Unsupervised Anomaly Detection in Energy Generation Time Series Using an LSTM Autoencoder with Spatiotemporal Analysis of Anomalies

Florian Sebastian Kienhoefer, Johannes Scholz

Paris Lodron University of Salzburg, Schillerstraße 30, 5020 Salzburg,
floriansebastian.kienhoefer@plus.ac.at, zgis.at

Abstract: The energy generated by photovoltaic systems is recorded as time series data, where technical issues may generate errors. Prior work shows that machine learning models can detect such anomalies; however, most studies neglect the spatial and temporal patterns of anomalies, risking the loss of crucial information about their origin and spread. We propose a two-stage method to detect and interpret these anomalies using machine learning and spatial statistics. First, we employ an LSTM autoencoder to detect anomalous sequences; second, we apply a Getis–Ord Gi* analysis to assess the spatial clustering of detected anomalies. Our results show that the method reliably detects anomalies and identifies significant anomaly hotspots that may indicate localised problems. These findings enable energy grid operators to prioritise inspections where anomalies are spatially clustered, isolate their likely origin, and accelerate the diagnosis of grid issues.

Keywords: anomaly detection, lstm autoencoder, spatial statistics, energy time series

1 Introduction

The energy generated by photovoltaic (PV) systems can be recorded as time series via smart meters [1], where technical issues may generate errors. When modelling energy systems, we are tasked with detecting abnormal patterns in data, referred to as anomaly detection (AD) [2]. For example, these anomalies can be caused by smart meter malfunctions, operational errors, fraud [3] or by interpolated values during communication interruptions [4].

A key limitation of existing anomaly detection methods is their lack of analysis of the spatial and temporal patterns of anomalies [5]. Hence, important information about the origin and spread of an anomaly may not be detected. In addition, not having this information can delay the identification and resolution of grid problems.

To address this research gap, we propose a method that uses machine learning to detect abnormal behaviour in PV generation data. We then apply spatial statistics to analyse whether the detected anomalies are spatiotemporally clustered, indicating localised issues. We demonstrate our method in a real-world case study.

In this paper, we formulate the following research question: How can machine learning and spatial statistical methods be combined to detect anomalies in energy generation time series and assess spatiotemporal clusters?

This paper contributes to a more holistic understanding of problems that affect multiple PV systems, indicating systemic issues in the energy grid, equipment failures, or localised disruptions.

2 Background

2.1 Anomalies in Energy Generation Data

The identification of anomalies in time series is critical for the early detection and resolution of system errors [5] to avoid energy losses and ensure consistent network stability [6]. Smart meters are calibrated devices and their measured energy generation should be accurate; however, communication errors between an electricity provider and smart meters can lead to missing values, which have to be estimated if the interruption persists for an extended period [4]. During these periods, abnormal measurements may occur in the data, which should be detected. Further causes of anomalies in smart meter data include meter malfunctions, operational errors, and fraud [3]. Unidentified anomalies can lead to billing errors [4] and should not be included in energy forecasts, as they would distort them [1].

2.2 Machine Learning for Anomaly Detection in Energy Generation Data

Anomaly detection is an active research field [7] and is gaining popularity through the emergence of machine learning-based approaches. While many unsupervised machine learning methods are successfully applied to the task of detecting anomalies in energy data [1], [3], [8], the analysis of the spatiotemporal distribution of these anomalies remains underexplored [5]. For example, Pereira and Silveira [8] proposed an unsupervised machine learning method based on an autoencoder to detect anomalies in energy time series. Their findings indicate that their model can effectively detect anomalies in time series by learning meaningful representations and using a probabilistic reconstruction assessment that takes the variability in the data into account. Further, the model can successfully learn the seasonal patterns in the energy data in its latent space for normal compared with abnormal sequences.

2.3 Getis–Ord Gi* Statistic to Analyse the Spatial Distribution of Anomalies

Njungwi et al. [9] extracted features from satellite imagery that showed the thermal infrared radiation of the Earth's surface, radiation emitted by, for example, asphalt, buildings, and vegetation, after being heated by the sun. These values were analysed for anomalies using the Getis–Ord Gi* statistic. This spatial analysis showed that the critical positive thermal anomalies in the Busan metropolitan area spatially increased from approximately 20% to over 24% from 2000 to 2020, indicating that urban planning needed to become more sustainable.

2.4 Summary

These findings confirm the potential of autoencoders to process energy generation time series and learn their distinct curve trends caused by factors such as weather effects or seasonal variations, and the capability of the Getis–Ord Gi* statistic to analyse the spatial distribution of anomalies.

3 Case Study: Anomaly Detection in Energy Generation Data with Spatiotemporal Analysis of Anomalies

The case study relies on a two-stage approach and uses machine learning for anomaly detection and spatial statistics to assess the spatial distribution of the detected anomalies.

3.1 Case Study Dataset

The case study uses real-world energy generation data provided by an Austrian energy distribution system operator to evaluate our proposed method. The dataset includes time series of 2,042 smart meters over a span of 52 weeks, with each timestamp representing the energy generation within the past 15 minutes. The meters are connected to 52 substations across Styria. We train the model on data from all 52 substations; however, for the analysis of spatial clusters within the detected anomalies, we use 46 substations due to missing location information. Substation locations are taken as the centroids of their polygon geometries.

3.2 Anomaly Detection Task

First, we utilise a machine learning approach to detect the anomalies within the real-world data.

3.2.1 Data Preprocessing

To prepare the data for the training process of the machine learning algorithm, several preprocessing steps are performed. First, smart meters with over 95% zero values or minimal variance defined by a coefficient of variation of less than 0.05 are excluded. Second, outliers are individually determined for each smart meter using a z-score analysis with a threshold of 5; if the number of outliers exceeds 2% of all data points of a smart meter, the smart meter is excluded. If the error rate is below this threshold, the peak values are capped to a z-score of 5. Third, z-score normalisation is used to bring all smart meters to a common – i.e. comparable – scale. The normalisation is performed separately for each smart meter to ensure that each has an average of 0 and a standard deviation of 1. Finally, sliding windows are created, in which 15-minute timestamps from the smart meters are individually aggregated into 24-hour windows (96 timestamps) with a 12-hour overlap (48 timestamps), resulting in 755 sliding windows per smart meter.

3.2.2 LSTM Autoencoder

For the anomaly detection within the preprocessed time series, we use an LSTM autoencoder [10]. Autoencoders consist of two parts: encoders that map the input data to a latent space, which is a compressed representation of it, and decoders that map back from the latent space to the input space. In anomaly detection tasks, autoencoders focus on normal patterns in the time series and try to reconstruct them as accurately as possible to the input sequence through minimising a loss function that measures the reconstruction error [8], [10]. Anomalies are usually rare in training data; therefore, the model can reconstruct normal patterns well but struggles with data that contain anomalies. As a result, if the reconstruction error is high, it indicates the presence of an anomaly. Autoencoders are well suited for our task as they are unsupervised and do not require labelled input data for training [11]. Further, they can perform pattern-based anomaly detection, which is ideal for learning the complex patterns in energy generation data.

3.2.3 Experimental Setup

The LSTM autoencoder in this work is implemented using the PyTorch library [12] and is divided into an encoder–decoder structure.

The encoder uses an LSTM layer (input_size=1, hidden_size=64) to transform the input sequences into a high-dimensional hidden state. To force an efficient feature representation, the last hidden state of the encoder is compressed into a latent vector of size 32 using a linear layer (nn.Linear). This vector serves as a bottleneck and contains only the most essential information from the original input.

The decoder mirrors the encoder. The latent vector (size 32) is transformed back to the original sequence length and mapped to the hidden dimension (64) via a linear layer. A decoder LSTM (input_size=64, hidden_size=64) processes this sequence, and a final linear layer maps each time step from 64 to 1 to reconstruct the input.

The model is trained for 28 epochs using the Adam optimiser (learning rate 0.001) with a batch size of 128. Mean Squared Error (MSE) is used as loss function; it measures the deviation between the original sequence and the reconstruction. The data are split into training and test sets in an 80/20 ratio. Cross-validation on the training set is used to select hyperparameters, and a final evaluation is performed on the held-out test set; these evaluations indicate good generalisation. Anomaly detection is then applied to the entire dataset. By minimising the reconstruction error, the model learns to accurately predict normal patterns, which allows significant deviations to be identified as anomalies. Since ground truth labels are unavailable, we do not report supervised detection metrics. Instead, domain experts qualitatively validate detected anomalies in the time series for plausibility. Based on this review, we use a global 98th-percentile threshold for the entire dataset; any value exceeding this limit is classified as an anomaly, as this setting yields the best results across all smart meters.

3.3 Spatial Analysis Task

In the second step of our method, we use spatial statistics to analyse the detected anomalies for spatial clusters.

3.3.1 Getis–Ord Gi* Statistic

For the analysis of spatial anomaly clusters, we employ the Getis–Ord Gi* statistic [13] since it allows us to statistically validate observations that would otherwise rely on visual inspection of spatiotemporal patterns. Getis–Ord Gi* identifies hotspots where substations with an unusually high number of anomalies are located close together. This suggests a localised, non-systemic issue with a common cause, rather than a general system-wide problem in the area. It also identifies coldspots, where substations with an unusually low number of anomalies are spatially clustered. This indicates a stable area. Further, it distinguishes these hotspots and coldspots from individual substations with high or low anomaly counts that are isolated coincidences, not reflected in nearby stations and therefore do not form significant clusters.

3.3.2 Experimental Setup

As a first step before the Getis–Ord Gi* analysis is carried out, a neighbourhood configuration is defined. For this, the k-nearest neighbours approach from the Python library libpysal [14] (libpysal.weights.KNN, k=4) is used. On the basis of Euclidean distances in a metric coordinate system (UTM, EPSG:32633), a spatial weights matrix is created, which assigns each substation its four closest neighbours and the weights are row-standardised. Furthermore, the anomaly counts are normalised by the total number of smart meters per substation to calculate a local anomaly rate, ensuring the results reflect spatial intensity rather than substation size. To include each substation in its own local calculation, the spatial weights are adjusted to incorporate self-influence. The subsequent Getis–Ord Gi* analysis [13] is carried out using the Python implementation esda.getisord.G_Local [14] with the previously defined libpysal weights to identify spatial patterns both weekly and over the entire data period. Statistical significance is determined using a p-value threshold of 0.05 and 95% confidence intervals; specifically, a substation is classified as a hotspot if it exhibits a positive z-score greater than 1.96 and as a coldspot if it exhibits a negative z-score less than -1.96.

Figure 1 shows the spatial neighbourhood of the substations and their connected neighbours. Arrows indicate directed connections between nodes.

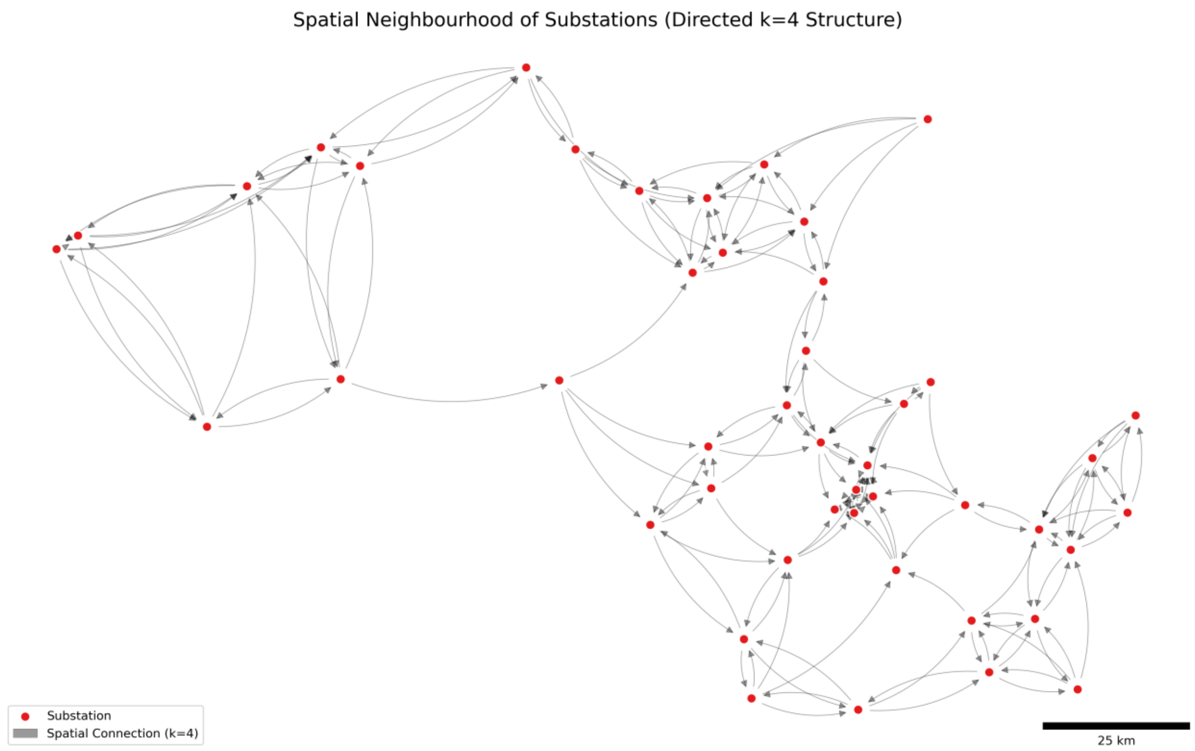

Figure 1: Spatial Neighbourhood Map: The graph illustrates the directed k=4 structure used for the analysis. Grey arrows represent connections between each substation and its four nearest neighbours.

4 Results and Discussion

The analysis of the detected anomalies reveals that it is possible to find spatial patterns in our case study data and to locate substations that form significant hotspots.

Figure 2 shows 46 substations in Styria over a 52-week period, with colours indicating the fraction of weeks each substation is a hotspot or coldspot. Larger coloured dots represent stations that are significant in at least one week, while smaller white dots indicate stations that are never significant over the measured timespan. As shown in the persistence map, no stations reach the dark red threshold for persistent hotspots, though several enlarged dots exhibit a light red colouration, indicating they are identified as frequent hotspots over the year. No coldspots are identified that would be marked in blue. Most stations appear as small, white markers, signifying they remain stable throughout the year.

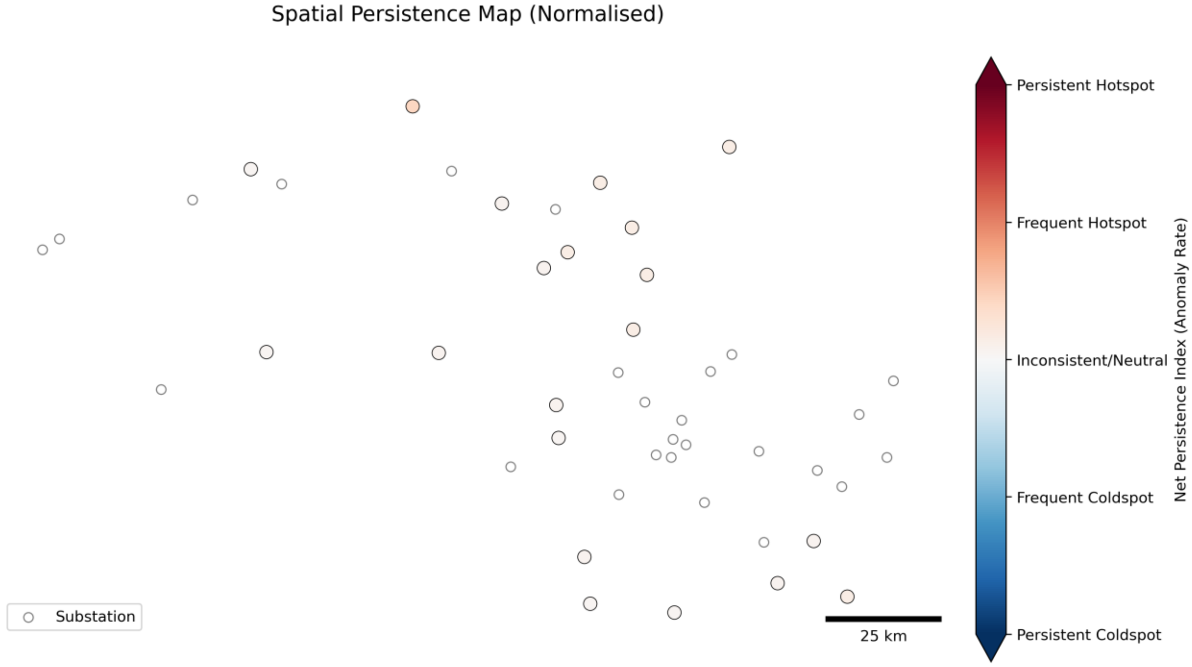

Figure 2: Spatial Persistence Map: The map shows the spatial persistence of 46 substations in Styria over 52 weeks: larger coloured dots are significant in at least one week; smaller white dots are never significant. Light-red hotspots are observed; no blue coldspots are observed.

Figure 3 displays hotspot/coldspot clusters for week 5. Larger dots represent statistically significant stations where a random accumulation of values is unlikely. Smaller white dots indicate stations that are not statistically significant. As shown in the plot, four large red hotspots are identified, indicating significant clusters of high anomaly rates. Conversely, no large blue dots appear, as no significant coldspots are found.

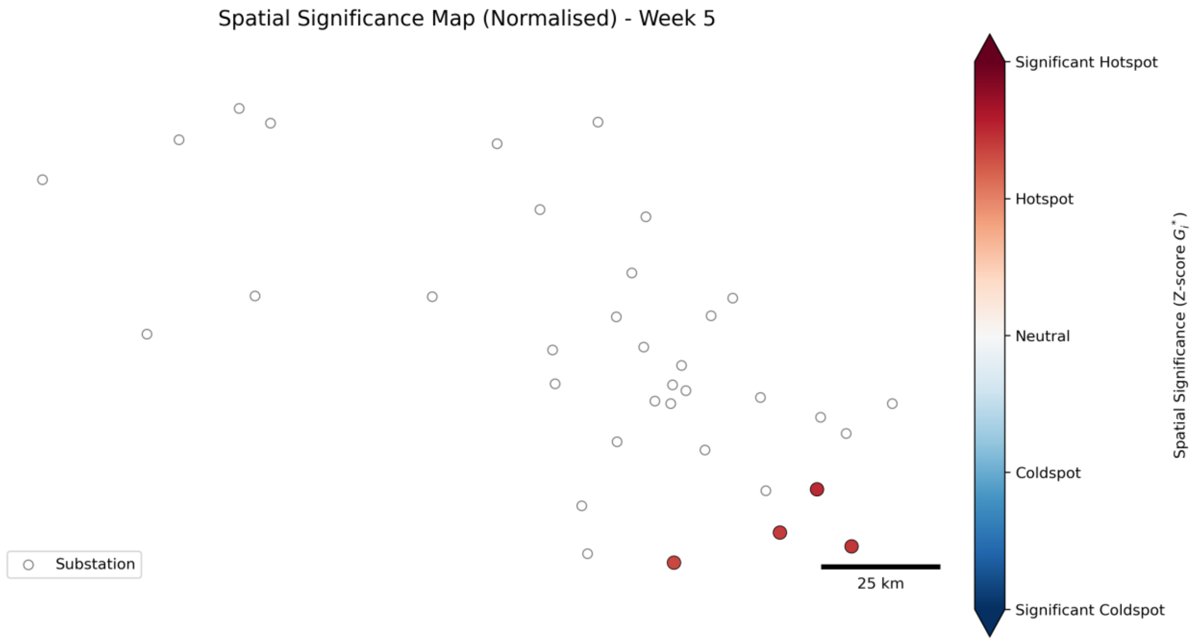

Figure 3: Spatial Significance Map: The map shows the spatial significance for week 5: larger coloured dots mark statistically significant clusters; smaller white dots are not significant. Four large red hotspots are observed; no blue coldspots are observed.

5 Conclusion

In this work, we show that the combination of LSTM autoencoders with spatial statistics (Getis–Ord Gi*) effectively detects anomalies in energy generation data and generates an interpretable spatial pattern. Our findings are backed by a real-world case study in which we identify substations that show significant spatial anomaly clusters with neighbouring substations. These clusters may indicate localised issues and provide a good starting point for further analysis to identify underlying causes. In summary, the proposed methodology has great potential to provide energy grid operators with actionable insights for prioritising inspections where anomalies are spatially clustered, isolating their likely origin, and accelerating the diagnosis of grid issues. The method can also be applied in other contexts where anomalies must be detected in time series and analysed for spatial patterns.

We propose several directions to expand upon this study. Our autoencoder solely models reconstruction-based deviations based on time series and does not take external variables into account. Future work should incorporate additional variables such as weather and operational data to understand anomaly causes more clearly. Further, it would be interesting to model the energy generation data as a spatiotemporal heterogeneous graph, where nodes represent substations and smart meters and edges encode the distance or physical connectivity between them, and to apply graph representation learning algorithms. This might be a more natural way to represent energy grid infrastructure and may yield insights that are potentially missed when using a non-graph representation.

6 Acknowledgements

This work is part of the project iKlimEt (FO999910627), which has received funding in the framework of “Energieforschung”, a research and technology programme of the Klima- und Energiefonds.

7 References

[1] M. Turowski et al., “Managing anomalies in energy time series for automated forecasting,” Lecture Notes in Computer Science, pp. 3–29, Dec. 2023, doi: 10.1007/978-3-031-48649-4_1.
[2] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly detection: A survey,” ACM Computing Surveys, vol. 41, no. 3, pp. 1–58, Jul. 2009, doi: 10.1145/1541880.1541882.
[3] Y. Shi, L. Xu, P. Wang, X. Wu, L. Wang, and Y. Hou, “Real-time anomaly detection of electricity time series data based on future-guidance network,” IEEE Signal Processing Letters, vol. 32, pp. 3660–3664, 2025, doi: 10.1109/LSP.2025.3598707.
[4] V. Gulisano, M. Almgren, and M. Papatriantafilou, “Online and scalable data validation in advanced metering infrastructures,” in Proc. IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe), Oct. 2014, pp. 1–6, doi: 10.1109/ISGTEurope.2014.7028740.
[5] S. Guan, B. Zhao, Z. Dong, M. Gao, and Z. He, “GTAD: Graph and temporal neural network for multivariate time series anomaly detection,” Entropy, vol. 24, no. 6, p. 759, May 2022, doi: 10.3390/e24060759.
[6] S. S. S. R. Depuru, L. Wang, and V. Devabhaktuni, “Smart meters for power grid: Challenges, issues, advantages and status,” Renewable and Sustainable Energy Reviews, vol. 15, no. 6, pp. 2736–2742, Aug. 2011, doi: 10.1016/j.rser.2011.02.039.
[7] A. Staffini, T. Svensson, U. Chung, and A. K. Svensson, “A disentangled VAE-BiLSTM model for heart rate anomaly detection,” Bioengineering, vol. 10, no. 6, Art. no. 683, Jun. 2023, doi: 10.3390/bioengineering10060683.
[8] J. Pereira and M. Silveira, “Unsupervised anomaly detection in energy time series data using variational recurrent autoencoders with attention,” in Proc. 2018 17th IEEE Int. Conf. on Machine Learning and Applications (ICMLA), Dec. 2018, pp. 1275–1282, doi: 10.1109/ICMLA.2018.00207.
[9] N. W. Njungwi, D. Lee, M.-J. Kim, C. Jin, and C. Choi, “Analyzing the evolution of summer thermal anomalies in Busan using remote sensing and spatial statistical tool,” Korean Journal of Remote Sensing, vol. 37, no. 4, pp. 665–685, Jan. 2021, doi: 10.7780/kjrs.2021.37.4.1.
[10] P. Malhotra, A. Ramakrishnan, G. Anand, L. Vig, P. Agarwal, and G. Shroff, “LSTM-based encoder-decoder for multi-sensor anomaly detection,” arXiv:1607.00148, Jul. 2016.
[11] O. I. Provotar, Y. M. Linder, and M. M. Veres, “Unsupervised anomaly detection in time series using LSTM based autoencoders,” in Proc. 2019 IEEE Int. Conf. on Advanced Trends in Information Theory (ATIT), Dec. 2019, pp. 513–517, doi: 10.1109/ATIT49449.2019.9030505.
[12] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," in Advances in Neural Information Processing Systems 32, 2019, pp. 8024–8035.
[13] A. Getis and J. K. Ord, “The analysis of spatial association by use of distance statistics,” Geographical Analysis, vol. 24, no. 3, pp. 189–206, Jul. 1992, doi: 10.1111/j.1538-4632.1992.tb00261.x.
[14] S. J. Rey and L. Anselin, “PySAL: A Python Library of Spatial Analytical Methods,” Review of Regional Studies, vol. 37, pp. 5–27, 2007.