International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 07 Issue: 08 | August - 2023 SJIF Rating: 8.176 ISSN: 2582-3930

Exploring Techniques for Poverty Level Prediction from Satellite Imagery

Vidya Marathe
Department of Computer Engineering
COEP Technological University
Pune, India
vidyanm21.comp@coep.ac.in

Dr. Vijay M. Khadse
Department of Computer Engineering
COEP Technological University
Pune, India
vmk.comp@coep.ac.in

Abstract: *Determining the poverty levels of various regions throughout the world is crucial in identifying interventions for poverty reduction initiatives and directing resources fairly. However, reliable data on global economic livelihoods is hard to come by, especially for areas in the developing world, hampering efforts to both deploy services and monitor/evaluate progress. Satellite images can be used to detect economic activity and infrastructure of a region in order to predict the wealth index (or poverty level) of a location. Research on poverty prediction from satellite images using deep learning has shown promising results. By leveraging the rich spatial information captured by satellite imagery, deep learning models can analyze various features and patterns to predict poverty levels. The aim of this exploratory survey is to comprehend and evaluate the work done in the field of poverty prediction and to understand how data captured using satellite can be mapped to various wealth and socio-economic wellbeing parameters. In this survey we examine different satellite image datasets that can be utilized for predicting poverty levels, while also reviewing previous research conducted in this domain. Ultimately, we propose an architecture that shows promise in delivering favourable outcomes for this task.*

Index Terms—Image processing, Satellite Imagery, Deep Learning

I. INTRODUCTION

The United Nations (UN) launched a list of 17 Sustainable Development Goals (SDG) in 2015, which provides a blueprint to achieve a better and more sustainable future for all. The first goal in the list is to "End poverty in all its forms everywhere" [1]. However, despite this being the foremost goal of the UN, there are extensive regions across the world where extreme levels of poverty are endemic.

The present difficulty for statisticians, officials, and organisations working on poverty statistics is that the data must be broken down by geography, ethnicity, gender, income class, and other pertinent factors. While household income and expenditure surveys have sample sizes large enough to estimate some levels of poverty, the datasets are not sufficiently large or detailed to produce estimates that would be accurate enough to meet SDG 1's requirements and enable development planners to concentrate on areas that need immediate aid and poverty interventions.

Along with the granularity level concerns, conducting family income and spending surveys is time-consuming, expensive, and labour-intensive, especially during the preparation, execution, and publication of results. Language hurdles, household educational status, dealing with people's expectations and lack of cooperation, weather conditions, and accessibility to remote villages/locations are challenges that research agencies and workers performing the survey also encounter. In addition to these difficulties, conducting these surveys requires an enormous financial investment, and the lack of financial assistance and support from governments and international organisations like the UN and World Bank affects the frequency and timeliness of the publication of updated poverty data.

There was a significant rise in the worldwide poverty rate due to COVID-19 pandemic, jumping from 8.3 percent in 2019 to 9.2 percent in 2020. This setback erased approximately three years' worth of progress in poverty reduction. Further, due to the emergence and spread of COVID-19 pandemic which led to the adoption of public health measures and social distancing protocols, data collection under several restrictions and health concerns for both enumerators and respondents was a challenge, especially in underdeveloped regions. E.g., The census of at least 18 countries was postponed from 2020 to 2021 while four countries postponed it to 2022 or beyond.

Hence, an alternate cost-effective method needs to be explored for meeting the growing demand and making the process of data collection, compilation and publishing of results easy and hassle-free becomes more important. Deep learning-based computer vision techniques may accurately and affordably identify a region's poverty status using satellite photos, bringing poverty detection closer to real-time and at a lower cost.

Planning and implementing policies and programmes to reduce poverty consume a considerable amount of resources in developing nations. If and when these schemes are employed, ground level surveys of socio-economic indicators, like the decennial census, will serve as the major sources of data. However, thorough data collection requires a lot of manual labour and financial resources, which leads to occasional sampling. As a result, it is frequently difficult to develop policy when timely and accurate data are lacking. This could result in inefficient execution and, occasionally, even needless spending.

The ready availability of timely, inexpensive, and accurate data sources can help address some of these issues. One inexpensive data source that provides a wealth of information about local development is satellite imagery. The increasing resolution of satellite imagery and the relative ease of accessing data in the public domain make satellite imagery a potential resource.

This survey offers a systematic study that helps understand and evaluate existing systems that predict poverty levels from images captured using satellites and map them to various wealth and socio-economic wellbeing parameters like population density, electricity, natural resources and availability of facilities – healthcare, sanitation and education etc. The study provides:

- A comprehensive survey of poverty prediction.
- Datasets available for study of poverty level estimation and classification.
- Identify challenges in estimation of wealth and socioeconomic status of a region from satellite images using deep learning.

II. LITERATURE REVIEW

In their research N Jean et al. [2] demonstrate an accurate, inexpensive, and scalable method for estimating consumption expenditure and asset wealth from high-resolution satellite imagery. They have chosen 5 African countries—Nigeria, Tanzania, Uganda, Malawi, and Rwanda and trained a convolutional neural network to identify image features that can explain up to 75% of the variation in local-level economic outcomes. They highlight the scarcity of labelled data, yielding data sets many orders of magnitude smaller than those typically used in deep learning applications, and how they overcame the challenge using a multistep “transfer learning” [3] approach. Despite being trained partially on nightlights, the model outperforms nightlights in 81.3% of trials, with an average increase in R^2 of 0.04. For clusters below two times the poverty line, model outperforms nightlights in 98.5% of trials, with an average increase in R^2 of 0.10, an 81.2% increase in explanatory power. For clusters below three times the poverty line, the model outperforms nightlights in 99.5% of trials, with an average increase in R^2 of 0.12, corresponding to a 54.2% increase in explanatory power.

Ayush K et al. [4] have worked on Fine-grained Object Detection on High-Resolution Satellite Imagery. Their Policy Network uses cheaply available Sentinel-2 low-resolution image (Sentinel-2 imagery is freely available) representing a cluster to output a set of 6 actions representing unique 1000 X 1000 px high-resolution tiles in the 34 X 34 grid. Then object detection is performed on the sampled HR tiles (black regions represent dropped tiles) to obtain the corresponding class-wise object counts (L-dimensional vectors). Finally, the class wise object counts vectors corresponding to the acquired. HR tiles are added element-wise to get the final feature vector representing the cluster. The reinforcement learning approach (Adaptive Tile Selection) dynamically identifies where to acquire high-resolution images, conditioned on cheap, low resolution data, before performing object detection. They have also considered Seasonality as a factor, which has high relevance in rural settings where crops are grown and agriculture is a major source of income. They acquired two sets of low-resolution imagery, one from dry-season (Dec - Feb) in Uganda and other from wet season (March-May, Sept-Nov) corresponding to the survey year. The policy network is trained on 1249 images and uses 42.3% HR images while approximating the fixed approach in mean Average Precision (mAP) and mean Average Recall (mAR) metrics (Redmon and Farhadi 2018). These results indicate that the policy network learns to successfully choose regions where there are objects of interest and eliminate the regions with no objects of interest. The model (wet season) achieves 0.61 R^2 and substantially outperforms the published state-of-the-art results (Ayush et al. 2020) (0.53 R^2) while using around 80% fewer HR images.

Huang Zixi [5] in their research merged data from two sources, Oxford Poverty and Human Development Initiative and Poverty Possibility Index. They considered factors like age, education level, financial activities, availability of phone and internet etc. for the prediction of poverty, concluding that Decision tree (0.8210) and gradient boosting (0.7853) provide high accuracy and interpretability.

Alsharkawi A et al. [6] take into account all the household expenditure and income surveys in their research which uses data collected by the DoS in five different national household expenditure and income surveys of Jordan, the compiled dataset contains 63,211 household responses with 47 features (17 categorical features and 30 numerical features). They apply multiple ML models. The LightGBM outperforms other algorithms and has achieved the best performance with 81% F1-Score.

S Pandey et al. [7] discuss a two-step approach for poverty prediction. They developed a multi-tasking convolutional NN to predict roof type, availability of electricity, roads, settlements, farmlands and water bodies from the satellite images. These parameters correlate with visual features and their values are available level in the Indian census of 2011, thus providing a larger dataset for training. Multi-tasking architecture is flexible, it involves learning multiple tasks simultaneously while exploiting the similarities and differences among the tasks, it can learn shared and independent representations for the different tasks and reduces the total computation time of training and prediction. Coordinates of the centre of a village are obtained from Google Geocoding API, and then Google Static Maps API is used to extract images from the determined geocodes. The 1920×1920 sized images, at zoom level 16, fully cover 67.46% (66135) villages. Each image spans a ground surface area of approximately 19 km^2 Area is classified as "poor" if the fraction of households in a sub-district having income less than Rs.5000 is greater than the threshold value. Accuracy of up to 96.9% is achieved for threshold value of 0.6

V Chitturi et al. [8] discuss that limited availability of data is often a barrier in poverty level prediction using satellite imagery. They experiment with both daytime and nighttime imagery and have developed a multitude of deep neural networks under various circumstances with the goal of predicting wealth index of a region (representing the poverty level). Flattened input image is passed through two multi-layer perceptrons, each with 512 nodes before finally outputting a wealth index score. The same architecture is used for daytime and nighttime images. An SGD optimizer with momentum is used and DNNs are trained using mean squared error (MSE) loss. The goal is to minimize RMSE values. All hyperparameters and training procedures were kept constant, with the exception of learning rate which differed for nighttime and daytime images. A DNN trained on nighttime images achieves RMSE's of 1.222 (95% CI: 1.033-1.423), 1.132 (95% CI: 0.963-1.354), and 1.106 (95% CI: 0.956-1.283) when trained on 800, 1600, and 2400 images respectively. A DNN trained on daytime images achieves RMSE's of 1.611 (95% CI: 1.399-1.832), 1.459 (95% CI: 1.275-1.633), and 1.477 (95% CI: 1.203-1.894) when trained on 800, 1600, and 2400 images respectively.

According to Babenko et al. [9] it is difficult and expensive to map the geographic distribution of poverty in emerging nations. They use high and medium resolution satellite photos from Planet and Digital Globe to train convolutional neural networks (CNNs) on the whole 2 million km^2 area of Mexico, with spatial resolutions of 3-5 m^2 and 50 cm^2, respectively. The 896 municipalities in the 2014 MCS-ENIGH are used to train CNNs. Following were the findings of their study: 1) the best models, which use satellite estimated land use as a predictor, explain roughly 57% of the variation in poverty; 2) across all MCS-ENIGH municipalities, explanatory power decreases to 44% in a CNN prediction and landcover model; and 3) predicted poverty from CNN predictions alone explains 47% of the variation in poverty in the validation sample and 37% across all MCS-E. Their best model explains5 7% up of the poverty variation in poverty in a sample set of 10% percent of municipalities in Mexico

III. AVAILABLE DATASETS

Publicly available datasets that provide satellite images for poverty / socio-economic studies:

- Google Earth Engine [10]: Google Earth Engine is a cloud-based platform that offers a wide range of satellite imagery and geospatial datasets. It provides access to datasets such as Landsat, Sentinel, and MODIS, which can be used for poverty index studies. Accessible through Google Earth Engine website and can be integrated through APIs.
- Nighttime Lights [13]: The Visible Infrared Imaging Radiometer Suite (VIIRS) collects data on nighttime lights, which can be used as a proxy for economic activity and poverty levels. The data is available from the National Oceanic and Atmospheric Administration

(NOAA) and can be accessed through the Earth Observation Group's website.

- Global Human Settlement Layer (GHSL) [11]: GHSL provides high-resolution satellite imagery that includes information on population density, built-up areas, and other socio-economic indicators. The data is available from the European Commission's Joint Research Centre (JRC).
- LandScan Global Population Database [14]: This dataset provides estimates of population distribution and density at a high resolution. It can be used to analyze population patterns in relation to poverty levels. The data is available from the Oak Ridge National Laboratory.
- WorldPop [12]: WorldPop offers high-resolution gridded population datasets that provide estimates of population distribution. The data is available for different years and can be useful for poverty index studies. WorldPop datasets can be accessed through their website.
- High-Resolution Settlement Layer (HRSL): HRSL provides detailed information on human settlements, including built-up areas and population density. The data is available for various countries and can be accessed through the Open Data Kit.

These datasets are used in combination with ground-level surveys and socioeconomic data as the ground truth data, for accurate analysis.

IV. RESEARCH CHALLENGES

The lack of labelled data and unavailability of large amounts of high-quality satellite image datasets to train a model to predict poverty level with sufficient accuracy. Obtaining / Creating labelled data is both time-consuming and tedious process. When labelled data is scarce, it can lead to overfitting or underfitting, which can ultimately result in poor performance of the model. High resolution satellite images can significantly increase the computational cost of developing and training machine learning models. High resolution images – Tag Image File format (TIF) images contain a large amount of data, which requires significant processing power and memory to handle. As a result, training models on these images can be both time-consuming and expensive. To address this issue, downsampling needs to be performed which reduces their size while still retaining key features. Another approach is to use specialized hardware, such as graphics processing units (GPUs) or tensor processing units (TPUs), which can handle large amounts of data more efficiently than traditional CPUs. Techniques such as transfer learning can also be used which involve using a pre-trained model as a starting point for training on a new task, this helps reduce the number of computational resources required. Availability of recent ground truth data for mapping is necessary because if there could be severe implications of the ground truth data being old, as the data may no longer reflect the current state of the system or environment being studied. The context or conditions may have changed significantly since the data was collected and it will no more accurately reflect the trends or patterns to be studied. Additionally, some of the datasets require registration or have usage restrictions.

V. PROPOSED SYSTEM

The ground truth data for location-based indicators such as wealth index can be acquired through a variety of surveys, including those conducted by organizations like DHS, the World Bank, government agencies, or NGOs. To complement these surveys, we can extract images based on the obtained coordinates. Additionally, data augmentation techniques like rotation, vertical flip and addition of gaussian noise etc, can be employed to enhance the dataset and make it robust. Our proposal involves designing CNN models that utilize both daytime and nighttime satellite images as inputs. By applying transfer learning techniques, we anticipate achieving superior results. Notably, nighttime imagery has demonstrated promising outcomes in previous studies. The regression models trained can predict the wealth index (poverty level) of the region as illustrated in Figure 1.

VI. CONCLUSION

Despite challenges such as the lack of a labelled dataset, restrictions on data access, the collection of reliable ground truth data, and difficulties in mapping ground truth data to satellite images, image processing on high-quality satellite images holds great promise for poverty prediction. Deep learning algorithms have demonstrated encouraging results in classification and prediction within this field.

To enhance the accuracy of poverty level prediction, several approaches can be adopted. By training neural networks to utilize both daytime and nighttime satellite

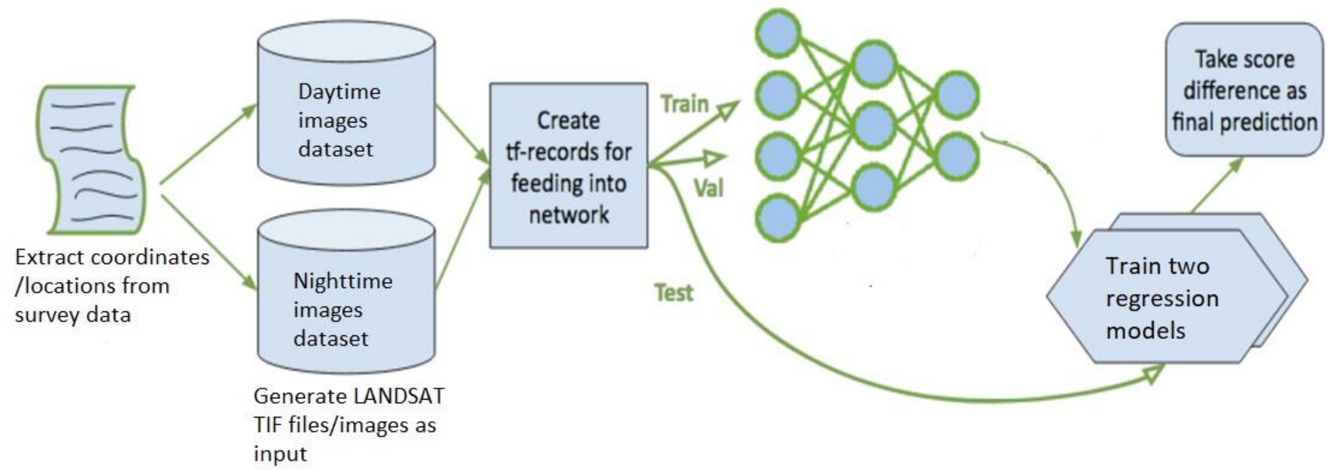

Figure 1:Proposed System Diagram

images as input, the benefits of both can be harnessed. Additionally, multiple satellite images of a single city are often accessible, allowing them to be provided as input to the network.

Furthermore, employing techniques like data augmentation, image enhancement, and image feature extraction, in combination with sophisticated deep learning models and transfer learning approaches, can yield promising results in poverty prediction. These techniques can help in increasing the diversity and quality of the training data and extracting meaningful features from the images, thereby enhancing the prediction accuracy.

REFERENCES

[1] UN.ESCAP (2021). SDG 1: No Poverty: end poverty in all its forms everywhere. Retrieved from: https://hdl.handle.net/20.500.12870/4538. https://www.un.org/sustainabledevelopment/poverty/.

[2] Jean N, Burke M, Xie M, Davis WM, Lobell DB, Ermon S. Combining satellite imagery and machine learning to predict poverty. Science. 2016 Aug 19;353(6301):790-4. doi: 10.1126/science.aaf7894. PMID: 27540167.

[3] S. J. Pan, Q. Yang, IEEE Trans. Knowl. Data Eng. 22, 1345–1359 (2010).

[4] Ayush, K., Uzkent, B., Tanmay, K., Burke, M., Lobell, D., & Ermon, S. (2021). Efficient Poverty Mapping from High Resolution Remote Sensing Images. Proceedings of the AAAI Conference on Artificial Intelligence, 35(1), 12-20. https://doi.org/10.1609/aaai.v35i1.16072.

[5] H. Zixi, Poverty Prediction Through Machine Learning, 2021 2nd International Conference on E-Commerce and Internet Technology (ECIT), Hangzhou, China, 2021, pp.314-324, doi: 10.1109/ECIT52743.2021.00073.

[6] Alsharkawi, A.; Al-Fetyani, M.; Dawas, M.; Saadeh, H.; Alyaman, M. Poverty Classification Using Machine Learning: The Case of Jordan. Sustainability 2021, 13, 1412. https://doi.org/10.3390/su13031412.

[7] Pandey, S., Agarwal, T., & C. Krishnan, N. (2018). Multi-Task Deep Learning for Predicting Poverty from Satellite Images. Proceedings of the AAAI Conference on Artificial Intelligence, 32(1). https://doi.org/10.1609/aaai.v32i1.11416.

[8] Predicting Poverty Level from Satellite Imagery using Deep Neural Networks https://doi.org/10.48550/arXiv.2112.00011.

[9] Babenko, B.; Hersh, J.; Newhouse, D.; Ramakrishnan, A.; Swartz, T. Poverty mapping using convolutional neural networks trained on high and medium resolution satellite images, with an application in

Volume: 07 Issue: 08 | August - 2023 SJIF Rating: 8.176 ISSN: 2582-3930

Mexico. arXiv 2017, arXiv:1711.06323.

[10] Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., & Moore, R. (2017). Google Earth Engine: Planetary-scale geospatial analysis for everyone. Remote Sensing of Environment, 202, 18-27. https://doi.org/10.1016/j.rse.2017.06.031.

[11] Pesaresi M., Politis P. (2022): GHS-LAND R2022A - Land fraction as derived from Sentinel2 image composite (2018) and OSM data.

[12] Open Spatial Demographic Data and Research. (n.d.). WorldPop. https://www.worldpop.org/

[13] Nighttime Lights — Earthdata. (n.d.). Earthdata. Retrieved June 27, 2023, from https://www.earthdata.nasa.gov/learn/backgrounders/nighttime-lights

[14] Dobson, Jerome Bright, Edward Coleman, Phillip Durfee, Richard Worley, Brian. (2000). LandScan: A Global Population Database for Estimating Populations at Risk. Photogrammetric Engineering and Remote Sensing. 66. 849-857.