are distinctly disadvantaged. This visual contrast illustrates the central argument of this section — that the GSI offers complementary and more geographically nuanced signals than GDP alone. The following subsections unpack these findings thematically, drawing on case studies from Greece, North Macedonia, Serbia, Bulgaria and Romania, Hungary, Croatia, Austria, and Slovenia.

Figure XI

The Regional Identification Power of the GSI

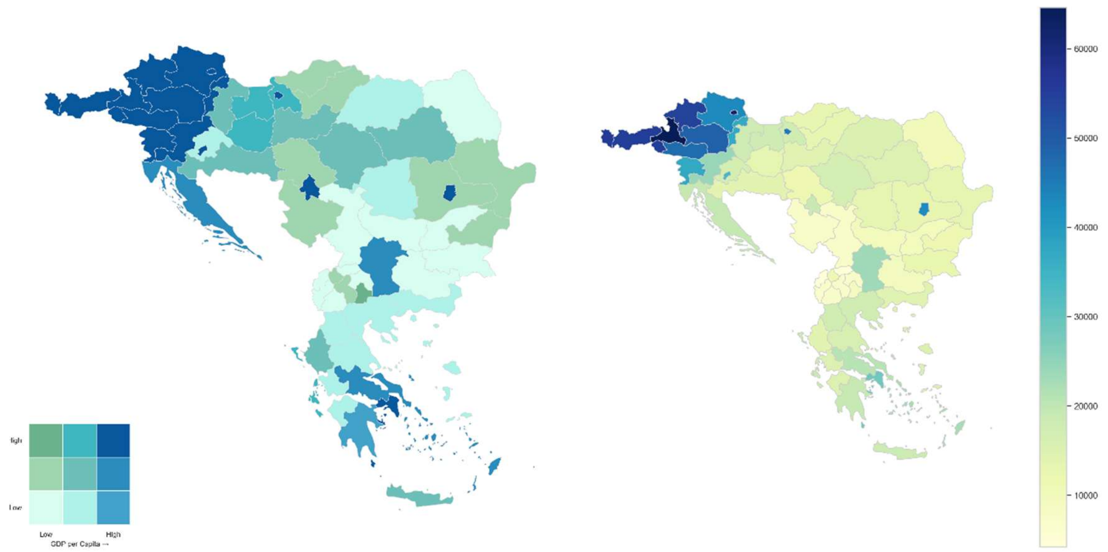

Note. Left panel shows bivariate GSI–GDP map; right panel shows GDP per capita only. GSI highlights transit corridors, capital city dominance, and peripheral underdevelopment more clearly than GDP maps alone.

While statistical models confirm the explanatory power of the GSI, its true value lies in how it maps spatial development logics onto real-world geography. The following subsections illustrate the different ways in which infrastructure patterns, as captured by the GSI, reveal persistent inequalities, growth opportunities, and governance outcomes across the case study regions.

Corridors and Transit Economies

One of the strongest signals in the GSI maps emerges along major transit arteries. Highways, cross-border routes, and industrial corridors create concentrated pockets of infrastructure that GDP alone cannot isolate. Regions along the Pan-European Corridor X in North Macedonia are perfectly outlined as it represents the nation's only major and developed trade and transit corridor. This highway is the main corridor connecting Greece to Central Europe and the regions in North Macedonia heavily benefit from cross-border trade and foreign investments. Corridor X regions in Serbia, Hungary and Croatia are also identified by the index, even though they are lacking in GDP in comparison to some other regions in the countries. Western Hungary's industrial belt and infrastructure heavy regions connecting to Austria are also adequately evaluated. This illustrates how strategic geography related to trade infrastructure and/or crossborder economic spillover, translates directly into visible infrastructure intensity, often preceding formal income growth, that is expected to equalize and provide returns in the future.

Tourism and Port Economies

Tourism and port regions show another distinct footprint in the GSI. Coastal areas and islands often appear overdeveloped in infrastructure and services density relative to their GDP per capita, reflecting the weight of seasonal demand, international mobility, and port-based commerce. Between the case study countries this is especially outlined in the regions of the Greek islands and coastline, that are tourism and/or port heavy, Varna in Bulgaria, Constanza in Romania, and the Croatian coastline. The GSI captures how ports and tourist hubs act as gateways of prosperity, sustaining dense service networks even when national accounts understate their economic importance. This isn’t the case in maps of regional GDP and average wages.

Capital City Dominance

The GSI reinforces a classic pattern of capital city bias, where national capitals hoard infrastructure and services far beyond what regional GDP maps suggest. Athens, Sofia, Belgrade, Zagreb and Budapest dominate their national landscapes, illustrating how political centralization and concentrated investment flows entrench spatial inequalities. Unlike GDP which partially identifies this, the GSI makes this bias immediately visible through everyday infrastructure.

Austria and Slovenia stand out as exceptions to the rule. Here, the GSI maps show relatively balanced regional performance, with no single city dramatically overshadowing the rest. This more equitable distribution underscores the impact of cohesive national policies and effective EU integration, offering a counter-model to the overcentralized structures of the post-crisis economies, the transitional economies of the Western Balkans and post Eastern Block countries

Peripheries and Stagnating Regions

Finally, the GSI highlights the “forgotten geographies”: depopulating eastern Hungary, the Romanian North-East region, peripheral Bulgaria, the infrastructure lacking regions of North Macedonia, inland Greece, and southeastern Serbia. These low-GSI, low-GDP regions form cumulative disadvantage clusters, where infrastructure neglect reinforces weak economic performance. By identifying these lagging areas with precision, the GSI functions as an early-warning indicator of regional stagnation, signaling where policy attention and investment are most urgently needed. Also through the identification of these regions, they may be regions where future infrastructure and service investments may yield the highest returns as confirmed in subsection 5.2. where diminishing returns are discussed.

Moreover, the GSI identifies regions where infrastructure development lags behind GDP per capita. For example, the northern regions of Greece (Dytiki Makedonia EL53, Anatoliki Makedonia and Thraki EL51, and Kentriki Makedonia EL52) record several times higher GDP per capita than neighboring regions in North Macedonia, yet their GSI scores are lower. This pattern suggests that such regions could strongly benefit from investments in infrastructure.

Overall, the case studies show that the GSI captures the structural forces that shape regional prosperity: the effects of centralized policymaking and capital dominance, the boost from transit corridors and trade flows, and the concentration of services in tourism and port economies. These patterns demonstrate that the GSI identifies real economic asymmetries that GDP alone often conceals. More importantly, the index highlights where growth potential is emerging, where regions are stagnating, and where policy interventions are most urgent. In this sense, the GSI should not be viewed only as a statistical proxy for GDP, but as a practical tool for understanding development dynamics and informing future economic policy. Beyond its value for governments and policymakers, the GSI also serves as a practical tool for private actors, from investment banks and infrastructure funds to multinational corporations, who seek fast, intuitive assessments of regional vitality. For example, a foreign investor evaluating new projects in Southeast Europe would gain a clearer understanding of local purchasing power, service intensity, and infrastructure quality from a single glance at the GSI than from headline GDP figures alone.

5.5. Robustness Checks

In this subsection, we evaluate the stability of our results through a series of robustness checks. Specifically, we: (i) re-estimate the models using a logarithmic transformation of regional GDP values, (ii) conduct a country subsample analysis using Hungary and Croatia only, (iii) implement a five-fold cross-validation procedure, and (iv) note that in the construction of the GSI, winsorization at the 5th and 95th percentiles was applied, effectively addressing the influence of extreme values.

Using the logarithm of regional GDP per capita as the dependent variable, both the linear and random forest regressions remain highly explanatory. The linear model achieves an R^2 of 0.636 with a cross-validated R^2 of 0.571 (Std: 0.090), while the random forest yields an in-sample R^2 of 0.942 and a cross-validated R^2 of 0.471 (Std: 0.159). These results confirm that the relationship between GSI and economic output is not dependent on the scale of the dependent variable.

Restricting the analysis to Hungary and Croatia (12 regions), the GSI continues to show strong explanatory power. The linear model achieves an in-sample R^2 of 0.696, while the random forest reaches 0.874. This indicates that the predictive capacity of the GSI holds even in smaller, more homogeneous regional samples.

Finally, we apply a five-fold cross-validation procedure to the linear regression, random forest, and XGBoost models. By averaging R^2 scores across multiple training and testing partitions, this approach minimizes the influence of any single data split and provides a more robust assessment of predictive performance. These values are reported above in the text.

Together, these robustness checks confirm that the explanatory power of the Gas Station Index is stable across model specifications, transformations, and regional subsamples, underscoring its validity as a proxy for regional prosperity.

6. Conclusion

This paper has introduced the Gas Station Index (GSI) as an infrastructure-based proxy for regional economic development. Using geospatial data on 19,033 stations across 62 regions in nine European countries, we show that the GSI alone explains 64 percent of the cross-sectional variation in GDP per capita. This explanatory power is notable for a single proxy variable and demonstrates that consumer-facing infrastructure encodes systematic information about regional economic conditions. At the same time, the unexplained variation underscores the structural difference of the two measurements of economic activity - ranging from sectoral specialization to mobility patterns—that are outside the scope of an infrastructure-based index, and the informal economic activity which is outside the scope of GDP. The GSI should therefore be viewed not as a substitute for national accounts, but as a complementary measure that highlights margins of development particularly relevant at the subnational level.

The findings of this study have important implications for both policy and practice. For governments and regional authorities, the GSI provides a practical tool to identify infrastructure deficits, highlight transit corridors and tourism hubs as drivers of growth, and prioritize investment in regions where neglect risks entrenching stagnation. For private actors such as investors, infrastructure funds, and development banks, the index offers a cost-efficient and real-time diagnostic of local purchasing power, service density, and accessibility—factors often obscured in headline GDP statistics. Nevertheless, while the GSI demonstrates robust explanatory power across diverse regional contexts, it remains partly constrained by open source data quality bias, which can limit comparability across countries.

A key limitation lies in the temporal and spatial coverage of data. While OpenStreetMap (OSM) provides strong contemporary snapshots, it offers limited historical depth, making it difficult to backcast reliable GSI scores for earlier decades. Equally important, OSM coverage is not uniform across the globe: countries outside Europe and North America are often undermapped, with fewer contributors and less complete metadata. This creates a quality bias that may understate infrastructure development in precisely those regions where alternative measures are most needed. As with other proxies such as night-lights, these constraints can be mitigated through mass-participation data-collection efforts. Expanding crowdsourced contributions to metadata (fuel types, service quality, amenities) would improve comparability, enhance reliability, and allow the GSI to evolve as a dynamic, living infrastructure indicator.

Future research should expand the GSI both geographically and temporally. Building a longitudinal (panel) version of the index would materially strengthen three dimensions of the framework: inference, prediction, and measurement. First, a time-series GSI enables within-region identification using fixed-effects, event-study, and difference-in-differences designs around shocks such as highway openings, station entry/exit, or rebrandings, helping separate persistent regional traits from true changes in infrastructure and services. Second, it would improve nowcasting and forecasting: by testing lead–lag structures (e.g., Granger causality, cross-lag regressions), one can evaluate whether changes in GSI precede recorded regional GDP per capita and detect turning points earlier than official statistics. Third, repeated measurement mitigates data-quality drift in open data (e.g., growth in mapping intensity): panel normalization (region-year baselines, vintage controls, seasonal filters for tourism regions) reduces bias and raises cross-time comparability. Many African and Asian countries either do not publish subnational GDP or release it with long lags and uneven quality; in such contexts, a near-real-time, panel GSI could provide the first consistent proxy for tracking regional prosperity as it evolves, not just as a static snapshot. At the other end of the spectrum—advanced economies with infrastructure saturation—a temporal GSI can be complemented by layers that move where GSI plateaus: innovation activity, environmental sustainability, and digital infrastructure. As global coverage expands, adopting threshold-based scoring (analogous to HDI) and dynamic benchmarks would standardize interpretation across contexts and over time, while preserving sensitivity to meaningful changes within regions. Taken together, these steps would position the GSI as a scalable, policy-relevant leading indicator that bridges academic research, investment decisions, and regional development management.

References

Benmelech, Efraim, Jun Yang, and Michal Zator. "Bank Branch Density and Bank Runs." *NBER Working Paper* No. 31462, National Bureau of Economic Research, 2023. https://doi.org/10.3386/w31462.

Calderón, César, and Luis Servén. *The Effects of Infrastructure Development on Growth and Income Distribution*. Policy Research Working Paper 3400, Washington, DC: World Bank, 2004. https://doi.org/10.1596/1813-9450-3400.

Clifton, Nick, Anita Füzi, and Gareth Loudon. "Coworking in the Digital Economy: Context, Motivations, and Outcomes." *Futures*, 135 (2022), 102439. https://doi.org/10.1016/j.futures.2019.102439.

Fujita, Masahisa, Paul Krugman, and Anthony J. Venables. *The Spatial Economy: Cities, Regions, and International Trade*. Cambridge, MA: MIT Press, 1999.

Gauger, Felix, Andreas Pfnür, and Jan-Oliver Strych. "Coworking Spaces and Start-Ups: Empirical Evidence from a Product Market Competition and Life Cycle Perspective." *Journal of Business Research*, 132 (2021), 67–78. https://doi.org/10.1016/j.jbusres.2021.04.008.

Gibson, John, Susan Olivia, and Geua Boe-Gibson. "Night Lights in Economics: Sources and Uses." *Journal of Economic Surveys*, 34 (2020), 955–980. https://doi.org/10.1111/joes.12387.

Glaeser, Edward L. *Cities, Agglomeration, and Spatial Equilibrium*. Oxford: Oxford University Press, 2008.

Goodwin, Phil, Joyce Dargay, and Mark Hanly. "Elasticities of Road Traffic and Fuel Consumption with Respect to Price and Income: A Review." *Transport Reviews*, 24 (2004), 275–292. https://doi.org/10.1080/0144164042000181725.

Henderson, J. Vernon, Adam Storeygard, and David N. Weil. "Measuring Economic Growth from Outer Space." *American Economic Review*, 102 (2012), 994–1028. https://doi.org/10.1257/aer.102.2.994.

Henderson, J. Vernon, Tim Squires, Adam Storeygard, and David N. Weil. "The Global Distribution of Economic Activity: Nature, History, and the Role of Trade." *The Quarterly Journal of Economics*, 133 (2018), 357–406. https://doi.org/10.1093/qje/qjx030

Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi. "Growth and Structural Transformation." In *Handbook of Economic Growth*, Vol. 2B, Philippe Aghion and Steven N. Durlauf, eds. Amsterdam: Elsevier, 2014, 855–941. https://doi.org/10.1016/B978-0-444-53540-5.00006-9.

Kireyeva, Akerke A., Alexandra Kredina, László Vasa, and Zhanar T. Satpayeva. "Impact of Financial Technologies on Economic Development: Theories, Methods and Analysis." *Journal of International Studies*, 14 (2021), 286–303. https://doi.org/10.14254/2071-8330.2021/14-4/19.

Pinkovskiy, Maxim, and Xavier Sala-i-Martin. "Lights, Camera … Income! Illuminating the National Accounts–Household Surveys Debate." *The Quarterly Journal of Economics*, 131 (2016), 579–631. https://doi.org/10.1093/qje/qjw003

Steckel, Richard H., and Jerome C. Rose, eds. *The Backbone of History: Health and Nutrition in the Western Hemisphere*. Cambridge: Cambridge University Press, 2002. https://doi.org/10.1017/CBO9780511549953.

Sutton, Paul C., Christopher D. Elvidge, and Tilottama Ghosh. "Estimation of Gross Domestic Product at Sub-National Scales Using Nighttime Satellite Imagery." *International Journal of Ecological Economics & Statistics*, 8 (2007), 5–21.

Zhu, Jie, Mengyao Zhu, Li Chen, Li Luo, Weihua Wang, Xueming Zhu, and Yizhong Sun. “Multi-Objective Optimization of Urban Gas Station Site Selection Under Territorial Spatial Planning Constraints.” *ISPRS International Journal of Geo-Information*, 13 (2024), 375. https://doi.org/10.3390/ijgi13110375.