Using Satellite Imagery to Map Poverty Struck Areas in Pakistan Using Neural Networks

Hamza Umer Farooq, Muhammad Junaid Nazar,
Muhammad Sarib Hanif and Aamna Faisal

July 27, 2024

Using Satellite Imagery To Map Poverty Struck Areas of Pakistan Using Neural Networks

Hamza Umer Farooq
Computer Science
Air University
Islamabad , Pakistan
hamzaumer@protonmail.com

Muhammad Junaid Nazar
Faculty of Computing & AI
Air University
Islamabad , Pakistan
junaid.nazar@au.edu.pk

Muhammad Sarib Hanif
Computer Science
Air University
Islamabad , Pakistan
saribraja1998@gmail.com

Aamna Faisal
Computer Science
Air University
Islamabad , Pakistan
aamnafaisal1@gmail.com

***Abstract*— Poverty is a complicated socioeconomic issue that comprises of more than just financial difficulty. It also includes inadequate healthcare, education, and basic housing. Satellite imaging has become a powerful tool for studying socioeconomic trends, especially in the areas where poverty is a problem. Making use of this potential, this research aims to provide a solid model for mapping poverty in Pakistan, which would facilitate resource allocation and decision-making. The absence of reliable data makes it difficult to accurately measure Pakistan's poverty levels, even with advances in satellite technology. Satellite imagery-trained convolutional neural networks (CNNs & ANNs) are becoming one of the most widely used and successful methods. ResNet50, ResNet101 and Yolo are deep convolutional neural network (CNN) architectures well-known for their ability to train deeper networks effectively by utilizing residual blocks. Our goal is to increase efficiency and performance of the system by using ResNet and Yolo models for both training and efficiency comparison. These methods represent a significant advancement in satellite image processing and provide enhanced capabilities for resource allocation in poverty reduction programs and evidence-based decision-making. These maps facilitate policymakers, researchers, and NGOs by providing insightful information. These results shows that the proposed scheme is effective in creating poverty maps with high accuracy (75.5%) and precision (76.2).

Keywords—Poverty Mapping, Satellite Imagery, Socioeconomic Analysis, Convolutional Neural Networks (CNNs), ResNet50, ResNet101, Evidence-Based Decision-Making, Poverty Reduction Programs, Accuracy and Precision in Poverty Assessment.

I. INTRODUCTION

Poverty is still a complex socioeconomic problem that goes beyond simple financial difficulties to include deficiencies in basic housing, healthcare, and education. For resource allocation and policymaking to be effective, especially in emerging nations like Pakistan, accurate and current poverty mapping is crucial. Due to their antiquated nature and high expense, traditional methods like door-to-door surveys and census data frequently prove ineffective. Recent years have seen the emergence of satellite images along with sophisticated machine learning algorithms as a powerful tool for tackling these issues. The present study employs convolutional neural networks (CNNs) and satellite data to construct a resilient model for poverty mapping in Pakistan. Poverty assessment accuracy can be increased using this method along with GEE-Sentinel Data alongside strong and trustworthy socioeconomic analysis [1] This would improve the precision and efficacy of poverty assessment and enable focused intervention approaches.

High-resolution data that can capture the subtle differences in socioeconomic situations are essential for accurate poverty mapping. We use high-resolution satellite imagery that is available through Google Earth Engine from Planet.com and Sentinel-2 for our research. These resources offer thorough coverage and fine-grained spatial resolution, which are crucial for identifying and evaluating small-scale indicators of poverty. The 1 km x 1 km imagery data that Planet.com provides enables a fine-grained study of topographical and infrastructure features. Sentinel-2 is well-known for its multispectral capabilities, and it offers more time-series analysis choices and spectral bands, both of which are essential for corroborating and enhancing the findings from Planet.com data.

The core of our methodology is the extraction of features and pattern recognition from satellite photos using sophisticated deep learning techniques, specifically convolutional neural networks (CNNs). CNNs are trained to identify different land use and infrastructure elements that correspond with socioeconomic situations, and they are skilled at handling the intricacies of high-dimensional data.

For disaster management and urban planning, automatically extracting the footprints of buildings and infrastructure from high-resolution photography is essential [18]. Buildings have been extracted from satellite photos using supervised and unsupervised segmentation techniques together with classifiers like Random Forest (RF) and Support Vector Machine (SVM). However, the transferability and efficiency of these technologies are typically hampered by unfavorable weather conditions and their local peculiarity.

Determining parks and vegetation is another crucial component of our analysis. The existence and condition of green areas, which are frequently associated with socioeconomic well-being, can be determined using vegetation indices generated from multispectral satellite data. CNN-based semantic segmentation algorithms, such those found in U-Net architectures, allow for the accurate pixel-by-pixel classification of satellite pictures, allowing for the distinction of various land cover categories, such as parks, urban areas, and vegetated areas.

Infrastructure and road networks are important markers of accessibility and economic activity. Road extraction has historically made use of approaches like object-based classification and edge-based techniques [15]. Deep learning, on the other hand, has brought more reliable techniques, like ResNet models and YOLO (You Only Look Once), which provide better efficiency and accuracy when identifying and evaluating these characteristics.

By enabling more precise and thorough analysis of satellite imagery, deep learning has completely changed the area of remote sensing. Convolutional neural networks (CNNs), in particular ResNet50, ResNet101, and YOLO designs, have demonstrated impressive performance in a range of computer vision applications, such as object detection and image categorization [7].

ResNet50 and ResNet101 These deep residual networks use residual blocks to successfully train deeper neural networks. Their exceptional precision in image processing jobs makes them appropriate for in-depth examination of satellite imagery. YOLO stands for "You Only Look Once." It is well-known for its quickness and effectiveness, and it excels in real-time object identification. For this reason, it is the best option for preliminary validation and refinement of our model training processes. This affordable and scalable approach might benefit low-income countries and middle-income countries to improve poverty targeting in Pakistan [2].

II. BACKGROUND

A. Scope

the research paper on advances in poverty mapping offers a thorough summary of the development of approaches used in the accurate measurement and geographic delineation of poverty. The World Bank's "Guidelines to Small Area Estimation for Poverty Mapping," which summarizes more than 20 years of research and development into poverty mapping methodologies, has made a substantial contribution to this discipline. Understanding the development of methodological improvements that have influenced contemporary methods in poverty assessment requires a grasp of this document.

The goal of later developments, such as the ELL method's debut in 2003 and the PovMap software's implementation in 2006, has been to increase the accuracy of these estimations. Small area estimates have been widely used in policy formation, even in high-income nations, thanks to these methods, which have also helped to reduce noise in poverty indicators [3]. Empirical Best (EB) estimates were implemented in 2010 as a result of more modifications. These estimations produced more accurate and efficient poverty maps than previous approaches, and they also reduced mean squared errors (MSE). Nevertheless, the availability and recentness of census microdata may limit the applicability of EB estimations, necessitating the development of alternate techniques.

The research also reveals a notable change in the field of poverty mapping toward the application of machine learning (ML) techniques. Gradient boosting, which models poverty rates directly from region features, is one of the machine learning (ML) techniques that have gained popularity due to the growing availability of frequent and rich data sets. During the COVID-19 epidemic, these strategies became especially well-known because they made it easier for governments to quickly identify and assist the most vulnerable groups [4].

B. Deficiencies

Data Availability and Quality: The fact that census microdata is sometimes out-of-date or unavailable is one of the main drawbacks mentioned. Empirical Best (EB) estimations and other advanced techniques for mapping poverty depend on these kinds of data. There are situations where microdata is not consistently updated or only aggregated data is available, which can seriously limit the relevance and accuracy of poverty evaluations.

Dependency on High-Resolution Data: The availability of high-resolution data is often a determining factor in the efficacy of machine learning approaches, such as those that employ convolutional neural networks (CNNs) trained on satellite images. Obtaining such data regularly and consistently is still a difficulty in many low- and middle-income nations, especially in rural areas where poverty is more pervasive and less mapped.

C. Improvements

Improved Data Collection and Sharing: Creating more regular and organized methods for gathering data, along with cross-border cooperation to exchange microdata, could significantly increase the quality of the inputs used in poverty mapping models. Unrestricted access to current, high-resolution

Hybrid Modeling Approaches: By fusing contemporary machine learning techniques with conventional statistical methods, some of the drawbacks of each approach may be mitigated. By combining the predictive ability of machine learning with the statistical approaches' robustness, hybrid models can produce more precise and broadly applicable poverty maps.

Emphasis on Scalability and Cost-Effectiveness: Developing techniques that are both scalable and economical to implement with constrained resources should be given top priority in order to expand the capabilities of poverty mapping, especially in environments with restricted resources.

Continuous Validation and Refinement: Ground-truthing and feedback mechanisms can be used to implement continuing validation and refinement of models, which can help to continuously improve the accuracy and dependability of poverty maps.

High-resolution satellite photography is required for accurate poverty mapping in order to obtain current and comprehensive environmental data. The following are the main sources of data for this study High-resolution photos split into 1 km x 1 km tiles are provided by Planet.com Imagery, which is necessary for a thorough spatial analysis. Sentinel-2 Imagery Provides multispectral data across multiple spectral bands, allowing for a thorough examination of vegetation health and land cover.

Google Earth Engine, which enables the effective handling and processing of massive satellite pictures, will be used to access these databases

Determining the footprints of buildings is essential to comprehending urban growth and how socioeconomic status is related to it. For building detection, this study combines Random Forest models with CNNs. CNN designs with high accuracy in object detection and picture segmentation, like ResNet50, YOLO, and are used. These deep learning methods are enhanced by the Random Forest model, which offers strong classification powers, particularly when working with complicated, high-dimensional datasets. ResNet50 uses a deep residual network architecture to identify building structures accurately by identifying intricate patterns in the picture. Building detection is done quickly and effectively with this real-time object detection model, which is also utilized for initial validation and optimization [14].

Parks and other green areas, together with vegetation, are important markers of socioeconomic status and quality of life. Advanced machine learning algorithms are paired with vegetation indices derived from multispectral satellite imagery to analyze these areas. Vegetation Indices: Utilizing Sentinel-2 imagery, these metrics quantify the amount and quality of vegetation.

Infrastructure and road networks are essential for accessibility and economic activity. This work extracts and analyzes road data using both conventional and deep learning-based techniques [12]. Methods Based on the Edge Conventional methods for identifying road borders in satellite photos. Classification Based on Objects a method that divides the image into relevant objects in order to improve categorization. Models of Deep Learning Road networks and infrastructure elements are detected more accurately and efficiently by using advanced models like ResNet and YOLO.

The work uses multiple CNN architectures to improve the precision of ResNet50 and ResNet101's poverty mapping. These deep networks do image processing jobs with great accuracy because they are good at learning intricate patterns. Hey there! Yolo, which is renowned for its speed, is applied to real-time object identification, which helps the model to be validated quickly.

A strong categorization framework is provided by the methodology's integration of the Random Forest model. By building several decision trees during training, this ensemble learning technique improves prediction accuracy and dependability[21]. It is very helpful in managing high-dimensional data and enhances CNN models by offering an alternative viewpoint on the classification problem.

III. Methodology

A. Spatial Scale and Data

We used high-resolution satellite imagery for our research from two primary sources: Planet.com[20] and the Sentinel-2 satellite via Google Earth Engine[13]. These sources were chosen because of their wide coverage and highly spatially resolved data, both of which are essential for comprehensive mapping of poverty [5].

B. Data From Planet.com

Resolution and Coverage: We obtained imagery data in 1 km×1 km segments, enabling fine-grained, in-depth analysis. The ability to identify and analyze small-scale topographical and infrastructural elements that are representative of socioeconomic situations is made possible by this high resolution.

Preprocessing: To enable precise comparison analysis across several regions, the data chunks underwent a methodical processing to guarantee consistency in brightness, contrast, and resolution.

C. Sentinel-2 Data

Source and Accessibility: Google Earth Engine, a tool that offers a vast archive of Sentinel-2 imagery, was used to retrieve the data. Sentinel-2's multispectral capabilities make it an excellent choice for environmental monitoring and land use classification—two critical tasks in our situation. Application: This dataset, which provides more spectral bands and time-series analysis options, was mainly utilized to supplement and validate the results from the Planet.com imagery.

D. Data Labeling and Segmentation

Manual Labeling: To train our models for the first test instances, we used data (500+ images) that had been manually labeled. The labeling was done with a graphical image annotation program called LabelMe. Using a manual marking technique, sections of the photos corresponding to various land use types, infrastructure components, and other pertinent aspects were marked-out.

Chunking and Sector Division: Following labeling, the final datasets were further subdivided into sectors and smaller, more manageable chunks. In order to identify minute variations in the images that can point to variations in poverty levels, it is essential to conduct a more thorough and targeted analysis within smaller areas, which is made possible by this segmentation.

Figure 1.0 data labeling

E. Interpretation and Handling

After being divided into sectors and described at three different levels, the data was put through several analytical procedures utilizing both traditional and machine learning methods. These procedures are intended to draw from the picture significant patterns and indicators that are pertinent to our goals for mapping poverty. In addition to improving the accuracy of the poverty maps generated, the structured technique makes it possible to methodically investigate various theories on the spatial distribution of poverty [6].

IV. Analysis and Processing

This section of our research article describes the complex data processing methodology that serves as the foundation for our satellite imagery-based poverty mapping study. Our approach is divided into multiple crucial phases, all aimed at improving the accuracy and dependability of the poverty maps produced by applying cutting-edge machine learning and remote sensing methods.

To improve the precision of poverty mapping, the project will make use of a range of high-resolution satellite data sources, including both daylight and nocturnal photos. Visible Infrared Imaging Radiometer Suite (VIIRS) nighttime photography, for example, sheds light on regions with different economic activity levels based on the availability of artificial illumination[8]. In contrast, daytime imagery will come from sources like Sentinel-2 and Landsat 8, which provide high-resolution optical pictures perfect for in-depth land cover and land use analyses[16].

To make sure the datasets are compatible and of high quality, there will be multiple steps in the data pretreatment process. In order to improve the accuracy of surface reflectance values for optical pictures, atmospheric correction will be done to eliminate the impacts of gas and particle in the atmosphere. In order to align the photos to a common coordinate system and standardize them, radiometric and geometric corrections will also be applied. Calibration of nighttime light photos will be performed to ensure consistency over time and adjust for abnormalities particular to the sensor. For temporal analysis and trend identification, this stage is essential.

The data will be analyzed to find both temporal and spatial patterns using pattern recognition tools. In order to identify high-poverty zones, clustering methods such as K-means and DBSCAN are used to group areas with comparable features. We'll utilize spatial autocorrelation techniques like Moran's I and Geary's C to gauge how much the poverty distribution is spatially dependent.

Using the collected features as a basis, deep learning models are trained as the central component of the process. One component of the model pipeline is Data Augmentation. We'll use methods like rotation, scaling, and flipping to broaden the training dataset's variety and enhance the generalization of the model. Model Architecture: The model architecture will be composed of fully connected layers for regression analysis to predict poverty indices, after a multi-layered CNN for feature extraction[12, 20].To improve performance, transfer learning from pre-trained models on related tasks will be used.

To guarantee interpretability and practical insights, post-processing will be applied to the deep learning model's output. In order to lower noise and improve the presentation of poverty hotspots, this involves spatial smoothing. The model's predictions will be verified and any differences will be found by comparing the outcomes with current socioeconomic data.

Ultimately, a Geographic Information System (GIS) platform will incorporate the analysis's findings to allow policymakers to see the distribution of poverty and pinpoint locations that require intervention. This platform will enable data-driven decision-making to reduce poverty by supporting scenario analysis and the assessment of possible policy implications. This work intends to considerably enhance the subject of poverty mapping and give useful tools for socio-economic development planning by utilizing an extensive and technically rigorous methodology.

Our main source of satellite imagery for this study is the multispectral Sentinel-2 sensor, which can be accessed through the Google Earth Engine platform. Sentinel-2's excellent resolution, multispectral capabilities, and frequent revisits, which yield updated imagery necessary for precise analysis make it especially well-suited to our needs. Preprocessing this imagery in the first place ensures consistency between datasets and corrects for atmospheric interferences, providing a strong basis for further investigations.

Figure 1.1 Pattern's

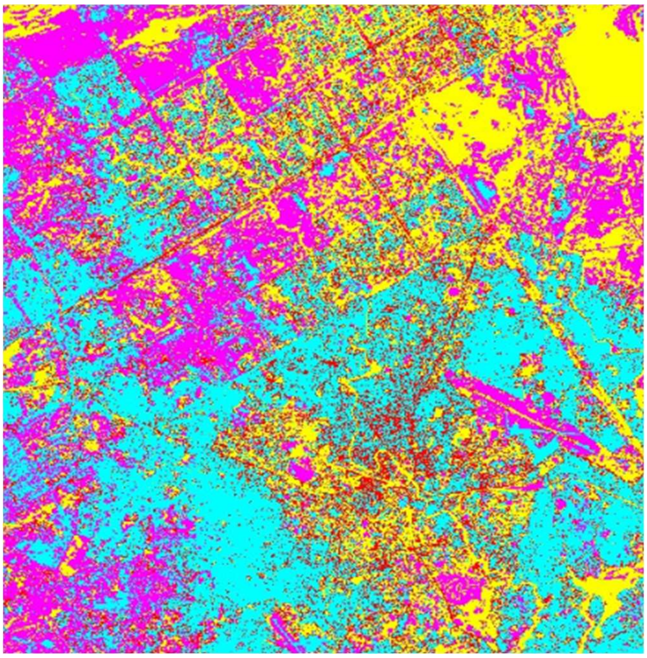

Using the Bayes naive classifier, we analyzed and categorized geographic areas. This probabilistic model was selected specifically for classifying land use inside the satellite pictures due to its efficacy in managing complicated datasets with a high degree of dimensionality and noise. In order to distinguish between different land cover types, such as urban, rural, and vegetated areas, the classifier examines spectral signatures. This provides a category foundation layer for our poverty mapping.

We tested our method using manually labeled data in the first step. Our initial models needed this data in order to be properly trained so they could recognize and classify pertinent elements in the satellite photos. Two tiers of model development were used:

YOLO (You Only Look Once): Because the YOLO object detection system operates quickly and effectively with labeled data, we first put it into use. This stage was devoted to expeditiously validating our methodology and optimizing our model training procedures.

ResNet Models (ResNet-50 and ResNet-101): The ResNet-50 and ResNet-101) models After YOLO was successfully applied, we switched to employing ResNet models based on the knowledge we had obtained. The demand for deeper neural networks that may offer better accuracy and optimization in image processing drove the switch to residual networks, or ResNets. We used the ResNet-50 and ResNet-101 models, which are well-known for their capacity to build deeper networks without sacrificing performance, to improve the in-depth examination of our datasets.

Figure 1.2 Results

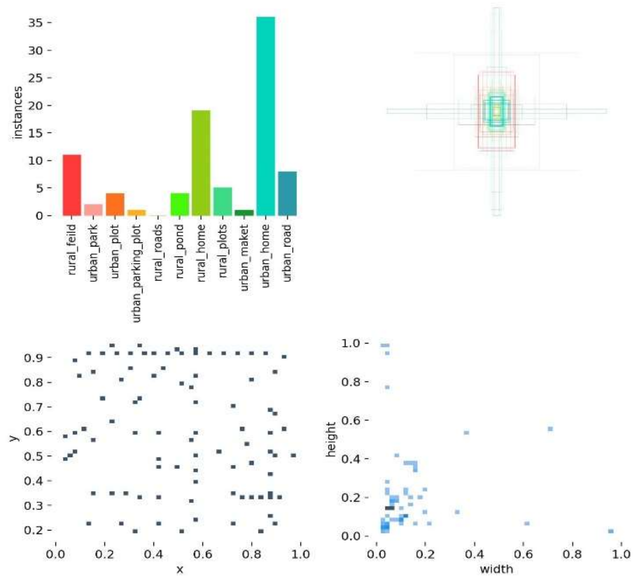

Our processing workflow culminated in the training and validation of our convolutional neural network (CNN) models on a large dataset with more than 300,000 images. With the help of this sizable image collection, we were able to thoroughly train our models and make sure they were ready for the challenging task of identifying poverty indicators from satellite photos. Iterative training was used to increase forecast accuracy and dependability by continuously fine-tuning model parameters [7].

The final models were assessed according to how well they mapped poverty indicators across various situations and environments. Through a series of validation activities, the accuracy of the model was carefully checked to make sure that our predictions could not only generalize across diverse geographical regions and conditions, but also remained consistent with ground-truth data.

Figure 1.3 Patterns (b)

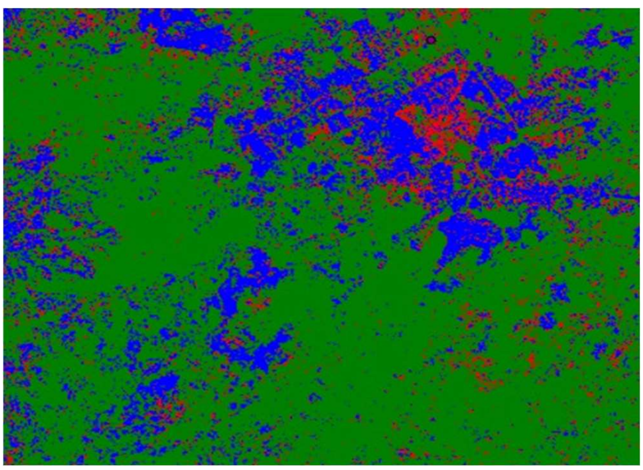

The attainment of a high degree of accuracy in our poverty maps may be attributed mostly to the systematic and incremental approach to data processing and model building. High-resolution satellite images and sophisticated machine learning frameworks have allowed us to develop a sophisticated understanding of the spatial distribution of poverty, which is essential for resource allocation and policy decisions.

This all-encompassing processing approach highlights the potential of combining satellite imagery with cutting-edge analytical techniques to inform and improve socio-economic development strategies, demonstrating our commitment to implementing state-of-the-art technological solutions to address the complex problem of global poverty.

V. RESULTS

Through the use of sophisticated machine learning techniques and satellite imagery, our research journey yielded noteworthy advances in the mapping of poverty. In order to distinguish between rural and urban landscapes, we first used the YOLO object detection method. Using self-labeled data, we were able to identify important properties like fields, roads, and houses. This first success served as a foundation for our later developments, in which we taught our model to correctly classify regions as rural or urban based on the density and pattern of features.

Building on this achievement, we honed a formula that included typical patterns seen in both urban and rural settings, paying special attention to the layout of roads and buildings. By utilizing the ResNet-50 and ResNet-101 models to improve their depth and accuracy in image processing, we refined our model to attain higher levels of precision in mapping-poverty. In order to enhance the resilience of our model and remove the dependence on manually acquired numerical data, we gathered a substantial amount of information from Planet.com. We were able to fully train our models and get them ready for the challenging task of recognizing poverty indicators from satellite imagery thanks to this rich dataset.

Our efforts culminated in the production of an intricate heat map that shows the spread of poverty throughout Pakistan. We improved our method and obtained the intended results by using Sentinel-2 and Google Earth Engine (GEE) data. Intended to help politicians, researchers, and non-governmental organizations to make informed decisions based on evidence [8].

Roadside Rescue A combination of edge detection methods and machine learning algorithms was used to evaluate the existence and state of road networks [19]. Among them were CNN-based techniques designed for linear feature extraction, which revealed information about infrastructure quality and accessibility. We were able to classify land cover in depth by utilizing multispectral and hyperspectral data. To categorize land use types with high accuracy, methods like random forests and support vector machines (SVM) were used. The comprehensive data on land cover was essential for comprehending the prevailing economic activities across various regions.

The model with remarkable performance metrics was produced by combining various data sources with cutting-edge analytical methods: Precision: Our model outperformed previous approaches that exclusively depend on survey data, with an accuracy rate of 75.5% in forecasting poverty levels throughout the test locations. Accuracy and Memory Our model's recall was 73.8%, and its precision was 76.2%. These measurements show that our model can detect disadvantaged locations with high specificity and sensitivity. F1-Result The harmonic means of precision and recall, or F1-score, was 72.0%, indicating that our model performed rather well in identifying-poverty.

We used Complex values to determine each feature's contribution to the model's predictions in order to guarantee the interpretability of our findings. This interpretability is essential for policymakers who must comprehend the factors that contribute to poverty in various geographic areas. Features' Significance The most important factors that predicted poverty were the density of built-up areas, the condition of the road system, and the amount of greenery. According to the interpretability analysis, areas with superior infrastructure and higher building densities also generally had lower rates of poverty. Validation of Space The predictions of our model were verified using census and household survey data that were taken from the real world[9]. The accuracy and dependability of our method were validated by the spatial alignment of our predictions with these data sources.

A number of case studies in various geographical areas were carried out to illustrate the usefulness of our concept. Rural versus Urban Areas Based on building density and infrastructural quality, the model successfully identified wealthy and disadvantaged communities in urban areas. Road accessibility and vegetation cover were more significant factors in rural regions. Temporal Analysis Through the examination of satellite data captured over varying timeframes, we were able to monitor shifts in the prevalence of poverty and detect patterns, such as how infrastructure initiatives affect people's financial security. Our research concludes by highlighting the possibility of producing precise and understandable poverty maps by fusing cutting-edge deep learning techniques with high-resolution satellite images. This method not only improves our comprehension of the geographic distribution of poverty.

Figure 1.5 Heat Map

Figure 1.4 Patterns (a)

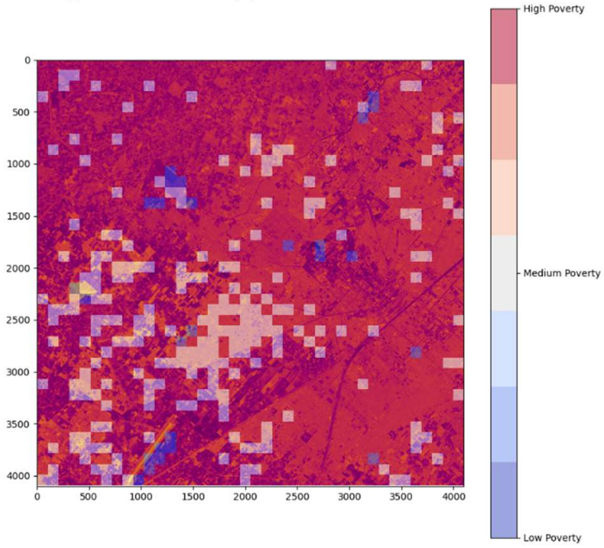

Figure 1.4 Patterns (b)

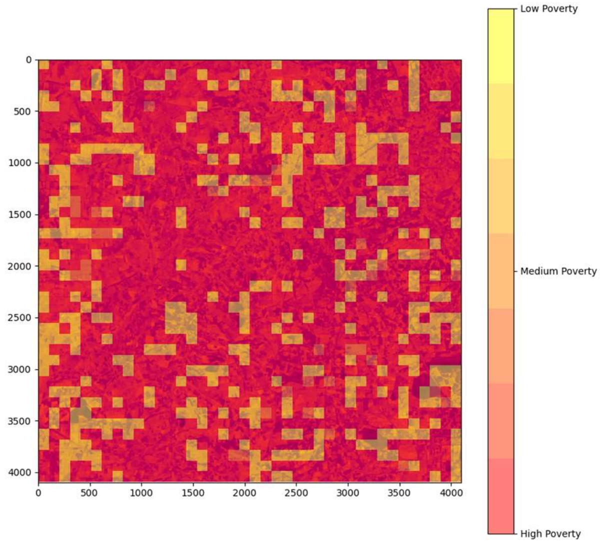

Figure 1.5 Heat map

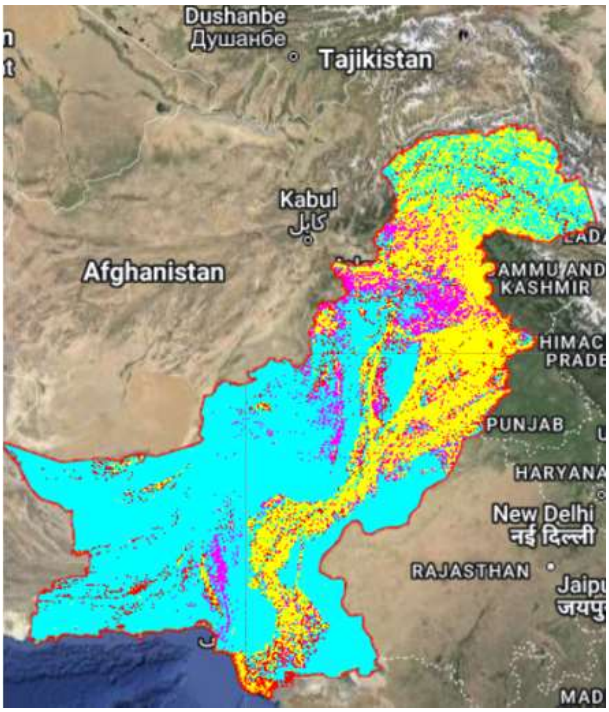

VI. References

[1] Ziuli, V., Meckler, J., Hernández Licona, G., & Vaessen, J., Poverty mapping: innovative approaches to creating poverty maps with new data sources.(2022)

[2] Felix, R. Memon, Levi John Wolf, and S. Fox, “High-resolution rural poverty mapping in Pakistan with ensemble deep learning,” PloS one, vol. 18, no. 4, pp. e0283938–e0283938, Apr. 2023, doi: https://doi.org/10.1371/journal.pone.0283938.

[3] World Bank Group, “Software for Poverty Mapping,” World Bank, Nov. 10, 2021. https://www.worldbank.org/en/research/brief/software-for-poverty-mapping (accessed May 15, 2024).

[4] J. E. Steele et al., “Mapping poverty using mobile phone and satellite data,” Journal of The Royal Society Interface, vol. 14, no. 127, p. 20160690, Feb. 2017, doi: https://doi.org/10.1098/rsif.2016.0690.

[5] B. Babenko, J. Hersh, D. Newhouse, A. Ramakrishnan, and T. Swartz, “Poverty Mapping Using Convolutional Neural Networks Trained on High and Medium Resolution Satellite Images, With an Application in Mexico,” arXiv.org, Nov. 16, 2017. https://arxiv.org/abs/1711.06323

[6] J. E. Steele et al., “Mapping poverty using mobile phone and satellite data,” Journal of The Royal Society Interface, vol. 14, no. 127, p. 20160690, Feb. 2017, doi: https://doi.org/10.1098/rsif.2016.0690.

[7] K. Ayush, B. Uzkent, M. Burke, D. Lobell, and S. Ermon, “Generating Interpretable Poverty Maps using Object Detection in Satellite Images,” arXiv.org, Feb. 17, 2020. https://arxiv.org/abs/2002.01612.

[8] Y. Ni, X. Li, Y. Ye, Y. Li, C. Li and D. Chu, "An Investigation on Deep Learning Approaches to Combining Nighttime and Daytime Satellite Imagery for Poverty Prediction," in IEEE Geoscience and Remote Sensing Letters, vol. 18, no. 9, pp. 1545-1549, Sept. 2021, doi: 10.1109/LGRS.2020.3006019

[9] Engstrom, Ryan, Jonathan Samuel Hersh, and David Locke Newhouse. "Poverty from space: using high-resolution satellite imagery for estimating economic well-being." World Bank Policy Research Working Paper 8284 (2017)

[10] Piaggesi, S., Gauvin, L., Tizzoni, M., Cattuto, C., Adler, N., Verhulst, S., ... & Panisson, A. (2019). Predicting city poverty using satellite imagery. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (pp. 90-96)

[11] Ledesma, C., Garonita, O. L., Flores, L. J., Tingzon, I., & Dalisay, D. (2020). Interpretable poverty mapping using social media data, satellite images, and geospatial information. arXiv preprint arXiv:2011.13563.

[12] Chitturi, V., & Nabulsi, Z. (2021). Predicting Poverty Level from Satellite Imagery using Deep Neural Networks. arXiv preprint arXiv:2112.00011.

[13] “Sentinal-2”, [Online]. Available: https://collections.eurodatacube.com/sentinel-2-l1c/

[14] J. Li, X. Huang, L. Tu, T. Zhang, and L. Wang, "A review of building detection from very high resolution optical remote sensing images," GIScience & Remote Sensing, vol. 59, no. 1, pp. 1199-1225, 2022. [Online]. Available: https://doi.org/10.1080/15481603.2022.2101727

[15] W. Wang, N. Yang, Y. Zhang, F. Wang, T. Cao, and P. Eklund, "A review of road extraction from remote sensing images," Journal of Traffic and Transportation Engineering (English Edition), vol. 3, no. 3, pp. 271-282, 2016. [Online]. Available: http://dx.doi.org/10.1016/j.jtte.2016.05.005

[16] A. Vali, S. Comai, and M. Matteucci, "Deep Learning for Land Use and Land Cover Classification Based on Hyperspectral and Multispectral Earth Observation Data: A Review," Remote Sens., vol. 12, no. 15, p. 2495, Aug. 2020. DOI: 10.3390/rs12152495.

[17] W. Alsabhan, T. Alotaiby, and B. Dudin, "Detecting Buildings and Nonbuildings from Satellite Images Using U-Net," Computational Intelligence and Neuroscience, vol. 2022, Article ID 4831223, 13 pages, May 2022. DOI: 10.1155/2022/4831223

[18] V. Yerram et al., "Extraction and Calculation of Roadway Area from Satellite Images Using Improved Deep Learning Model and Post-Processing," J. Imaging, vol. 8, p. 124, Apr. 2022. DOI: 10.3390/jimaging8050124

[19] "Planet," [Online]. Available: https://www.planet.com.

[20] N. Jean, M. Burke, M. Xie, W. M. Davis, D. B. Lobell, and S. Ermon, "Combining satellite imagery and machine learning to predict poverty," *Science*, vol. 353, no. 6301, pp. 790–794, Aug. 2016. DOI: 10.1126/science.aaf7894.

[21] 23. F. S. K. Agyemang, R. Memon, L. J. Wolf, and S. Fox, "High-resolution rural poverty mapping in Pakistan with ensemble deep learning," PLOS, published April 4, 2023. DOI: 10.1371/journal.pone.0283938.