University of the Witwatersrand, Johannesburg

Master of Science in Engineering by Research

Measuring Quality of Life in Gauteng using satellite images and machine learning

Emily-Rose Simões Steyn

School of Electrical Engineering and Information Engineering

A dissertation submitted to the Faculty of Engineering and the Built Environment, University of the Witwatersrand, Johannesburg, in fulfilment of the requirements for the degree of Master of Science in Engineering.

Johannesburg 2025

Abstract

Quality of life is a multidimensional concept that accounts for various socioeconomic factors, including wealth, safety, and access to education, healthcare, and basic services. Understanding and improving the quality of life for people around the world is a global endeavour supported by several of the United Nations Sustainable Development Goals. Governments and organisations rely on socioeconomic data to inform the quality of life of a particular country or region. Traditionally, socioeconomic data is collected via surveys and censuses. However, these are expensive to run and require significant work and coordination.

Prior research investigates using machine learning and big data sources like satellite images to predict various socioeconomic indicators. In recent years, studies on several sub-Saharan countries have shown the potential to measure poverty via indicators such as asset wealth and consumption expenditure. However, these studies often focus on the narrow scope of economic wellbeing. Limited work has been undertaken in estimating proxies of multidimensional quality of life from satellite images, especially in a South African context.

This research investigates the question: "how well can a machine learning model estimate a quality of life index from satellite images in Gauteng, South Africa?" By adapting the existing gold standard methods for predicting poverty from satellite images, this research investigates predicting a multidimensional quality of life index from daytime satellite images. The quality of life index for wards across Gauteng, South Africa is obtained from data provided by the Gauteng City-Region Observatory. This dataset is used to train a convolutional neural network to extract features from daytime satellite images (obtained from Planet Labs). From these features, the selected quality of life index is predicted.

Results obtained are suboptimal with a negative Coefficient of Determination, which is below baseline performance for poverty prediction and worse than predicting the sample mean of the test set. A critical analysis of the methods and results concludes that the methods applied are not suitable for a multidimensional index such as the one collected by the Gauteng City-Region Observatory. Primary conclusions suggest that this result is due to the subjective nature of factors in the index, the appropriateness of the machine learning model selected, and the physical disparities of South African urban environments within small spatial areas. A primary recommendation is further exploration of the boundaries of using machine learning methods with satellite data for measuring quality of life.

Declaration

I declare that this dissertation is my own unaided work. It is being submitted for the degree of Master of Science in Engineering to the University of the Witwatersrand, Johannesburg. It has not been submitted before for any degree or examination to any other university.

....

(Signature of Candidate)

...8th... day of ...Sept... year ...2025...

Acknowledgements

The journey to completing this dissertation has been unpredictable, challenging, exciting, and eye-opening. I am fortunate to have had so many people who have helped me along the way, believed in me, and cheered me on.

I am incredibly grateful to my supervisors, Professor Ken Nixon and Dr Martin Bekker, for their endless support, guidance, and patience as I navigated the ups and downs of my master's research. At every turn, you have encouraged me to not only push the boundaries of my own work but also explore additional opportunities (conferences, workshops, journal articles) that have given me a taste of academic life and truly enhanced my experience. I would like to also acknowledge Professor Ivan Hofsajer, who so freely shared his insights into academic writing and research at different phases of my studies providing me with perspective and motivation to simply write. Additionally, I would like to acknowledge the broader community in the School of EIE and at Wits, for their academic, administrative, and peer support.

This journey would have been impossible without my group of believers who have sustained me physically, emotionally, mentally, and spiritually over the last few years. I am especially grateful to my loving, hard-working and inspiring family - Jo, Pieter, Christina, Teresa, Richard, and Heather. And, of course, to Liam - my dear, funny, rock of a husband who has seen it all and relentlessly stayed on my team - **thank you**.

This work was supported by the University of the Witwatersrand Postgraduate Merit Award; Professor Barry Dwolatzky through funds gifted to the School of EIE; the SIGCHI Gary Marsden Travel Awards 2023; the IC2S2 Travel Grants 2024; and the Deep Learning Indaba 2024. This work was also supported by the project Focus on Africa Space Science and Technology for Future Development (FAST4Future). FAST4Future received funding from the EC under the grant agreement number 101082487 — ERASMUS-EDU-2022-CBHE.

Contents

List of Figures

List of Tables

Glossary

r^2 Coefficient of Determination. 31, 39, 47, 50, 51, *see*

Activation Maps Visualisations showing which parts of an input image most strongly influence a neural network's "decision-making" process, useful for understanding and interpreting model behaviour. xiii, 49

Adam Optimizer An adaptive learning rate optimisation algorithm widely used for training neural networks due to its efficiency and robustness. 38

Basemap Quad A standardised geographic tile that divides satellite imagery into uniform grid cells, typically 4096 × 4096 pixels each, used by satellite data providers like Planet Labs to organise and distribute imagery coverage. xii, 34–36, 41, 42, 44, 45

Coefficient of Determination A statistical measure that represents the proportion of variance in a dependent variable that is explained by an independent variable or variables. In this context, a metric used to evaluate the performance of a regression model. 31

Confirmatory Factor Analysis A statistical technique used to test whether a predetermined theoretical model of relationships between observed variables and underlying factors fits the data, unlike exploratory factor analysis which discovers these relationships empirically. xx, 11

Convolutional Neural Network A type of neural network specifically designed to process grid-like data such as images by using filters (kernels) to detect spatial features like edges, textures, and patterns through convolution operations. xx, 21

Coverage Bias Systematic exclusion of certain population groups from a survey's sampling frame, often occurring in online surveys due to unequal internet access (digital divide) across demographic groups. 2

Data Augmentation Techniques for artificially expanding training datasets by applying transformations (rotations, flips, crops, brightness changes) to existing samples, improving model generalisation and reducing overfitting. 39

Deep Residual Architecture A neural network design that uses skip connections to allow information to bypass certain layers, enabling the training of very deep networks without performance degradation. The ResNet architecture is an implementation of this design. 37–39

Dense Layer A fully connected neural network layer where each input neuron connects to every output neuron, typically used in the final layers of neural networks for classification or regression tasks. 31, 38, 54

Early Stopping A regularisation technique that halts model training when performance on validation data stops improving for a specified number of epochs, preventing overfitting and reducing computational costs. 38

Exploratory Factor Analysis A statistical method that identifies underlying relationships between observed variables without predetermined assumptions about their structure, used to discover latent constructs like quality of life dimensions. xx, 2, 3

Feature Extraction The process of automatically identifying and quantifying relevant patterns, textures, shapes, or other characteristics from raw data (such as satellite images) that can be used as inputs for machine learning models. 23, 37, 38, 40, 53–55, 61, 62

GeoTIFF A standard file format that combines geographic coordinate information with raster image data, allowing satellite imagery to retain spatial reference information for accurate mapping and analysis. 36

ImageNet A dataset containing 14 million labeled natural images across 1,000 categories. xvii, xix, 18, 21, 38

ImageNet Weights Pre-trained neural network parameters learned from the ImageNet dataset, widely used as initialisation for computer vision tasks. 37, 38, 58

k-fold Cross-Validation A model evaluation technique where data is divided into k subsets (folds), with each subset used once for testing while the remaining k-1 subsets train the model, providing more robust performance estimates. Especially useful when working with small datasets. 31, 37–39

Mean Absolute Error A regression evaluation metric that measures the average absolute difference between predicted and actual values, where lower values indicate better model performance. xx, 31

Mean Squared Error A regression evaluation metric that measures the average squared difference between predicted and actual values, commonly used as a loss function in neural network training. xxi, 31

MOSAIKS A method that extracts random convolutional features from satellite images that can be reused across multiple prediction tasks without retraining. The "kitchen sinks" reference comes from the machine learning concept of "throwing everything including the kitchen sink" at a problem by using many diverse, randomly generated features . xviii, xxi, 23

Multispectral Satellite Images Imagery captured across multiple wavelengths of electromagnetic radiation, including bands beyond the visible spectrum such as near-infrared and short-wave infrared, providing additional information about surface materials and vegetation. 22

Object Detection Computer vision techniques that identify and locate specific objects within images (such as buildings, roads, or vehicles), typically providing both classification labels and bounding box coordinates. 22, 25

Principal Component Analysis A statistical technique that reduces data dimensionality by identifying orthogonal components that capture the maximum variance in the original variables, commonly used to create asset wealth indices from household possession data. xxi, 9

Probabilistic Sampling Survey methodology where each member of the target population has a known, non-zero probability of being selected, enabling statistically valid inferences about the broader population. 2

Random Convolutional Features Image characteristics extracted using randomly initialised filters rather than learned patterns, as used in the MOSAIKS approach to create task-agnostic features for multiple prediction tasks. 23

Sigmoid Activation A mathematical function that constrains neural network outputs between 0 and 1, often used in the final layer of binary classification or regression models. 38

Skip Connections Direct pathways in neural networks that allow information to bypass certain layers, helping preserve important features during training and enabling the training of very deep networks. xvi, 37

Spatial Autocorrelation The tendency for geographically nearby locations to have similar characteristics. In satellite imagery analysis, this can bias machine learning models if training and test data contain spatially adjacent areas, as the model may learn location-specific patterns rather than generalisable features. 37

Spatial Join A geospatial operation that combines datasets based on their geographic location or spatial relationships (such as containment, intersection, or proximity) rather than common attribute values. 36

Statistical Performance Indicators and Index A World Bank framework that measures countries' statistical capacity across five dimensions: data use, data services, data products, data sources, and data infrastructure, providing a standardised assessment of national statistical systems. xxi, 15

Transfer Learning A machine learning technique where a model pre-trained on one large dataset (such as ImageNet) is adapted for a different, but related task, leveraging previously learned features to improve performance on the new task with limited data. An example of transfer learning is Transfer Learning with Nighttime Light Data. 21–23, 25, 38

Transfer Learning with Nighttime Light Data A two-step machine learning process where: (1) a model is first trained to predict nighttime light intensity from daytime satellite images, then (2) the learned features from this model are used to predict indicators like asset wealth or consumption expenditure. xix

Vanishing Gradient Problem A challenge in training deep neural networks where gradients become exponentially smaller as they propagate backward through layers, making it difficult to update the parameters in earlier layers and limiting learning effectiveness. 37

Visual Spectrum The portion of electromagnetic radiation visible to the human eye, typically captured in satellite imagery as Red, Green, and Blue (RGB) bands. 21, 22, 29, 40, 61

Weighted Least Squares A regression technique that assigns different weights to observations based on their reliability, precision, or importance, used when observations have unequal variances or when some data points are more trustworthy. xxi, 10

List of Abbreviations

MDB South African Municipal Demarcation Board. xii, 31, 33–35, 41, 42, 45

MICS Multiple Indicator Cluster Surveys. 13

MOSAIKS Multi-task Observation using Satellite Imagery and Kitchen Sinks. 23

MSE Mean Squared Error. 31, 39

PCA Principal Component Analysis. 9

RGB Red, Green, and Blue. xix

SPI Statistical Performance Indicators and Index. xii, 15

Stats SA Statistics South Africa. 9, 13, 41

UN United Nations. 7, 16, 20

USAID United States Agency for International Development. 9, 41

WLS Weighted Least Squares. 10

1 Introduction

On a flight inbound to Johannesburg, one sees an aerial visual tapestry of human life in Gauteng, South Africa. From the sky, contrasting living conditions are evident to the naked eye - urban centres rise amidst the rural outskirts of cities, and densely populated townships juxtapose leafy, affluent suburbs. As South Africa's economic powerhouse and the most populous province, Gauteng is a compelling case study for novel methods of measuring quality of life.

Understanding the quality of people's lives within a community is vital in modern society. "Quality of life" is a broad concept that fundamentally refers to how good or bad people's lives are. The exact definitions and contributing factors of a "good" quality of life have been widely debated [1]. Historically, income, asset wealth, and statistics like Gross Domestic Product (GDP) are the simplest proxies for quality of life. However, today, there is consensus that quality of life measures should extend beyond economic factors alone to account for multidimensional objective and subjective indicators of wellbeing. Factors such as access to essential services (water, electricity, internet, refuse removal, sanitation), healthcare, education, employment opportunities, safety, and subjective experiences of happiness and satisfaction all contribute to a more robust description of quality of life [1–5].

Regardless of the formulation, data about the various contributing factors is critical to measuring quality of life, and the benefits of measurement are multifold. Quality of life data attempts to reflect the wellbeing of residents and offers a baseline reference to identify challenges and opportunities for improvement at city, regional, and national levels. From this data, governments, advocacy groups, civil society, and international organisations can report on and address the state of people's lives.

What's more, meaningful action can be taken to improve individual and community livelihoods through evidence-based policies that enhance service delivery, promote infrastructure development, and support social needs [6, 7]. Over time, measuring quality of life provides an important feedback loop to monitor interventions taken by the government and for those in power to be held accountable [7, 8].

Traditionally, governments and organisations collect quality of life data through surveys and censuses. A census aims to enumerate every person within a defined enumeration area, while a survey aims to produce representative estimates for a target population by applying a sampling design and appropriate weighting [9]. For representative household surveys, probabilistic sampling methods are typically used and, where possible, data is collected via door-to-door interviews to maximise coverage and response rates [1, 7, 10]. Online surveys can lower cost and shorten field time compared with door-to-door approaches, but they pose risks to representativeness because of coverage bias, self-selection, and difficulties in verifying respondent identity or characteristics [11]. Although censuses are intended to achieve full enumeration, they nonetheless commonly face undercoverage and nonresponse and are undertaken less frequently because of their substantial cost and logistical complexity [9].

Once survey data is collected, responses can be aggregated in different ways to form meaningful models of quality of life. In practice, this typically involves the construction of a composite index that combines various reported dimensions into a single summary measure. However, the methodological approaches to creating such indices vary considerably and involve important theoretical and statistical decisions that can significantly impact the resulting measure.

The construction of composite indices requires several key methodological choices: the selection and weighting of indicators, the approach to handling missing data, the statistical techniques used for dimension reduction, and the methods for aggregating dimensions into a final index score [12]. For example, indices can be constructed using theory-driven approaches where dimensions are defined in advance based on existing literature, or through data-driven Exploratory Factor Analysis (EFA) using techniques such as factor analysis to empirically determine how indicators cluster together. A prime international example is the Human Development Index (HDI) [13], which uses a predetermined theoretical framework to combine health, education, and standard of living indicators. In the context of this research, a key local example is provided by the Gauteng City-Region Observatory (GCRO) who produces the GCRO Quality of Life Index/Survey/Dataset (GCRO QoL) index - a seven-dimensional, 33-indicator model. This model is derived through Exploratory Factor Analysis (EFA) that allows the South African context to influence the index construction, rather than imposing externally-derived theoretical frameworks [10, 12, 14].

Traditional survey-based methods of measuring quality of life are often costly, time-consuming, and face several challenges, especially in low- and middle-income countries (LMICs). An emerging field of research explores alternative methods for estimating select socioeconomic indicators from readily available data sources such as satellite images, social media, or mobile phone activity. This data is typically large and unstructured, rich in information about human behaviour and patterns. These big data sets can be used in conjunction with existing survey data to train machine learning models to estimate a given indicator.

1.1. Overview of Research

This research explores alternative methods for measuring quality of life, specifically focusing on the GCRO QoL index. The central research question is:

How well can a machine learning model estimate a quality of life index from satellite images in Gauteng, South Africa?

This question extends the broader research domain of using remote sensing technologies to assess socioeconomic development. While existing work has explored machine learning methods applied to remotely sensed data to measure a limited set of socioeconomic indicators, this research investigates the extent to which these methods can be applied to measuring multidimensional quality of life.

1.1.1 Research Objectives

The primary goal of this research is to design a tool that estimates the GCRO QoL index in Gauteng per ward from satellite images. The following research objectives are identified:

1. Identify high-quality data sources for this research.
2. Develop a reproducible data processing pipeline.
3. Develop a machine learning model suited to estimate the GCRO QoL from remote sensing data.
4. Evaluate the model using machine learning standards and baseline performance metrics.
5. Assess the feasibility and scalability of this approach.

1.2. Roadmap of the Dissertation

The remainder of this dissertation documents and discusses relevant work conducted to answer the research question. Chapter 2 provides background of the problem domain. Thereafter, a literature review of remote sensing applications for measuring quality of life is presented in Chapter 3. This review highlights the research significance and identified gaps inform the research scope, including assumptions and constraints taken into consideration. Chapter 4 documents the research scope, along with the research methods applied, including data sources, data processing methods, the machine learning model employed, and evaluation techniques applied. The results of the applied research methods are presented in Chapter 5. Chapter 6 provides an interpretation and analysis of the results and methods. This chapter discusses the implications of the findings against the broader backdrop of using remote sensing for measuring quality of life is shared. Finally, Chapter 7 summarises the contributions, reflects on the overall study, and proposes recommendations for future research.

1.3. Research Significance

This research explores whether satellite imagery can supplement traditional survey methods for measuring quality of life. The work is motivated by a vision of a future in which frequent, reliable, low-cost data about the quality of people's lives is readily available and can be used to meaningfully improve livelihoods.

While traditional survey-based methods remain the gold standard for quality of life measurement, their limitations in resource-constrained environments create opportunities for complementary approaches. This work investigates what satellite imagery can reveal about multidimensional human wellbeing and whether remote sensing might enhance conventional data collection.

The significance extends beyond methodology alone. In places like Gauteng, where rapid urbanisation meets persistent inequality, more frequent, spatially comprehensive quality of life data could strengthen evidence-based policymaking and resource allocation.

2 Understanding the Problem

This chapter contextualises the problem domain, providing necessary background to the conducted research. The definitions of quality of life are further elaborated on, and more details are provided on traditional methods of measurement. Challenges with these traditional methods are highlighted. From these challenges, a clear problem statement emerges: LMICs face a significant data gap compared to their wealthier counterparts. Closing this gap with traditional, survey-based methods alone is not easy or cheap. As an alternative, novel methods of measurement from remotely sensed data have emerged, which are discussed further in Chapter 3.

2.1. A Brief History of Quality of Life Measurement

Quality of life is a broad, subjective concept that fundamentally describes how good or bad people's lives are. But what does a "good" quality of life look like? This is a question that has occupied philosophers for millennia. The experience of a "good" life is dependent on one's material living standards and external environment. Does one have access to water, electricity, sewage and waste services? What kind of roof and building material is one's dwelling made of? What is one's employment status and income bracket? However, beyond these objective factors, there are also subjective factors of what wellbeing means to an individual or community, influenced by social norms, culture, beliefs, and personal values and aspirations. A single, consistent definition for quality of life, therefore, does not exist [1]. Instead, predominant theories describing quality of life have emerged and today, there is consensus that the concept describes multidimensional human wellbeing [1, 6, 15, 16].

The measurement of quality of life has evolved over the past century as broader societal values have shifted to prioritise the desire for wellbeing, equality, equity, and basic living standards for all [1, 17]. In the early 20th century, economic indicators such as GDP and income levels were proposed as objective measures of a nation's development and progress. Thinkers of the time drew direct correlations between a nation's prosperity and material living standards. Higher income levels typically meant better access to goods and services, improved housing, enhanced healthcare, and greater educational opportunities. As a result, economic indicators were seen as proxies for overall quality of life [1, 6, 16, 17].

By the mid-20th century, GDP was the de facto measure for comparing economic performance of nations. It was widely adopted by governments and international organisations like the World Bank and International Monetary Fund [16–18]. The metric presented an opportunity for quantifying quality of life through a measurable, numerical proxy. However, GDP, and its byproduct, GDP per capita, is limited by the emphasis on economic output and prosperity through production. It does not account for inequality (how income is distributed), the quality of goods and services, environmental sustainability of production, or non-economic factors that impact perceptions of the quality of life [16, 18].

The gaps in economic measures to describe quality of life motivated the development of the capabilities approach led by Amartya Sen [1, 15]. His work represented a fundamental shift away from evaluating prosperity and wellbeing based on the current economic success of individuals within the society. Instead, Sen proposed evaluating quality of life according to what people **can** be and do [1, 15]. In this context, quality of life becomes about the opportunities and resources an individual could have and not only about what they currently own or earn, offering a more nuanced and human centred-perspective on development.

Sen's capabilities approach, developed in the 1980s, was key in informing the HDI [13]. Developed by the United Nations (UN) [17, 19], the HDI is constructed by aggregating proxies of access to health, education, and goods. Although Sen criticised HDI as another crude composite index [1], the index acted as a revolutionary alternative to GDP per capita and popularised the need for measuring multidimensional quality of life.

2.2. Traditional Methods of Measuring Quality of Life

Traditional methods for measuring quality of life predominantly rely on quantitative and qualitative data gathered through surveys, censuses, and administrative records. From these primary data sources, researchers can derive various socioeconomic indicators that capture different dimensions of wellbeing. This data may be collated in different ways—depending on the approach taken—to construct a composite index that summarises quality of life.

However, there are fundamental differences in how these indices are constructed and applied spatially. For example, the HDI mentioned previously is calculated at an aggregate national level, using pre-aggregated administrative data and survey statistics relating to health, education, and income [13]. In contrast, the GCRO QoL index is constructed at the individual respondent level from survey responses, which can then be aggregated to different spatial scales such as wards, municipalities, or provinces [12]. This distinction between individual-level and aggregate-level index construction has important implications for spatial analysis and the potential for satellite-based measurement, as individual-level indices can be flexibly aggregated to match the spatial resolution of remotely sensed data.

2.2.1 Surveys

Surveys are a cornerstone in assessing quality of life—whether uni- or multivariate. They collect detailed information from a sample of the population on income, employment, education, health, and housing conditions. These surveys can also include subjective wellbeing questions, where individuals self-report their satisfaction with different aspects of life, providing valuable insights into personal perceptions of livelihoods [4, 20].

Importantly, surveys employ carefully designed sampling frameworks that aim to accurately represent demographic, socioeconomic, and geographic diversity within the target population. Without this intentional design, results may systematically exclude certain groups or overrepresent others, introducing bias that undermines the validity of findings [7]. The questionnaire itself must also be developed to ensure questions are culturally appropriate, accessible across education levels, and capable of capturing the full spectrum of experiences within heterogeneous populations [7].

2.2.2 Censuses

Censuses are a particular type of survey where the sample aims to include all members of a population [21]. Due to the large scope of this method, censuses are usually conducted by national statistic offices once every five to ten years. They capture essential information such as age, gender, income levels, housing conditions, and employment status across large geographic areas. This extensive data allows for the analysis of trends and disparities within regions, enabling policymakers to identify areas needing improvement and allocate resources effectively.

2.2.3 Administrative Data

In addition to surveys and censuses, administrative data can also be used to measure quality of life. Various economic metrics can be gathered from the public and private financial sectors and unemployment rates, birth rates, life expectancy, and literacy rates can be collected from government departments (e.g. Department of Home Affairs, Department of Basic Education, Department of Labour). Together, this administrative data contributes to a multifaceted view of a population's wellbeing with less intervention than surveys or censuses. However, administrative data alone does not paint a full picture and is heavily dependent on what and how data is collected by government departments.

2.3. Relevant Survey Examples

Specific examples of surveys relevant to this dissertation are discussed below.

2.3.1 Demographic and Health Surveys

The Demographic and Health Surveys (DHS) programme, funded by USAID, is a primary source of data on population, health, and nutrition in LMICs [22]. The programme supports a local agency, for example, Statistics South Africa (Stats SA), in conducting the DHS survey and consolidating the results. One component of the DHS datasets is the asset wealth index, which is constructed by applying a Principal Component Analysis (PCA) to survey responses on household ownership of various assets, such as televisions, bicycles, and housing materials. The DHS asset wealth index is used in previous research in the field (see Section 3.3).

2.3.2 Living Standards Measurement Study

The Living Standards Measurement Study (LSMS) is a survey initiative by the World Bank, designed to collect detailed data on various aspects of household well-being in LMICs [23]. A primary indicator of economic welfare and living standards in the LSMS surveys is consumption expenditure, a measure of consumer spending and purchasing power. Together with asset wealth, consumption expenditure can be useful in assessing poverty [6, 24]. As with the asset wealth index above, consumption expenditure is used in previous research in the field (see Section 3.3).

2.3.3 Gauteng City-Region Observatory Quality of Life Index

Established in 2008, the GCRO is a partnership think tank between local universities and government in Gauteng, South Africa. The organisation's primary objective is to support data collection and use for policy and research in the country's most densely populated province [25]. The GCRO conducts a survey on multidimensional quality of life every two years (GCRO QoL surveys), covering a broad range of topics including household circumstances, service delivery, health, education, safety, employment, and life satisfaction. Using the survey data, a single composite index is calculated (hereafter the GCRO QoL index) [10, 12]. Given the importance of this index for the remainder of the dissertation, a deeper dive into the index methodology is provided below.

Index Construction and Methodology

The GCRO QoL index represents a tailored approach to measuring multidimensional quality of life, specifically designed for the South African context. Unlike many international quality of life indices that apply predetermined theoretical frameworks developed primarily in the Global North, the GCRO employs an explicitly context-sensitive approach that allows local conditions and priorities to influence index construction [12].

The GCRO derives a 33-variable, seven-factor model of quality of life from the GCRO QoL survey data, which has 58 questions. The approach employs EFA and Weighted Least Squares (WLS) to generate the model. The seven dimensions that emerge from the factor analysis are: infrastructure and services, socioeconomic status, government satisfaction, life satisfaction, safety, health status, and political engagement [12, 26]. The dimensions are then aggregated into the composite GCRO QoL index, weighted by their eigenvalues rather than being equally weighted, and scaled to run from 0 to 100 [12]. This approach ensures that dimensions with greater explanatory power contribute more substantially to the overall index.

Objective and Subjective Components

The GCRO QoL index incorporates both objective and subjective indicators as reported by respondents [10, 12]. Objective measures include tangible aspects such as access to services, employment status, housing conditions, and income levels. Subjective measures capture individuals' experiences and perceptions, including life satisfaction, feelings of safety, and satisfaction with government services. This combination creates a comprehensive measure that extends beyond purely economic indicators to capture the multidimensional nature of wellbeing.

Validation and Temporal Consistency

The GCRO has validated the index model through Confirmatory Factor Analysis (CFA) and tested for invariance across multiple survey iterations (Rounds 3 to 5), confirming that the pattern of indicators among factors remains stable over time. This validation demonstrates the model's ability to generate longitudinally comparable results, making it suitable for tracking changes in quality of life across survey cycles [12].

Methodological Significance

The GCRO's approach represents a significant methodological contribution to quality of life measurement in middle-income contexts like Gauteng, South Africa [12]. By prioritising empirical exploration over predetermined theoretical frameworks, the methodology allows for the emergence of locally relevant dimensions that might be overlooked by indices developed in different contexts. This context-sensitive approach provides a model for quality of life that is locally significant and relevant for applied research, as in the case of this dissertation.

2.4. Challenges of Traditional Methods of Survey-Based Measurement of Quality of Life

Survey-based methods for measuring quality of life, or sub-indicators thereof, face a number of challenges, especially in LMICs. These challenges are discussed below, setting the scene for the problem statement that this research addresses.

2.4.1 Cost

Representative surveys that produce high-quality data are expensive and time-consuming, often requiring substantial financial investment over multiple years [27–29]. For example, the GCRO QoL survey, despite being geographically limited to a single province in South Africa, allocated a budget of R 16.5 million (excluding VAT) for the 2023/24 iteration to conduct a minimum of 12,000 face-to-face interviews across Gauteng's 529 wards [30]. This substantial cost encompasses fieldworker training and remuneration, survey management, data collection infrastructure, quality control processes, and analysis. The financial burden is particularly challenging for LMICs where resources for data collection are already constrained. Even with significant investment, there remains considerable risk that the high costs may not necessarily result in high-quality data, as evidenced by the GCRO's own experience of budget overruns and the need to discard substantial numbers of interviews due to quality issues in previous survey iterations [26, 31].

2.4.2 Logistics

Beyond the financial and time cost of conducting surveys, the logistics are also complex, even in the case of smaller-scale surveys. For example, the GCRO QoL survey requires extensive coordination and planning across Gauteng to ensure effective sampling and household visits. Hiring and training surveyors requires preparation and investment.

2.4.3 Data Processing and Analysis

Once survey data is collected, the consolidation and analysis of survey results presents additional challenges. This process involves complex statistical procedures including data cleaning, weighting to ensure representativeness, handling missing values, and applying appropriate analytical techniques. For composite indices like the GCRO QoL index, this extends to methodological decisions about indicator selection, dimension reduction, and aggregation methods. These analytical processes require specialised expertise and can be time-consuming, adding to the overall complexity and cost of survey-based quality of life measurement.

2.4.4 Frequency and Timeliness

Due to the cost and logistics of surveys, the frequency of data collection is often constrained. This is true in most LMICs in sub-Saharan Africa, including South Africa. While survey programmes often maintain regular schedules, such as the GCRO QoL surveys (every two years) [26], Multiple Indicator Cluster Surveys (MICS) (every three to five years) [32], and DHS surveys (every five years) [22], there are often delays and irregular implementation cycles. For example, the GCRO QoL survey, originally planned as a biennial survey, has experienced increasing timing challenges due to survey size and resource constraints [31]. Irregular frequency of representative household surveys may result in outdated data being used to inform policy decisions and assess quality of life trends. Additionally, even when surveys are conducted on schedule, the substantial lag time between data collection, processing, and analysis can delay the implementation of necessary interventions by months or even years, reducing the overall impact of policy initiatives. The timeliness challenge is particularly acute for rapidly changing urban environments like Gauteng, where socioeconomic conditions can shift significantly between survey cycles, making older data less relevant for contemporary decision-making.

2.4.5 Sampling

Ensuring that sampling accurately reflects the population's diversity is another challenge of survey-based methods. National statistics offices like Stats SA and research organisations like the GCRO take precautions through careful stratification and weighting techniques when sampling [7]. However, disparities in literacy rates, language barriers, and varying levels of trust in governmental institutions can still lead to sampling biases and non-response rates. Marginalised groups, including those in informal settlements or remote rural areas, may be underrepresented, skewing the results and limiting the generalisability of findings.

2.4.6 Data Quality and Accuracy

Given the resource and sampling constraints in LMICs, a pertinent question arises about the quality of socioeconomic statistics collected in these environments. In his book, *Poor Numbers*, Jerven [28] argues that developing economies are poorly modelled and reported on, because of limited funding and low state capacity. Although Jerven‘s argument applies to national statistics, similar questions are posed about localised surveys as well. For example, examination of the GCRO QoL survey methodology reveals tensions between administrative boundaries and natural population clusters. While the survey employs a sophisticated multistage approach using probability proportional to size sampling within enumeration areas [33], the overall sampling framework remains constrained by ward boundaries - South Africa's lowest administrative division [26]. These divisions are dependent on the number of voters within a particular region. However, voter clusters may not necessarily align with natural population clusters, potentially affecting the precision of local estimates. This is a consideration that must be kept in mind when using representative surveys.

2.4.7 Safety of Fieldworkers

The safety of fieldworkers is a relevant challenge to traditional based survey methods. In a South African context, high crime rates in certain areas are a risk to researchers and surveyors, potentially limiting their ability to collect data effectively [34]. For example, questionnaires are often completed with smartphones or tablets, prime targets for robbery. Fieldworkers may also face animosity and harassment depending on community perceptions [34].

2.5. The Widening Data Gap

Against the backdrop of complex historical, institutional, economic, and political nuance, the challenges of traditional survey-based methods leads to a widening data gap between developed and developing nations. Many of the challenges disproportionately affect developing nations [28, 29]. LMICs often face more severe resource constraints, making costly census and survey operations difficult to fund and sustain [28]. Logistical challenges are amplified by inadequate infrastructure and transport, particularly in remote regions. Sampling biases may be compounded in countries where population registries are incomplete or outdated, or where rapid urbanisation creates informal settlements that are difficult to enumerate systematically. Additionally, safety concerns for fieldworkers can be heightened in areas with political instability or conflict, a common feature in less developed regions. These amplifications of traditional survey methods set many developing nations behind in their data capabilities [28, 29]. A key consequence of this data gap is limited insights into quality of life in LMICs.

A useful metric for reporting on the data gap is the World Bank's Statistical Performance Indicators and Index (SPI), which is a proxy for a country's statistical capacity [35, 36]. SPI considers five dimensions: data use, data services, data products, data sources, and data infrastructure. The average statistical capacity of sub-Saharan Africa ranks significantly lower than most European, North American, and Middle Eastern countries (see Figure 2.1) [37]. Improving SPI scores requires more than simply collecting additional data. It demands systemic improvements in statistical infrastructure, institutional frameworks, and human capital development that extend well beyond any single measurement challenge.

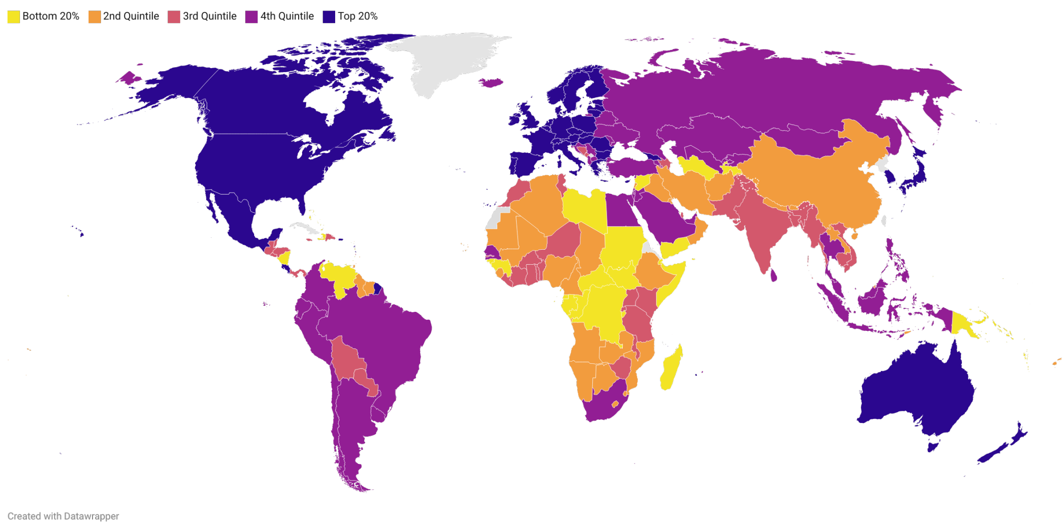

Figure 2.1: Global SPI score 2023 [37]

Addressing this data gap fundamentally requires strengthening institutions, developing statistical skills, and building capacity. These are investments that are valuable for countries regardless of data collection needs. However, these institutional developments require substantial time and sustained investment. As Jerven [29] notes, the financial and opportunity costs of comprehensive data collection through traditional survey methods are considerable, particularly given the exponential increase in data demands from initiatives like the UN Sustainable Development Goals. While building strong statistical institutions remains the long-term solution, there is value in exploring complementary approaches that could help bridge data gaps in the interim, particularly for countries where institutional capacity building is still underway.

2.5.1 Potential Solution: Novel Methods of Quality of Life Measurement

To address the widening data gap, a field of research has emerged to explore alternative methods for measuring quality of life. Big data sources like satellite images, social media, and mobile phone activity are rich in information about patterns in human lives and wellbeing. These data sources offer potential advantages including more frequent updates and broader spatial coverage compared to traditional surveys. At the same time, the scaling of global computational capacity has made machine learning methods increasingly viable for analysing and processing large, complex datasets. At this intersection, an interdisciplinary field of research exists where machine learning methods are applied to big data sources for measuring socioeconomic development, quality of life, and wellbeing.

2.6. Chapter Summary

This chapter has described the evolution of measuring quality of life, and the challenges faced with traditional methods. In middle-income contexts like South Africa, there is growing recognition of the need for high quality, frequent data about populations' multidimensional quality of life, as evidenced by initiatives like the GCRO QoL surveys. While basic poverty and development indicators remain priorities in many LMICs, understanding broader dimensions of wellbeing becomes increasingly important as countries develop institutional capacity and seek to address more complex aspects of human flourishing. However, traditional methods are slow and costly. Novel data collection methods using remotely sensed data offer a promising avenue to explore for supplementing survey-based quality of life measurement.The extent to which these methods can capture multidimensional aspects of quality of life remains an open research question that this dissertation seeks to address. This dissertation is situated within this context. The next chapter will explore how satellite images and machine learning techniques have been applied in prior work in the field.

3 Literature Review

The field of measuring quality of life from remotely sensed datasets is relatively new, having emerged in the last 30 years [27, 38, 39]. Although there is relatively limited prior research that explicitly references quality of life [27, 38], much work has been completed in the domain of measuring various socioeconomic development indicators that can be construed as proxies for quality of life [27]. For this reason, this literature review focuses on machine learning methods applied to satellite images to predict proxies for human wellbeing—whether economic or multidimensional. Additionally, some data fusion methods are touched on that incorporate additional data sources.

Satellite imagery is a primary remote sensing dataset that is readily available and frequently collected and used as inputs to machine learning models for measuring socioeconomic development indicators [27, 38, 40]. The application of deep learning to satellite imagery analysis became feasible in the mid-2010's, following the development of ResNet architecture [41] and the availability of large-scale labeled datasets like ImageNet. Other data sources include mobile network data, Wikipedia articles, land cover maps, OpenStreetMaps, and Facebook marketing data [40]. Much of the research in this domain has focused on satellite images as the primary data source and then expanded this to incorporate additional datasets. This literature review will therefore highlight work conducted with satellite images and touch on additional data sources that are relevant in this field.

The themes of prior work are presented through the lens of what is measured and how measurement is conducted. By discussing what has been done, a few gaps in the literature emerge. These include scope to explore the prediction of multidimensional indexes, the focus on nighttime light data as a primary methodological approach, and the lack of research in a South African context.

3.1. Overview

Prior research in the field explores estimating a range of socioeconomic variables from satellite images. Generally, work focuses on economic activity and poverty [27]. However, some multidimensional proxies for human development and quality of life have also been considered [42–44]. This section analyses key themes of variables measured, along with the machine learning methods used for estimation from remote sensing data.

3.2. Measuring Economic Activity

Early work in the field, beginning in the late 1990s, reported a high correlation between nighttime light data and GDP [45–47]. This finding emerged from research into estimating the spatial distribution of population from satellite images. Elvidge et al. [45] showed that although population is not well-correlated with nighttime satellite images, GDP is, opening the doors to measuring economic development from remotely sensed data.

Correlations between nighttime light data and economic activity laid the foundation for much of the subsequent work in this field [27, 47]. However, the underlying premise of this research is that global, frequent daytime satellite imagery is less available and high cost. Historically, daytime satellite images were expensive and high-effort to obtain, especially for global coverage [48]. In contrast, nighttime light data is a fraction of the cost because it can be collected at lower resolution using simpler sensors, originally designed for weather monitoring [49]. This data requires much less data storage due to smaller file sizes, and is freely available from sources like the NOAA Defense Meteorological Satellite Program [49]. This made nighttime lights data a preferred data source for early poverty mapping research. However, programmes like Landsat have revolutionised the use of daytime satellite imagery by providing archival data at no cost [48]. Many other data sources (e.g. Planet [50]) provide high coverage and often at a low cost, especially for research purposes [27]. This has sparked lines of research that broaden the scope of remote sensing applications. Nonetheless, given the importance of nighttime light data in the field, it is valuable to review methods that rely on this datasource.

A similar prominent economic variable is average household income [51–53]. Input datasets from which predictions are made, include daytime satellite images, nighttime light data, Google Street View, and land use.

Although GDP and income are valuable economic measures, they do not sufficiently reflect multidimensional quality of life (see Chapter 2). Estimating GDP from remote sensed data supports valid needs to understand economic activity and growth of countries or regions. However, in the context of assessing quality of life, GDP does not account for income distribution or environmental, social, or subjective factors that impact wellbeing.

3.3. Mapping Poverty

Poverty mapping is a significant area of study for estimating indicators of socioeconomic development from remote sensing data [27, 38]. One of the primary UN Sustainable Development Goals is to end poverty by 2030 [54]. This has driven efforts to reliably map poverty frequently, especially from data sources like satellite images, based on the success of understanding economic activity from these sources [27, 55, 56]. Previous work in the field has focused on estimating poverty metrics from a range of remote sensing data sources using machine learning models. This includes daytime satellite images, nighttime light data, geolocated Wikipedia articles, and mobile phone network data [27, 55–57].

3.3.1 Defining Poverty

The definition of poverty is complex and contested. As Everatt [58] notes, despite poverty being a central focus in South Africa's post-apartheid development agenda, "the country has spent years fighting over how to define poverty, and thus how to measure and understand it." While poverty is related to quality of life, it represents a specific lack of resources required to meet people's needs [58]. In this definition, multidimensional poverty extends beyond purely economic resources, though these remain important. Various approaches to poverty measurement include income and asset poverty, calorific intake measures, and sustainable livelihoods approaches that focus on "capacities and potentialities the poor have, rather than on the assets and income they lack" [15, 58]. Furthermore, poverty also crosses into the "psychosocial realm" [58], where factors like alienation, anomie, and social exclusion play a significant role, but are not necessarily directly observable.

Despite these multidimensional considerations, it is common to assess poverty according to an economic baseline of income or wealth [59]. One such tool for this is poverty lines, which define a minimum income necessary to achieve an adequate standard of living. For example, the international poverty line is currently USD 2.15 per day, below which individuals are considered extremely poor [60]. Similarly, poverty can also be assessed relative to asset wealth indices and consumption expenditure [24].

3.3.2 Transfer Learning with Nighttime Light Data

Beginning with Jean et al. landmark 2016 study, one primary method of measuring asset wealth and consumption expenditure is a multistage transfer learning pipeline using both day and night satellite images [27, 55, 61, 62]. First, a pre-trained Convolutional Neural Network (CNN), typically initialised with ImageNet [63] weights, is fine-tuned to predict nighttime light intensity from satellite images in the visual spectrum. This machine learning model is used to extract features from the same satellite images. These features are then input as training data to a regression model to predict household expenditure and an asset wealth index. The approach has been found to explain up to 80 % in the overall variation of these outcomes [27, 38, 40].

Importantly, the transfer learning approach does not rely on manually engineered features for estimation. The model is trained to predict nighttime light intensities from daytime satellite images without requiring human annotations to guide the model's "attention". This approach significantly reduces the need for labour-intensive preprocessing and annotation of satellite imagery. Instead, the model autonomously learns relevant features that correlate with nighttime light patterns.

3.3.3 No Transfer Learning

Building on poverty mapping with transfer learning approaches, Yeh et al. [56] trained machine learning models to estimate asset wealth across Africa from satellite imagery. However, nighttime light data was used as an additional data input for training rather than an intermediate target. Additionally, multispectral satellite images were used, including several bands beyond the visual spectrum. This model performed similarly to the transfer learning based approach even with limited training data availability [56].

3.3.4 Object Detection

Some research has explored object detection methods for measuring economic indicators useful for poverty mapping [57, 64]. Although these methods are comparable to the transfer learning and end-to-end methods described above, they rely on intensive satellite image annotation. In fact, Engstrom et al. [64] outsourced the annotation of the satellite images in their dataset to a private company.

3.3.5 Additional Combination Methods

Several multimodal machine learning methods have also been explored for poverty mapping, predominantly in LMICs. These include incorporating datasets such as mobile phone data [65, 66], geolocated Wikipedia articles [67], Facebook data [68, 69], land cover data [70], and Open Street Maps data [40, 71]. Although these methods are interesting, the scope of this dissertation is constrained to satellite images, as discussed further in Chapter 4.

3.4. Beyond Poverty

A primary focus of existing research in the field has been on measuring one or two economic indicators from satellite imagery, such as GDP per capita, asset wealth, and consumption expenditure [27, 38, 47, 55, 56]. In these cases, there is often a claim of measuring “poverty”, which is problematic, given the nuances in defining poverty, as described in Section 3.3.1. Nonetheless, while measuring economic indicators from remotely sensed data provides valuable insights, quality of life encompasses multiple dimensions beyond material wealth. Some prior research into measuring composite indices beyond poverty has been undertaken, as discussed below. However, this research is limited, providing opportunities to consider multidimensional quality of life measurement from remotely sensed data.

3.4.1 MOSAIKS: Generalising Feature Extraction

To generalise beyond poverty, Rolf et al. [72] considers a novel approach to feature extraction from daytime satellite images. Specifically, they designed Multi-task Observation using Satellite Imagery and Kitchen Sinks (MOSAIKS)—a method of extracting random convolutional features from satellite images. These features are task agnostic and can be used to train a regression for various outcomes. This is a big step forward in reducing computing requirements for training a deep learning model to detect objects or extract features based on the specific target variable. The approach has been shown to estimate nighttime lights, household income, population density, and housing price [72].

However, an important constraint is that MOSAIKS is trained once off on a single set of satellite images out of which the features are used across different timings [72]. In other words, the time of the features (2019 for the original paper) might not match the times of the actual target. Retraining could be undertaken but at a similar cost to training other more traditional deep learning models [72].

3.4.2 Indicators of Human Development

One extension of MOSAIKS is an application to estimating a global map of HDI [73]. However, at the time of writing, the research is in a working paper format and has not been peer-reviewed to the author's knowledge. Additionally, the working paper relies on the 2019 MOSAIKS features and does not retrain the feature extraction on new satellite images [73].

An alternative approach to measuring multidimensional human development from satellite images attempts to apply transfer learning methods with nighttime light data to estimating a range of indicators in the DHS surveys [42]. Although economic indicators were consistently predictable, most other indicators that would contribute to multidimensional quality of life were found to perform poorly [42].

3.4.3 Inequality

Similar research has been conducted in the United Kingdom, where multiple socioeconomic outcomes were estimated from the same dataset and algorithm [44, 74]. This allowed for comparisons between various outcomes within the same spatial and temporal resolution. However, in Suel et al.'s research [44, 74], socioeconomic data for training the machine learning model was more readily available for the indicators estimated at a granular spatial resolution. Access to granular spatial micro-estimate data is a constraint in South Africa and many other LMICs.

The author has explored analogous preliminary work into visualising inequality by mapping street names based on country of origin [75]. However, the goals in that work were not to measure proxies of inequality. Instead, the approach aimed to uncover patterns in the identity of places and how those patterns correlate to population density and average income.

3.5. Research Gaps and Opportunities

In reviewing the existing literature, several research gaps and opportunities are identified. These gaps highlight the limitations of prior work and pave the way for contributions to research in the field in the context of Gauteng, South Africa.

3.5.1 Gap Analysis in Existing Literature

- **Limited Focus on Multidimensional Quality of Life:** Most prior research concentrates on economic indicators such as GDP and income, with limited attention paid to multidimensional aspects of quality of life that encompass health, education, and environmental factors.
- **Uniform Socioeconomic Landscapes:** Prior studies have focused on relatively uniform socioeconomic landscapes, including either developed regions like the United Kingdom or the United States of America, or rural developing regions (e.g. Burundi or Rwanda) [27, 44, 56]. There is limited research in more complex and unequal urban landscapes like Gauteng, South Africa.
- **Over-reliance on Nighttime Light Data:** The predominant use of nighttime light data as a proxy for economic activity may overlook other relevant features available in daytime satellite images.
- **Limited Integration of Local Surveys and Data Sources:** Current methodologies primarily use large-scale surveys such as DHS, which may not fully capture localised or niche socioeconomic variables pertinent to specific communities or regions.

- **Temporal Constraints:** Most studies provide static snapshots of economic indicators, neglecting the dynamic nature of socioeconomic changes over time, which is crucial for understanding trends and the impact of policy interventions.

3.6. Chapter Summary

This literature review has examined the emerging field of measuring quality of life and socioeconomic indicators from satellite images using machine learning approaches. The review highlights how research has progressed from early correlations between nighttime lights and GDP to more sophisticated methods including transfer learning, object detection, and multimodal approaches incorporating additional data sources.

While significant work has been done in mapping economic activity, asset wealth, and consumption expenditure, the dominant focus on these indicators allows room to explore measuring multidimensional aspects of quality of life instead. Furthermore, by exploring this research in Gauteng, South Africa, an opportunity arises to probe the boundaries of existing machine learning methods in localised contexts.

The continued contribution to research in this domain has significant implications for policymaking and development in both local and global contexts. Exploring novel methods of measuring quality of life tests the possibilities of moving beyond costly traditional data collection methods. In developing regions, this could empower governments to rapidly monitor and rollout meaningful change for citizens.

4 Research Methods

This chapter describes the methods and experiment setup used to evaluate the question: "how well can a machine learning model estimate a quality of life index from satellite images in Gauteng, South Africa?" The chapter begins with a brief discussion of the overarching approach taken, the scope of research, assumptions, and constraints. Thereafter, the methods for producing results are described. This includes details of the collection and processing of data, the specifications of the design and training of the machine learning model, and the methods and metrics used to evaluate the machine learning model. Finally, the chapter describes the system environment and tools used to execute these methods.

4.1. Research Approach and Methodology

This study primarily follows a quantitative research approach. Data is collected and processed systematically from verified data sources. Modifications to the data are documented in the codebase and version control is used for code, data, and results [76]. The model and experiments are designed to be reproducible. Numerical and statistical results are collected and compared to answer the research question.

To complete this research in good time, an adapted agile methodology is used. Agile is a set of principles guiding modern software development that prioritises functional software over comprehensive documentation, collaboration over processes and tools, and responsiveness to changing requirements over a rigid plan [77]. In this context , functionality is analogous to producing reproducible and valid results. Additionally, collaboration has driven meaningful progress and learning through the course of the research. As with any research project, responding and adapting to change is a critical practice to complete the work required.

4.1.1 Ethical Considerations

There are several ethical aspects to consider when using human survey data and satellite images for the purposes of improving understanding of quality of life.

First, there is a risk of human survey data revealing personal and private information about individual and household social and economic conditions. However, in this case, the original data source is anonymised by the publisher such that all directly personal identifying information is removed. Additionally, the survey data does not reveal exact survey locations. Instead, survey responses are only spatially identified by a ward code corresponding to the ward in which the response is collected. This means that the lowest level by which one can locate survey responses is at the ward-level, the smallest of which can be seen in Figure 4.1 with an area of 0.227 km^2. Reverse engineering survey responses to specific individuals or households within a ward remains a difficult task, even in a ward of this size. In fact, there would be higher risk of identifying survey response in wards with low population density, where households are far and few between. Nonetheless, by removing individual identifying information from the survey data, survey data publishers ensure that collected data does not reveal personal private information about individuals or specific households.

Second, a trade-off exists between the ethics of using satellite image data and the benefits that these use cases may bring to society. Applications using satellite image data may infringe on the privacy of individuals and groups depending on the spatial resolution and frequency of image capture [78, 79]. Currently, available satellite image data is not of a high enough resolution to identify individuals [80]. However, identifying information about households or groups of people **can be** extracted from satellite image data. This information could be used to discriminate against individuals or groups of people.

Legislation for collecting satellite image data is typically handled by national governments and are not consistent across countries [79, 81]. Obtaining individual consent for collecting satellite image data is not a legal requirement anywhere in the world [79]. Thus, visual or inferred information about individuals or communities that is extracted from satellite image data is not voluntarily disclosed.

Figure 4.1: Ward 7980059: Spatially the smallest ward in Gauteng. Image ©2024 Planet Labs PBC.

In the case of this research, understanding the possible solutions for regular, low-cost estimation of aggregated quality of life from satellite images outweighs the risk of infringing on individual or group privacy. There are no direct human participants, obtained human data is anonymised, and individual identifying features from satellite images are not extracted. The satellite image data is used as an input for estimating quality of life, but no action is taken from that estimation.

This research is categorised as a *"no risk"* project according to the risk table in Appendix B. The University of the Witwatersrand's Human Research Ethics Committee reviewed this categorisation and granted an ethics waiver for the research (see Appendix B). Additional ethical analysis is required if the results of this research is used to take action (e.g. target aid) in particular regions.

4.2. Scope of Research

Based on the identified gaps in the literature, the research scope is narrowed to designing a machine learning model that predicts a specific quality of life index from satellite images in Gauteng, South Africa. The scope is reflected in the research question: "how well can a machine learning model estimate a quality of life index from satellite images in Gauteng, South Africa?" The primary hypothesis is that some correlation exists between the selected GCRO QoL index and visual features in satellite images.

4.2.1 Assumptions

Several assumptions are made that underpin the conducted research. These assumptions influence various decision points in the research, and thus, are important to highlight.

First, it is assumed that unsupervised computer vision algorithms can identify patterns in satellite images that meaningfully correlate with quality of life, without being explicitly trained on target features. This is similar to a human's interpretation of an image where one might not explicitly identify all the factors observed, but characteristics of an image are extracted. For example, in a satellite image with roads and greenery (see Figure 4.2), human observation can pick out roads, greenery, and buildings, without explicit labelling. The assumption is that a machine learning model can extract this kind of relevant information, without explicitly being trained to do so. The assumption is based on previous methods, which have shown that machine learning models can extract features relevant to a target indicator, without supervised training of what features to look for.

A related assumption is that there is sufficient information contained in the visual spectrum of a satellite image to discover relevant patterns that correlate to quality of life. This assumption is based on the premise that satellite images have visual features that correlate to economic wellbeing, as found in prior literature (Chapter 3). Therefore, satellite images should have at least some visual information that correlate to multidimensional quality of life, which incorporates aspects of economic wellbeing.

Although quality of life and visual features being correlated is assumed, there is an expected limitation of the extent to which visual features can describe multidimensional quality of life. Some critical factors, such as social cohesion, happiness, and education may not be directly observable from visual information about a place. The extent to which these factors may affect quality of life measurement from satellite images is a complementary question to the research conducted.

Figure 4.2: An example of a satellite image with various features visible to the human eye. Image ©2024 Planet Labs PBC.

4.2.2 Constraints

A constraint on the machine learning methods applied is imposed to ensure the research scope is limited. Specifically, the conducted research applies machine learning methods, rather than designing entirely novel methods for predicting quality of life. Therefore, the research is focused on the intersection of quality of life measurement from satellite images with existing machine learning methods.

Additionally, the research is constrained by the data available. For example, the satellite images available have a predetermined spatial resolution that enforces the lowest level of detail visible per pixel. Similarly, the GCRO QoL index is provided per individual response, but spatially, the lowest administrative region that this can be mapped to is the ward-level, based on the data available.

A noted practical constraint is time and resources. The research is conducted part-time and is required to be completed within a maximum of three years. A limited number of methods and datasets can be experimented with during this period. Additionally, a single researcher is assigned to this work, which also limits the scope of experimentation. Alternatives seen in prior work, such as outsourcing of specific components (e.g. labelling/algorithm design) or a research group dedicated to components of the overarching problem, is either not permitted or not feasible in the research context.

4.3. Methods Overview

A machine learning model is designed to estimate the GCRO QoL index given a satellite image. Figure 4.3 illustrates an overview of the pipeline.

Two key datasets are required: the GCRO QoL survey data [82, 83] and daytime satellite image data obtained from Planet [84]. To collate these two datasets, geographic data mapping ward boundaries in Gauteng is required. This geographic data is obtained in shapefile format from the South African Municipal Demarcation Board (MDB) [85]. The three datasets are processed independently and then combined into a new dataset that includes satellite image tiles for each ward with the corresponding GCRO QoL index for that ward.

The combination dataset is split into training, and test sets, as is the standard for machine learning pipelines. This is an important step to mitigate the risk of overfitting the data during training. A pre-trained machine learning model is fine-tuned to extract features from the satellite image tiles. The output feature vectors are passed directly into a fully-connected Dense Layer to estimate the GCRO QoL index for the corresponding satellite image tile.

Experiments are run with the machine learning model under different conditions and evaluated using k-fold cross-validation. Mean Squared Error (MSE), Mean Absolute Error (MAE), and the Coefficient of Determination (r^2) are consistently collected for each experiment to produce results. Specifically, the MSE in this end-to-end machine learning pipeline is used as the loss function for training. The r^2 value on a hold-out test set is used to compare to baseline metrics identified in previous work.

The code used to execute these methods and reproduce results is available [76].

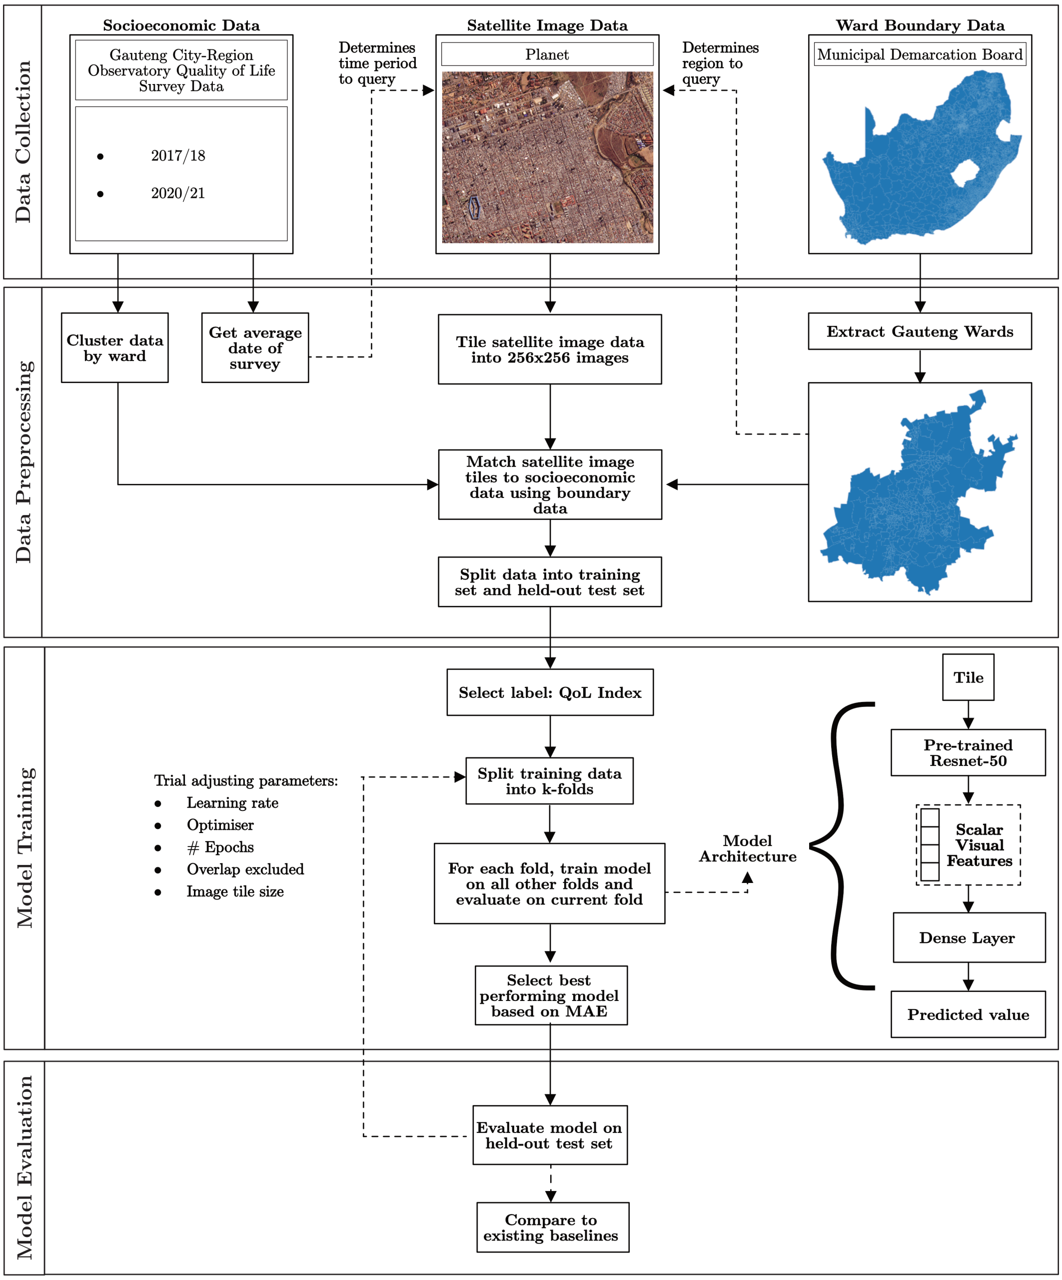

Figure 4.3: An overview of the methods used to predict the GCRO QoL index from satellite images. Satellite Image ©2024 Planet Labs PBC.

4.4. Input Data

Three third-party data sources are relied on for the machine learning methods.

4.4.1 GCRO QoL Survey Data

The GCRO QoL survey data [10] is selected as the ground-truth data source for training and validation of the machine learning pipeline. This data provides a composite GCRO QoL index that summarises 33 variables across different components, including socioeconomic status, services, health, safety, participation, life satisfaction, and government satisfaction [12]. These subcomponents and the final weighted GCRO QoL index are provided in the datasets made available by the GCRO [10]. In the subsequent data processing and training, the pre-calculated GCRO QoL index is used as the selected target variable.

The GCRO makes their data available for public use via the DataFirst platform. To download data from the DataFirst platform, one has to be registered with the platform, and a reason for use must be provided before the dataset can be downloaded. The GCRO QoL survey data is downloaded in DTA format for Round 5 (2017-2018) [82] and Round 6 (2020-2021) [83]. These surveys are selected due to sharing the same GCRO QoL index methodology [10, 12].

4.4.2 Municipal Demarcation Board Ward Boundary Data

The MDB provides public geographic data in shapefile format that includes the different administrative boundaries (ward, municipal, and provincial) across South Africa. In this research, the 2016 demarcations are used, because these were the boundaries in place at the time of the selected 2017/18 and 2020/21 GCRO QoL surveys [86]. The shapefile data can be downloaded directly from the MDB website [86].

4.4.3 Planet: Satellite Image Data

Critical to this research is satellite image data matching the two GCRO QoL surveys. Daytime satellite images of Gauteng are obtained from Planet, a private satellite image data provider, under their Education and Research programme [84]. This programme makes available a limited amount of satellite data per month at no cost to individual students and academics.

Two different satellite image data products from Planet are considered for this research: Planet Monitoring and Planet Basemaps [84]. Planet Monitoring delivers high frequency satellite image data. However, a substantial amount of data processing would be required to obtain cloudless, normalised satellite image data covering the entire region of Gauteng, South Africa during the corresponding time periods of the GCRO data. The Planet Basemaps product offers these requirements out the box with limited additional processing required. Specifically, the Basemaps product makes composited satellite image data available on a monthly, quarterly, and annual basis for the entire planet. This data has been pre-processed by the company to deliver normalised image data with satellite data merged over the relevant period to select the least noisy pixels.

The Basemap Application Programming Interface (API) provides data from PlanetScope/RapidEye satellites. These capture images of the earth with a resolution of between 3 and 4.5 metres per pixel. The images available are divided into a grid, where each grid cell - a Basemap Quad - is 4096 x 4096 pixels. To select relevant Basemap Quads for the two GCRO QoL datasets, additional data processing is performed as described in 4.5.4.

4.5. Data Processing

Once the input data is downloaded, several processing steps are followed to ensure the data is in a suitable format for training the machine learning model. Note that the GCRO QoL and the MDB data do not have any dependencies. However, the Planet data requires spatial and time-based parameters that are obtained from the independent datasets. This is documented below.

4.5.1 Quality of Life Data

Each GCRO QoL survey dataset is processed to obtain a mean interview date. The dataset is ingested with *pandas* [87] and the in-built function for calculating mean is applied to the interview date column. This mean interview date is used to inform the time range in which to select satellite image Basemap Quads from Planet. Given that changes over time are not a primary consideration of this research, a single date period in which to obtain satellite data simplified the design.

4.5.2 Ward Boundary Data

The MDB data is processed independently of the GCRO data. Using *geopandas* [88], the shapefile for the administrative boundaries of South Africa is filtered down to extract Gauteng wards boundaries. Additionally, an available dissolve function is used to extract the Gauteng provincial boundary, which is required to retrieve the correct spatially located satellite image Basemap Quads from Planet.

4.5.3 Clustering Quality of Life Data

Each GCRO QoL survey contains thousands of interview records collected across Gauteng. To use this data for training, the survey data is summarised at the ward level. There are 529 wards in Gauteng and each response in the survey data has a ward number assigned. Using *pandas* [87], the GCRO QoL survey data is grouped by ward and survey year (2017/18 and 2020/21) and the average GCRO QoL index for that ward and survey is obtained using in-built *pandas* functions. The resulting dataset contains 1058 records or clusters.

For each cluster, the corresponding bounding polygon from the processed MDB ward boundary data is appended. In this case, *geopandas* [88] is used to ingest and filter intermediate Gauteng shapefile and to create a new *GeoDataFrame* object that includes the cluster and its boundary.

It should be noted that the aggregation of the GCRO QoL survey data is performed using simple arithmetic means without applying the survey weights provided with the dataset. The GCRO provides weights to ensure representative estimates when aggregating survey responses, as the sampling design employs stratification and clustering that requires weighting adjustments for unbiased population estimates. Excluding the weights during clustering represents a methodological limitation of the current analysis that may have introduced bias in the ward-level quality of life estimates used as target variables for the machine learning model.

4.5.4 Downloading Correct Planet Data

With the processing the survey and MDB data complete, the mean interview date and Gauteng spatial boundaries are used to select the correct Basemap Quads from Planet. For each of the two GCRO QoL surveys, a query is run via the Planet

Basemap web user interface to find Basemap Quads within or overlapping with the Gauteng boundary. Additionally, the Basemap Quad query includes a filter for the month of the average date of the corresponding survey. The identifiers of the resultant Basemap Quads are saved and stored in the codebase, and the Planet API is used to download the exact Basemap Quad satellite images in GeoTIFF format via a series of HTTP requests.

4.5.5 Tiling Satellite Image Data

The satellite image Basemap Quads are further divided into tiles with a crop size of 256 x 256 pixels for processing during training, validation, and testing. The tile size is set by a parameter in the pipeline and therefore, different tile sizes can be experimented with.

To tile the GeoTIFF images, the *windows* module from *rasterio* [89] is used. A grid with cells of dimensions x x x pixels - where x is the parameterised tile size - is used to iterate across the full Basemap Quad. The value of x must be a divisor of 4096 - the maximum width of the square Basemap Quad. For each grid cell in the grid, the corresponding portion of the Basemap Quad corresponding is saved as a standalone image tile.

4.5.6 Matching Clustered Socioeconomic Data to Satellite Image Data

The final data processing requires satellite image tiles to be matched to their corresponding ward data. A spatial join is performed using *geopandas* [88] on the physical coordinates of the tile bounding box and the clustered GCRO QoL data. The final dataset contains records with the following features:

- A satellite image tile file name,
- The ward code to which the tile belongs,
- The GCRO QoL index for the tile based on the ward, and
- The survey year matching the tile and index.

Importantly, this dataset contains duplicate tiles because there are tiles that overlap ward boundaries and therefore belong to two wards.

4.5.7 Training and Validation Data Split

In preparation for training, the matched dataset of satellite image tiles with corresponding GCRO QoL values is divided into five folds or subsets. This is to allow for k-fold cross-validation, which is a useful technique for data limited machine learning pipelines. The dataset is partitioned into k equal subsets, where each subset serves as a testing set once while the remaining $k-1$ subsets form the training data, with the final performance being the average across all k iterations. Given that there are only 529 unique wards, this is a suitable approach to take.

Additionally, the folds are grouped spatially so that satellite image tiles from the same ward do not appear in different folds. To address the duplication challenge described previously, overlapping tiles are dropped. These two steps are intended to mitigate spatial autocorrelation - the tendency for nearby locations to have similar characteristics. This is a key challenge in satellite image machine learning tasks where spatial information can cause biases in the model if there is "leakage" between training, validation, and test sets [90]. In other words, if training and test data contain spatially adjacent areas, the model may learn location-specific patterns rather than generalisable features, leading to overly optimistic performance estimates.

4.6. Machine Learning Model

The overarching machine learning approach relies on two steps: feature extraction from the satellite images and a regression model to predict the GCRO QoL index. This two-stage approach is common in remote sensing applications where visual features must be transformed into meaningful representations before being used for prediction tasks. Refer to Figure 4.3 for a high-level overview of the model.

4.6.1 Feature Extraction

A CNN is selected as the base model for extracting features from satellite image tiles. Specifically, a ResNet-50 architecture is selected with pre-trained ImageNet Weights. ResNet-50 is chosen for several reasons. First, its Deep Residual Architecture allows for effective learning of complex visual patterns in satellite images while mitigating the vanishing gradient problem through skip connections [41, 56]. Second, based on prior literature, pre-training on ImageNet provides a strong initialisation for transfer learning in poverty mapping, even though satellite imagery differs from natural images [27, 51, 56].

Typically, the ResNet architecture is used in classification tasks. Given that this research is a regression task, the top layer (fully connected classification layer) is removed so that the output of the feature extraction step is a 2048-dimensional feature vector representing the visual features identified by the neural network from the input image.

The base model is fine-tuned using the GCRO QoL training data as target labels for each input tile. This allows the model to adapt to the specific characteristics of the local satellite images while retaining the general feature detection capabilities learned from ImageNet.

4.6.2 Regression

After extracting features from satellite images, a standard approach is to train a regression model to predict the target variable from the extracted features. This approach allows for more complex relationships to be captured between extracted features and the target variable.

In the case of the conducted research, the decision is taken to attempt to regression with the CNN architecture alone. A fully connected Dense Layer with a Sigmoid Activation is selected.

4.7. Training

Compilation and training of the machine learning model relies on *Tensorflow* [91]. The ResNet-50 model is loaded with pre-trained ImageNet Weights. k-fold cross-validation is then used to fine-tune the model using each fold's training set with input satellite image tiles and corresponding GCRO QoL target values. The following configuration is used:

- Batch Size: 32.
- Optimiser: Adam Optimizer with an initial learning rate of 1e-3.
- Early Stopping: Training terminates if validation loss does not improve for

ten consecutive epochs.

- Data Augmentation: Random horizontal and vertical flips to reduce overfitting risks.

Each fold's training process is monitored using the corresponding validation set to prevent overfitting, with MSE as the loss used to evaluate each training epoch. Once the full k-fold cross-validation has run, the best performing model is selected for final evaluation based on MSE.

4.8. Final Model Evaluation

The final model is evaluated using r^2. This indicates the proportion of variance in the predicted GCRO QoL index explained by the model. This is visualised by plotting actual GCRO QoL index value against their predicted value.

The model is evaluated based on the corresponding fold's hold-out test set. An analysis of all test tiles is conducted as well as a comparison of ward-level predictions which rely on averaging all tile predictions in a ward.

Additionally, qualitative evaluation is conducted by visualising activation maps from the ResNet-50 model to understand the aspects of the satellite images (e.g., building density, vegetation, road networks) the model "focuses" on when making predictions. This provides an initial look at the interpretability of the model and the relationship between visual features and socioeconomic indicators.

4.9. System Environment and Tools

This section documents the specific hardware and software specifications used to process data and train the deep learning model. Additional details are available in the Github repository [76].

4.9.1 Hardware Specifications

Model training is performed on a high-performance computing cluster equipped with a NVIDIA GeForce RTX 4090 GPU with 128 GB of VRAM. At least 100 GB of storage space is recommended.

4.9.2 Software Specifications

The machine learning pipeline is implemented using the following key software libraries:

- Python 3.10 [92]
- TensorFlow 2.18 [91]
- Rasterio 1.3 [89]
- Pandas 2.0 [87]
- GeoPandas 0.14 [88]
- Matplotlib 3.8 [93]

4.9.3 Development and Deployment Tools

The research implementation employs several tools to support reproducibility and code tracking:

- Git for version control of code and documentation
- GitHub for repository hosting and collaboration
- Data Version Control (DVC) for tracking large datasets and model artifacts [94]
- Poetry for dependency management and creating isolated virtual environments [95]
- Jupyter Notebooks for exploratory data analysis and result visualisation
- ArcGIS for spatial data exploration and analysis

4.10. Chapter Summary

This chapter has presented the applied methods for predicting the GCRO QoL index from satellite images in Gauteng, South Africa. The selected approach tests feature extraction and regression of the target index using only satellite images in the visual spectrum and a CNN. The results of these methods are documented in the next chapter.

5 Results

This chapter presents the results of the methods applied as documented in Chapter 4. The results are organised into three key sections: understanding the input data, intermediate outputs of data processing, and the outcomes of the machine learning models used to measure the GCRO QoL index in Gauteng, South Africa.

5.1. Understanding the Input Data

Samples of input data are presented below. The full datasets are available for download from the relevant data sources.

5.1.1 Planet Satellite Basemap Quads

Figure 5.1 shows an example of the Basemap Quads available from Planet. These Basemap Quads are 4096 x 4096 pixels. The displayed Basemap Quad is from April 2018, the average month for the GCRO QoL survey Round 5 (2017/18).

5.1.2 MDB Ward Boundary Data

Figure 5.2 displays the MDB data for 2016. This provides ward boundaries for the entire area of South Africa.

5.1.3 Quality of Life Data

Figure 5.3 illustrates the different data source choices available in a South African context with their corresponding sample sizes specifically for Gauteng. The 2016 USAID DHS survey has the lowest sample size, at only approximately 2500 records. The General Household Surveys (GHS) are a series of surveys conducted by Stats SA, and also have relatively low sample sizes. In comparison, the GCRO QoL Surveys have much higher sample sizes with both 2018 and 2021 including over 10,000 records.

Figure 5.1: A sample Basemap Quad from Planet in April 2018 - the average month for the GCRO QoL survey Round 5 (2017/18) [84]. Image ©2024 Planet Labs PBC.

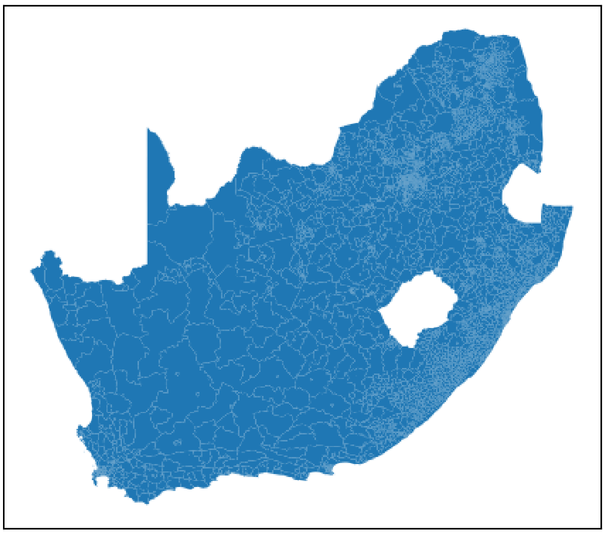

Figure 5.2: Map of raw South African MDB data for 2016

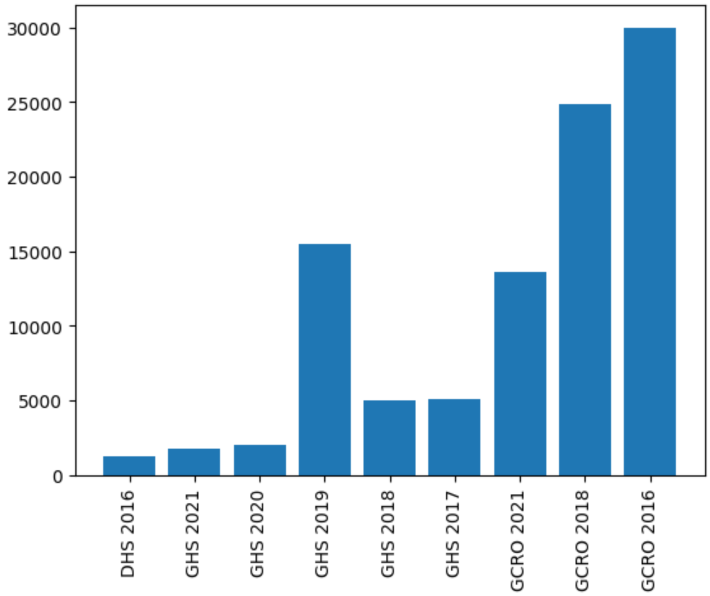

Figure 5.3: Comparison of sample sizes in Gauteng for different surveys over different years: DHS , GHS, and GCRO QoL

Statistical summaries of the 2018 and 2021 GCRO QoL surveys are presented in Tables 5.1 and 5.2.

Table 5.1: Statistical Summary of 2017/18 GCRO QoL Index Data

Statistic	Interview Date	QoL Index
Count	24,889	24,889
Mean	2018-04-10	63.797958
Min	2017-10-31	11.7711
25%	2018-02-07	56.9283
50%	2018-05-03	64.7718
75%	2018-06-16	71.9428
Max	2018-09-07	96.1394
Std Dev	N/A	12.0369

5.2. Intermediate Outputs of Data Processing

Intermediate outputs are obtained from processing the input data, as described in Chapter 4.5.

Table 5.2: Statistical Summary of 2020/21 GCRO QoL Index Data

Statistic	Interview Date	QoL Index
Count	13,616	13,616
Mean	2021-01-12	61.3763
Min	2020-10-06	15.1310
25%	2020-11-15	54.3398
50%	2021-01-16	62.2891
75%	2021-03-01	69.7130
Max	2021-05-27	94.1068
Std Dev	N/A	12.0167

5.2.1 Satellite Image Tiles

Satellite image tiles are obtained by processing the Basemap Quads into smaller chunks of 256 x 256 pixels. A sample tile is displayed in Figure 5.4. This is one of the outputs of processing the previously shown Basemap Quad in Figure 5.1.

Figure 5.4: A sample tile extracted from a downloaded Planet Basemap Quad from April 2018. Satellite Image ©2024 Planet Labs PBC

5.2.2 Extracted Gauteng Ward and Provincial Boundaries

The output of processing the MDB shapefile to extract all wards in Gauteng is displayed in Figure 5.5a. The outline of the province in Figure 5.5b is important for identifying Basemap Quads to download from Planet.

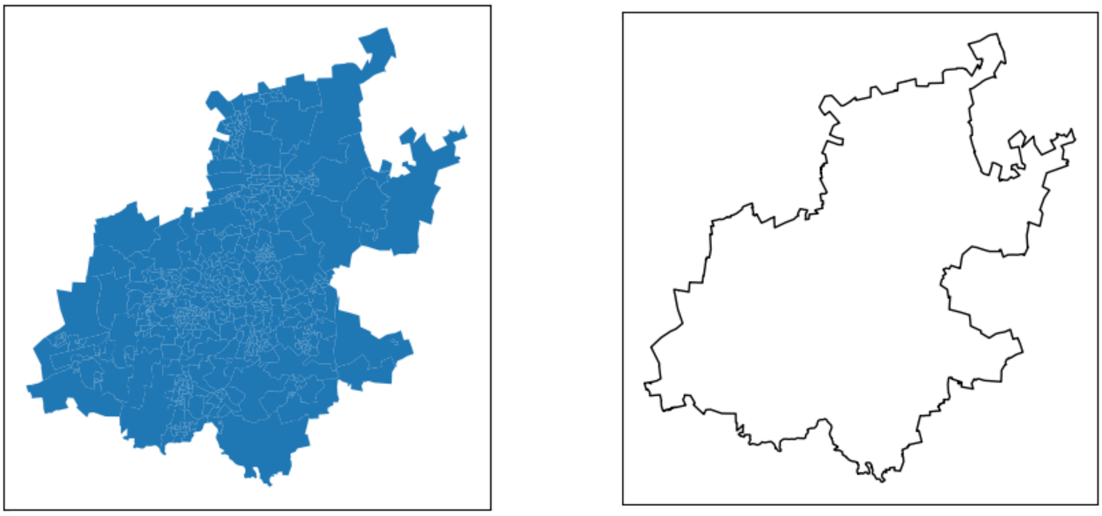

(a) Gauteng wards **(b) Gauteng boundary**

Figure 5.5: Maps of processed South African MDB data for 2016

5.2.3 Matched spatial and survey data

The outputs of matching the clustered GCRO QoL index data to the corresponding wards are displayed in Figure 5.6.

As described in Chapter 4.5, satellite image tiles that overlap multiple wards are discarded to avoid spatial auto-correlation errors during training. This results in the final mapped areas as illustrated in Figure 5.7.

5.3. Machine Learning Model Outputs

The machine learning model outputs are displayed in this section. Results are selected based on relevant discussion points in the following chapter.

5.3.1 Actual vs Predicted Values Test Set

The standalone test set is used to evaluate the final selected model. The test set is not used in the training or validation of the model.

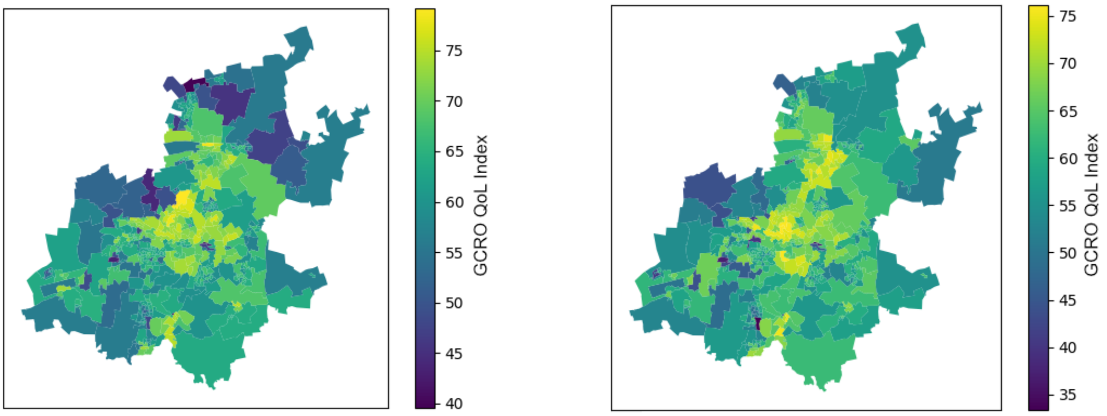

(a) Round 5 (2017/18) (b) Round 6 (2020/21)

Figure 5.6: Chloropleth maps of the 2017/18 and 2020/21 GCRO QoL Index aggregated at the ward level

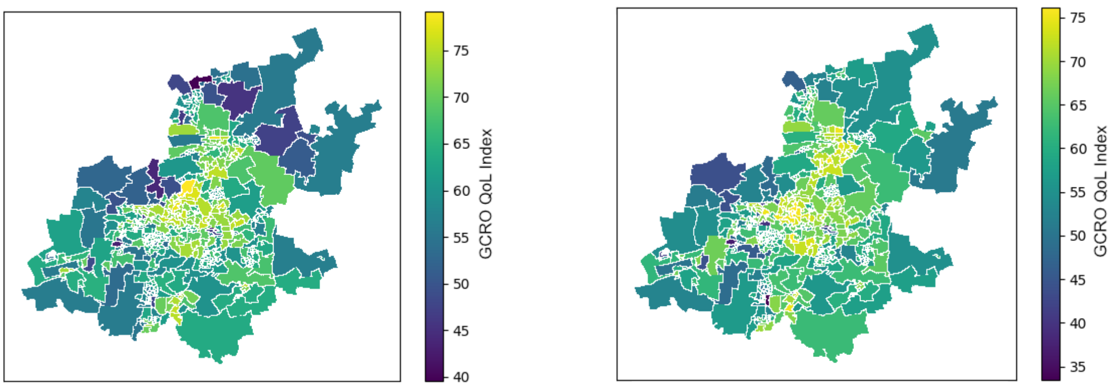

(a) Round 5 (2017/18) (b) Round 6 (2020/21)

Figure 5.7: Chloropleth maps of the 2017/18 and 2020/21 GCRO QoL Index aggregated at the ward level with tiles overlapping ward boundaries removed from the dataset

When grouped by date and ward (see Figure 5.8), r^2 for this result is negative indicating a model that performs worse than predicting the sample mean.

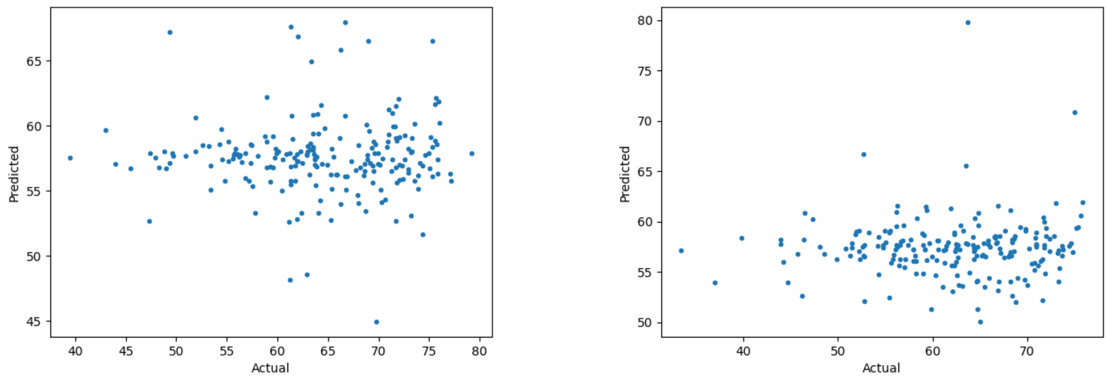

(a) Actual vs Predicted GCRO QoL Index values for all test wards in 2018

(b) Actual vs Predicted GCRO QoL Index values for all test wards in 2021

Figure 5.8: Actual vs predicted GCRO QoL Index values for 2018 and 2021 data grouped by ward

Figure 5.9 presents maps comparing actual and predicted GCRO QoL Index values for wards.

5.3.2 Activation Maps of Neural Network

The output of the different layers of the neural network is seen in Figure 5.10. The selected sample satellite tiles are from the test set and illustrate how features of different urban contexts are extracted at different convolutional layers.

5.4. Summary

This chapter has presented the results obtained from different stages in the machine learning pipeline. These results will be interpreted in the following chapter to draw a conclusion to the research question posed.

(a) Maps of actual and predicted GCRO QoL Index values for all wards in 2018

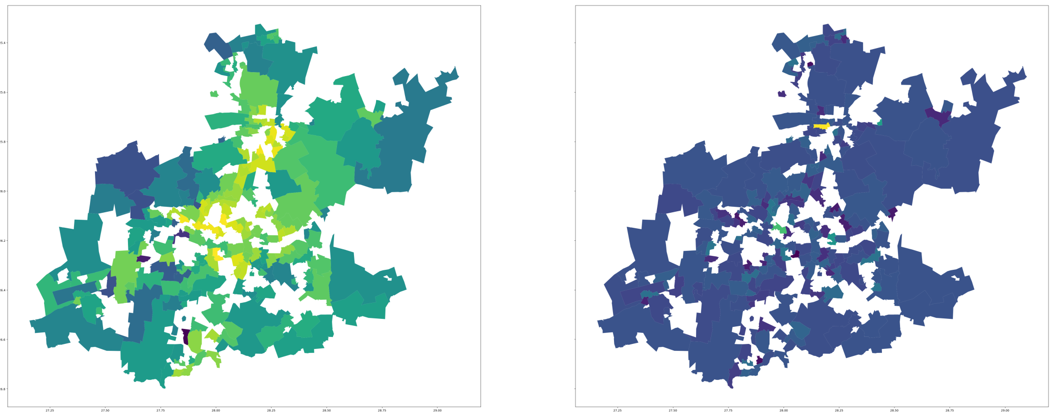

(b) Maps of actual and predicted GCRO QoL Index values for all wards in 2021

Figure 5.9: Maps of actual and predicted GCRO QoL Index values for 2018 and 2021

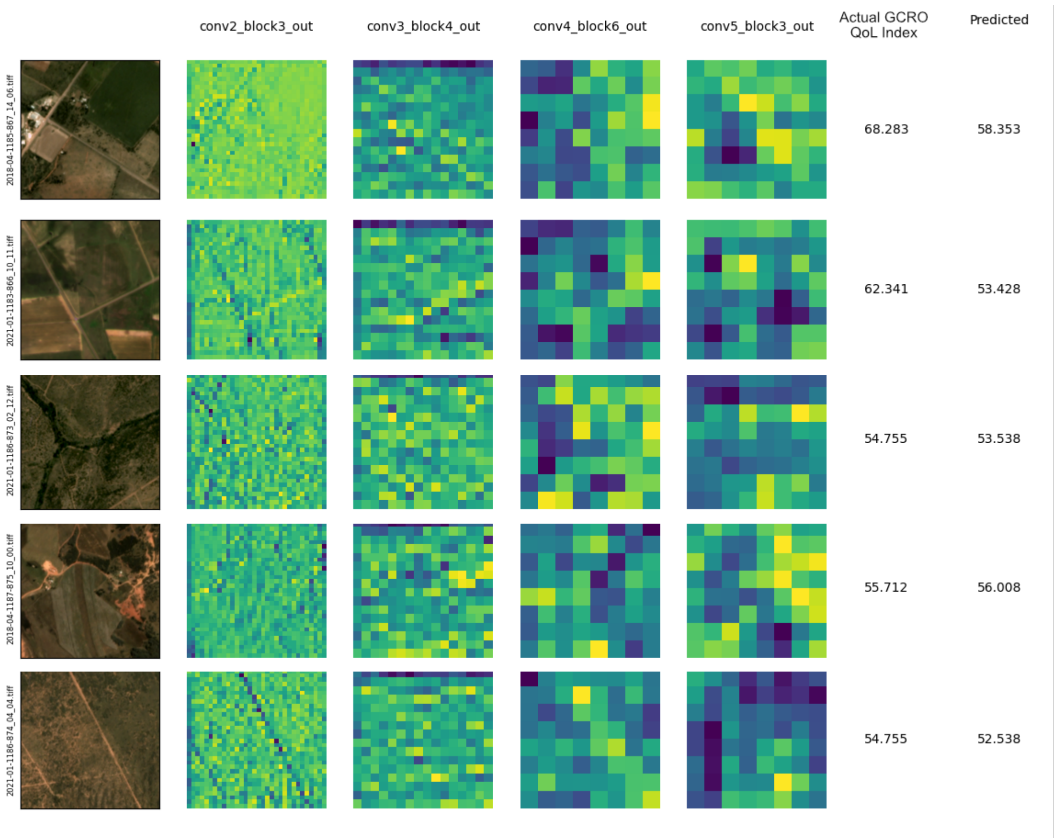

Figure 5.10: Activation Maps of selected model

6 Discussion

Envision a future in which data about the quality of people's lives, in cities and rural communities alike, is continuously updated and used in decision-making to meaningfully improve livelihoods. Making this future a reality requires frequent, reliable, accurate, low-cost multidimensional quality of life data that could be synthesised from readily available big data sources, like satellite images.

To contribute to this vision, the conducted research asks: "how well can a machine learning model estimate a quality of life index from satellite images in Gauteng, South Africa?" Specifically, the GCRO QoL index in Gauteng, South Africa is predicted from satellite images using the data processing and machine learning methods described in Chapter 4. Based on the results obtained, the research question can be answered as follows: the machine learning model cannot reliably estimate the GCRO Quality of Life Index/Survey/Dataset (GCRO QoL) index from satellite images in Gauteng, South Africa. With negative r^2 values, the model performed worse than baseline predictions, demonstrating no meaningful correlation between satellite imagery features and the selected multidimensional quality of life index.

The hypothesis of the research is that there would be, at minimum, a baseline correlation between visual satellite image features and the selected GCRO QoL index. A null hypothesis is concluded. The causes of this finding appear to stem from three key factors, and likely, the interaction between them. First, the complexity of quality of life as a multidimensional measure encompassing both quantitative and subjective aspects of human wellbeing proves difficult to capture through visual satellite data alone. Second, the designed model extracts features from daytime satellite images without supervision, which proves to be ineffective. Finally, the findings suggest that methods developed primarily for poverty mapping, do not transfer well to multidimensional quality of life measurement in South Africa's unique socioeconomic context, especially in a highly unequal urban landscape such as Gauteng.

The remainder of this chapter delves into each of these themes, with a discussion and analysis of the research methods and results. The impact of these findings on the research outcomes is reviewed and opportunities for future work are proposed.

6.1. Summary of Results

Before unpacking the null hypothesis conclusion, this section summarises the results and provides an interpretation of what is observed.

In Figure 5.8, scatter plots of target and predicted values for the 2018 and 2021 GCRO QoL index are shown. In both cases, there is no definitive correlation between predicted and target values and a negative r^2 is reported. As seen in Figure 5.9, the model's predictions are fairly homogeneous across the map, differing from the variations seen in the actual GCRO QoL index values per year.

Similarly, in the mapping of this same data, one observes that the predicted values are relatively uniform compared to the input data. This suggests that no meaningful patterns are discovered in the satellite image data with respect to the selected GCRO QoL index. Further discussion on why this might be the case is provided below.

6.2. Complexity of Quality of Life as a Multidimensional Measure

Quality of life is a complex, multidimensional concept that is challenging to measure and quantify in ordinary socioeconomic development studies. The GCRO QoL index incorporates various objective and subjective factors that influence quality of life. Despite this complexity, the assumption of this research was that it would be possible to predict the index, at least to some extent, from satellite images. In the case of the methods applied, this assumption has proven to not hold true.

Based on analysis of the results and similar work in the field, a primary conclusion is that quality of life's subjective variables significantly affect the model's performance. The expectation was that certain tangible features like roofing or dwelling type would counter some subjective measures in the multidimensional, GCRO QoL index. Further investigation is required to better understand, which specific variables may reduce performance and if any of the objective measures of the GCRO QoL index can be estimated reliably from satellite images using the same machine learning methods.

6.2.1 Decision to Use GCRO QoL Index

The choice to predict the composite GCRO QoL index, rather than individual variables or dimensions, reflects different considerations and decision points during the course of this research. A brief reflection on the reasoning for selecting this dataset is presented below.

At a high-level, the GCRO QoL surveys has higher sample sizes. Figure 5.3 illustrates a comparison in survey counts between different survey sets and different years. The DHS survey is of interest, given its wide use in the domain of measuring development indicators from satellite images. However, in a South African context, the last survey was released in 2016 and had less than 5000 household records. In comparison, the 2020/21 and 2017/18 GCRO QoL surveys had over 10,000 and 20,000 survey responses respectively in a much smaller spatial area 5.3. Through this lens, the GCRO QoL datasets are a good choice given the increased amount of data. That being said, when the GCRO data is clustered, it only has 529 records. This is a trade-off where locally relevant data is prioritised. Additionally, a primary intention of this research is to understand how existing methods apply to a local context. For this reason, the GCRO QoL index is well-suited for the research at hand as well, in spite of also being a good practical choice.

Why not individual GCRO QoL variables?

The GCRO dataset includes individual variables across seven dimensions that feed into the GCRO QoL composite index. Using the available individual dimensions as target variables might have yielded better predictive performance, particularly for objective measures like infrastructure access that are expected to correlate more obviously with visual features. However, the composite index was intentionally chosen for several reasons.

First, the GCRO QoL index aligns with the research goal of understanding satellite imagery's potential for measuring locally relevant, multidimensional quality of life. Testing this boundary seemed more valuable than optimising performance on individual dimensions that might not reflect the complexity of quality of life measurement in practice.

Second, the GCRO QoL index represents a sophisticated, locally-derived measure. Unlike international indices that impose predetermined frameworks developed in the Global North [12], this index emerges from factor analysis of South African data, allowing local conditions to shape its construction. The resulting seven-factor model incorporates both objective indicators (housing, services, employment) and subjective measures (life satisfaction, government satisfaction) tailored to the Gauteng context.

Third, the composite index provides a more rigorous test of satellite imagery's capabilities. Individual objective dimensions might correlate predictably with visual features, but a validated multidimensional index incorporating subjective wellbeing represents the kind of complex measure most valuable for policy applications. If satellite imagery could estimate such an index, it would demonstrate genuine potential for complementing traditional survey methods.

Finally, practical constraints influenced this approach. The research was designed to start with the GCRO QoL index and then further explore individual dimensions if results indicated poor performance. This would be valuable in identifying specific variables impacting model performance. However, time constraints prevented extending the investigation to individual factors, creating an opportunity for future work.

6.3. Ineffectiveness of Unsupervised Feature Extraction from Daytime Satellite Images

The second core theme believed to underpin the poor performance of the machine learning model is the selected feature extraction techniques applied, along with how these features are used. Specifically, a decision is taken to apply an unsupervised learning approach for extracting features from satellite images, and the finding is that these features are not sufficiently correlated to quality of life alone.

6.3.1 Possible Flaw?

Although quality of life is a complex metric, its complexity alone cannot explain why there is such limited correlation between satellite image inputs and outputs. Evidence supporting this is observed in parallel experiments where other sub-indexes from the GCRO QoL data that contribute to quality of life are also not correlated to input satellite images. Specifically, economic indicators like wealth or income, which are included in the GCRO surveys, have no improved performance when trying to extract these products from satellite images.

A possible flaw in the machine learning approach is that multiple image tiles are used to train for the same ward and the same GCRO QoL index. This means that in the case that a single ward has high visual feature disparities (e.g. a golf course next to a township), significant noise is likely introduced into the model training. The feature representations are vastly different and yet, the GCRO QoL index is the same.

To resolve this, future work should consider selecting only one tile per ward or a more complex regression model could be considered as described in the next subsection. Additionally, there could be value in further investigation into whether model performance varies across different regions or ward types, particularly in more socioeconomically homogeneous areas.

6.3.2 Extracted Features used for Regression

Apart from the feature extraction method, the approach for correlating these features with the GCRO QoL index is constrained. By using a Dense Layer as the regression step, more complicated and nuanced patterns in the data are likely undetectable. A prime opportunity for future work is exploring alternative regression models that can be used on the extracted features. In the case of multiple satellite tiles per ward, features could be extracted for all ward tiles and then this multidimensional feature representation of a ward could be used in a regression model (see Figure 6.1 illustrating the proposed solution for such a regression). Key in this approach would be considering factors of dimensionality reduction, while still resolving the initial questions of the appropriateness of the applied feature extraction methods.

6.3.3 Interpretability

When it comes to feature extraction from satellite images using an unsupervised algorithm, interpretability is a major limitation. Future work would require investigating and prioritising interpretability over functionality. In other words, designing a model that is understandable so that regardless of whether it performs, there are clear human-interpretable reasons why. In the case of this dissertation, the reasons for the model not performing are difficult to assess given the type of model that is applied.

6.3.4 Strengths of Unsupervised Feature Extraction

Despite the critical review of unsupervised approaches to feature extraction, it is worth noting two strengths. First, the minimal requirement for annotation of data is a primary benefit in resource constrained environments. Second, the potential to uncover non-obvious visual patterns that humans might overlook is an interesting and valuable line of research. Although these strengths did not lend themselves to the conducted research, the question remains - what can be measured from satellite images when it comes to quality of life? Unsupervised feature extraction is a valuable tool for exploring this particular question.

6.4. Lack of Transferability of Poverty Mapping Methods to South Africa's Context

The final core theme that emerges, based on the analysis of the results as well as prior literature, is that machine learning models used for mapping poverty do not necessarily transfer well to the South African context. One explanation is that the existing methods have been shown to perform more poorly in indistinguishable urbanised contexts [96]. Instead, poverty mapping methods, similar to the one applied in this research, perform better at differentiating wealth between urban and rural areas.

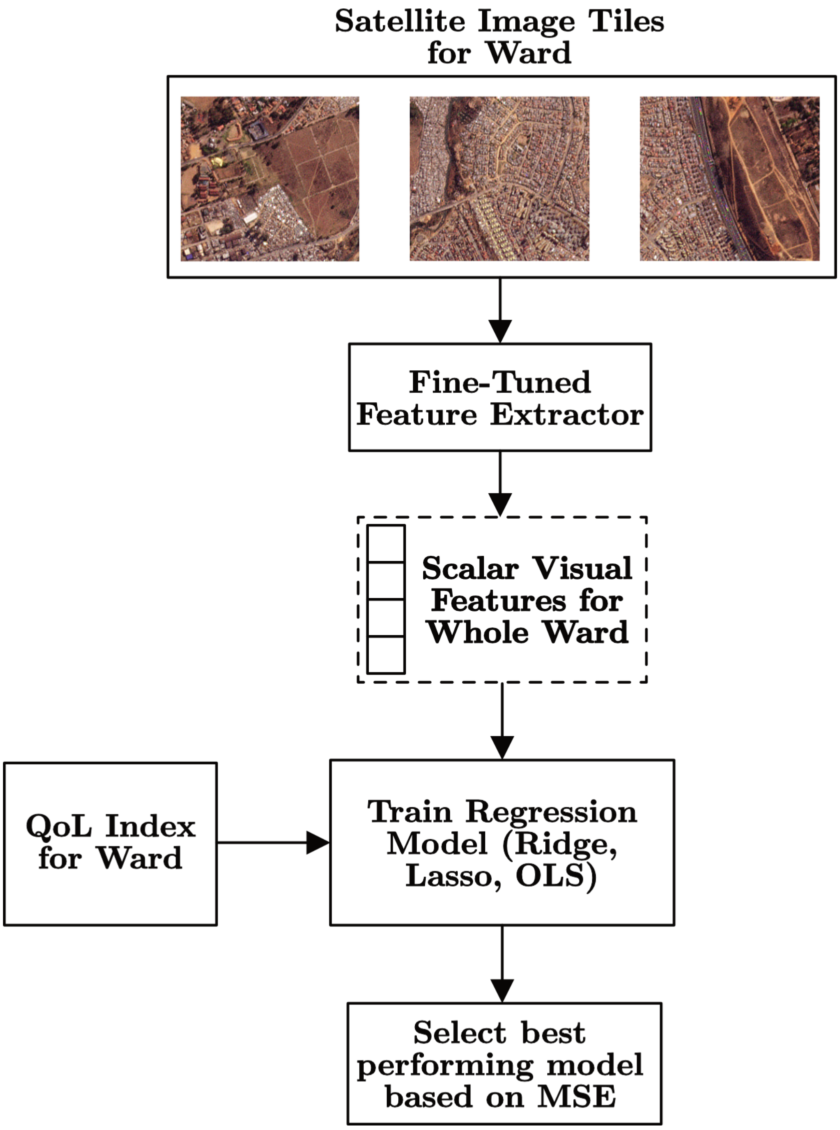

Figure 6.1: A proposed solution to training a regression on extracted features for all satellite image tiles of a ward to predict the ward GCRO QoL index

In Gauteng, South Africa, the density of urban landscape is relatively uniform, while at the same time, the region faces high inequality between different population groups. A logical conclusion is that the standard methods for poverty mapping from satellite images may not be well suited to highly unequal areas like Gauteng. It follows that if wealth is highly unequal and quality of life is highly subjective, that perhaps these methods would more generally not perform well in highly unequal contexts.

There is broad scope to explore this further and to investigate the suitability, transferability, and interpretability of these machine learning methods for complex socioeconomic development measurement.

6.5. Critical Analysis of Methods

This section discusses some of the technical decision points and their implications on the findings. Opportunities for future work are concluded from this critical analysis as well as the discussion above.

6.5.1 Applied Methods

A key assumption was that existing methods would be suitable for this particular problem if visual features and quality of life are somewhat correlated. This assumption did not hold. Future work should seek to design novel machine learning models that solve the particular challenge of measuring quality of life from satellite images. Bearing in mind that this could limit the work's relevance in other contexts, a balance must be struck between a generalised model for estimating particular indexes, while catering for requirements of fine-tuning for unique local factors.

6.5.2 Data Decisions

Two key data decisions influence results:

1. Using GCRO QoL data with inconsistent ward sizes, chosen for its local relevance despite potential statistical challenges.

2. Selecting Planet satellite imagery for its high resolution, though future work might benefit from using free public datasets like Landsat or Sentinel.

There are plenty of opportunities to experiment with alternative datasets or to incorporate multimodal data. However, the recommendation is to only action this once a baseline model exists for measuring quality of life from a single data source.

6.5.3 Pre-trained Weights

Different pre-trained weights are not compared in this study. Future work should explore alternatives to ImageNet Weights, particularly domain-specific weights like xView that are designed for satellite imagery applications [97].

6.5.4 Survey Weighting Oversight

A methodological limitation of this research was the failure to apply survey weights when aggregating individual GCRO QoL responses to the ward level. The GCRO provides weights designed to account for the complex sampling design (stratification, clustering, and probability proportional to size selection) and ensure representative population estimates [10]. By using simple arithmetic means rather than weighted averages, the ward-level quality of life estimates used as target variables may be biased and not representative of the true ward populations. This oversight could significantly impact model performance and the validity of results, as the fundamental target variable used throughout the analysis may not accurately reflect ward-level quality of life. Future work should prioritise proper application of survey weights, requiring recalculation of all ward-level aggregates and rerunning the analysis pipeline.

6.6. Broader Opportunities for Future Work

Despite the limitations encountered, measuring quality of life with satellite imagery remains a compelling pursuit, particularly in a South African context. Building on the discussion above, this section outlines key opportunities for future research.

6.6.1 Refining Approaches to Quality of Life Measurement

Quality of life measurement remains fundamental to effective urban planning and policy development in South Africa. Understanding multidimensional wellbeing is critical not only for building infrastructure correctly but also ensuring that development meets the diverse needs of communities across different regions. While traditional survey-based methods capture valuable information, complementary approaches using big data sources like satellite imagery could provide more frequent, cost-effective monitoring capabilities to support informed decision-making.

Future work should explore alternative approaches and datasets that might yield more successful results. This includes investigating different spatial resolutions, temporal frequencies, and combinations of satellite and additional data sources that may better capture the nuances of quality of life indicators. Additionally, exploring how satellite data might supplement rather than replace survey data could lead to more robust hybrid measurement systems.

6.6.2 Enhancing Model Interpretability and Understanding Limitations

Future work should prioritise interpretability in machine learning approaches to better understand what specific aspects of satellite images relate (or fail to relate) to quality of life dimensions. Enhanced interpretability would help identify precisely what visual features are being captured by machine learning models.

This improved understanding would help establish clearer boundaries regarding what aspects of quality of life can and cannot be measured from visual features in satellite imagery. Such knowledge would be invaluable for guiding future research directions and setting realistic expectations for satellite-based quality of life assessment methods, especially in real-world environments.

6.6.3 Developing Independent Models of Quality of Life from Satellite Images

A significant opportunity exists in shifting from "modeling the model" to developing independent frameworks for understanding quality of life directly from satellite images. As Alfred Korzybski coined, "the map is not the territory" - survey data itself is an abstraction of lived experience, and attempting to model this abstraction introduces compounded complexity and potential error.

Rather than trying to replicate survey results, future research could focus on what unique quality of life indicators can be directly extracted from visual data. This approach suggests creating new abstract representations of quality of life that reflect reality, potentially yielding novel insights into human development patterns that surveys might miss.

6.7. Chapter Summary

By pursuing these three interconnected research directions, the field can advance beyond current limitations while maintaining focus on the ultimate goal: developing reliable, accessible methods for understanding and improving quality of life across diverse communities. The next chapter will reflect on the conducted work, the findings, the opportunities for future work, and conclude the dissertation.

7 Conclusion

This dissertation has described the research conducted to answer the question: "how well can a machine learning model estimate a quality of life index from satellite images in Gauteng, South Africa?" The conclusion is that the applied methods are not well-suited to the task of measuring the GCRO QoL index from daytime satellite images.

The key findings are as follows:

1. A multidimensional, composite quality of life index, when aggregated to the ward level, does not have trivial features correlated to visual spectrum satellite images.

2. Unsupervised feature extraction from daytime satellite images is not an appropriate method for this particular task.

3. Adapting previous methods used for mapping poverty from satellite images, do not transfer well to estimating multidimensional quality of life in a local South African context. In particular, this is perhaps because of the highly unequal socioeconomic landscape within the urban environment of Gauteng.

These findings suggest that there is more than meets the eye when measuring quality of life from satellite images. Just as humans may face challenges quantifying on-the-ground quality of life from simple aerial observations, so does the selected machine learning approach presented in this dissertation. Further exploration into the boundaries of this and other machine learning algorithms for measuring human development would be valuable.

7.1. Identified Limitations

Several limitations affected the outcomes of this research:

Methodological constraints: The selected feature extraction and regression methods are not optimal to measure the specific characteristics of the GCRO QoL index, especially at the ward-level, where fine-grained spatial linkages in the GCRO QoL data are unavailable.

Conceptual challenges: The GCRO QoL index incorporates subjective dimensions that may not manifest in physical features visible from above. Components such as social relationships, governance satisfaction, and subjective wellbeing have limited physical expressions in satellite imagery.

Contextual factors: Gauteng's historical development patterns, characterised by extreme spatial inequality from apartheid planning, create unique challenges for satellite-based analysis. Areas with similar physical appearances may have vastly different quality of life profiles due to socioeconomic or historical factors invisible from above.

Survey weighting: The analysis failed to apply the survey weights provided with the GCRO QoL data when aggregating responses to ward level, potentially introducing bias in the target variables used for model training and evaluation.

7.2. Opportunities for Future Work

This research opens several promising avenues for future investigation:

Domain-specific feature engineering: Developing feature extraction techniques specifically designed for quality of life and the South African urban context, possibly incorporating relevant data about settlement patterns and infrastructure.

Component-specific models: Rather than predicting a composite quality of life index, future work could focus on specific dimensions (e.g. housing quality, service delivery) that may have stronger correlations to features in satellite image data.

Temporal analysis: Examining changes in satellite imagery over time, aligned with quality of life surveys, could reveal relationships between infrastructure development and changing wellbeing. This approach might better capture quality of life than static analysis.

Interpretable techniques: Employing methods that provide greater interpretability could help identify which features contribute most to predictions, potentially revealing unexpected relationships between visual features and quality of life components.

Multimodal approaches: Combining satellite imagery with other data sources such as street-level imagery, mobile phone data, or social media activity could provide a more comprehensive view of quality of life. These complementary data streams might capture dimensions invisible from a satellite perspective alone.

Developing independent models of quality of life: Exploring opportunities to create novel abstractions of quality of life derived directly from big data sources, rather than attempting to replicate existing survey-based indices. Such approaches could establish alternative frameworks for understanding wellbeing that reveal new dimensions of quality of life not captured by traditional survey methods. This paradigm shift would move beyond using machine learning to approximate conventional measures and instead pioneer distinct measurement techniques that operate independently of survey-based historical data.

7.3. Significance of Work

Despite the limitations encountered, this research makes several contributions to the field.

First, it provides a methodological exploration of the boundaries of satellite image analysis for quality of life measurement in a South African context. While previous studies have demonstrated success in estimating poverty measures from satellite images in various regions [55, 56], this work highlights the challenges of applying similar techniques to measure multidimensional quality of life in Gauteng's unique urban landscape.

Second, this research contributes to understanding the relationship between physical features visible in satellite imagery and complex socioeconomic constructs. The findings suggest that quality of life, as a multidimensional index incorporating subjective wellbeing components, may not be as directly observable in physical infrastructure as simpler economic indicators.

Third, this work establishes a foundation for future research by identifying specific challenges and limitations in the current approach, thereby informing more effective methodological designs for subsequent studies in this domain.

7.4. Final Remarks: Continuing Interdisciplinary Quality of Life Research

This research underscores the inherent complexity of measuring quality of life. While machine learning approaches offer promising tools for complementing traditional survey methods, this study demonstrates that they are not yet capable of replacing ground-level data collection, particularly for multidimensional quality of life.

The value of survey enumerators and direct community engagement remains evident, capturing nuanced aspects of quality of life that remain invisible to sophisticated remote sensing technologies. Rather than seeking to replace these methods, future research should aim to develop complementary approaches that leverage the strengths of both human-centered data collection and computational analysis.

The boundaries of machine learning applications in development measurement continue to evolve. This research contributes to understanding these boundaries, helping to define where current methods succeed and where they require further refinement. By identifying limitations, this work helps guide the development of more effective approaches to measuring and understanding multidimensional quality of life.

In conclusion, while this study did not achieve its original objective of estimating quality of life from satellite imagery alone, it has advanced understanding of the challenges involved and laid groundwork for more sophisticated interdisciplinary approaches. The future of quality of life research likely lies not in replacing traditional methods with artificial intelligence, but in thoughtful integration of different methods to develop richer, more comprehensive measurements of human wellbeing.

References

[1] M. Nussbaum and A. Sen, *The Quality of Life*. Clarendon Press, 11th Mar. 1993, 466 pp., ISBN: 978-0-19-152136-2.

[2] F. Stewart, 'Capabilities and human development,' *UNDP (United Nations Development Programme)*, 2013.

[3] OECD, *How's Life?: Measuring Well-being*. Paris: OECD Publishing, 2011. DOI: 10.1787/9789264121164-en. [Online]. Available: https://www.oecd-ilibrary.org/content/publication/9789264121164-en.

[4] OECD, *How's Life? 2020: Measuring Well-being* (How's Life?). OECD, 9th Mar. 2020, ISBN: 978-92-64-65467-9 978-92-64-97671-9 978-92-64-78116-0 978-92-64-72844-8. DOI: 10.1787/9870c393-en. Accessed: 4th Jul. 2022. [Online]. Available: https://www.oecd-ilibrary.org/economics/how-s-life/volume-/issue-_9870c393-en.

[5] Gauteng City-Region, Observatory and Gauteng City Region Observatory. 'Understanding Quality of Life | GCRO,' Accessed: 29th May 2022. [Online]. Available: https://gcro.ac.za/research/research-themes/detail/quality-life/.

[6] A. Deaton, 'Measuring and Understanding Behavior, Welfare, and Poverty,' *American Economic Review*, vol. 106, no. 6, pp. 1221–1243, Jun. 2016, ISSN: 0002-8282. DOI: 10.1257/aer.106.6.1221.

[7] A. Deaton, *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*. World Bank Publications, 1997, 492 pp., ISBN: 978-0-8018-5254-1.

[8] UN-Habitat, Ed., *World Cities Report 2020: The Value of Sustainable Urbanization* (World Cities Report 2020). Nairobi, Kenya: UN-Habitat, 2020, 377 pp., ISBN: 978-92-1-132872-1.

[9] R. M. Groves, F. J. F. Jr, M. P. Couper, J. M. Lepkowski, E. Singer and R. Tourangeau, *Survey Methodology*. John Wiley & Sons, 14th Jul. 2009, 487 pp., ISBN: 978-0-470-46546-2. Google Books: HXoSpXvo3s4C.

[10] Y. Naidoo and J. de Kadt, 'GCRO Quality of Life Survey 6 (2020/21): Quality of Life Index methodology,' Gauteng City-Region Observatory (GCRO), Johannesburg, 2021. [Online]. Available: https://cdn.gcro.ac.za/media/documents/2021.09.09_GCRO_Quality_of_Life_Survey_6_Quality_of_Life_Index_Methodology.pdf.

[11] K. B. Wright, 'Researching Internet-Based Populations: Advantages and Disadvantages of Online Survey Research, Online Questionnaire Authoring Software Packages, and Web Survey Services,' *Journal of Computer-Mediated Communication*, vol. 10, no. 3, pp. 00–00, 2005, ISSN: 1083-6101. DOI: 10.1111/j.1083-6101.2005.tb00259.x.

[12] S. Katumba, J. de Kadt, M. Orkin and P. Fatti, 'Construction of a Reflective Quality of Life Index for Gauteng Province in South Africa,' *Social Indicators Research*, 29th Jun. 2022, ISSN: 1573-0921. DOI: 10.1007/s11205-022-02945-2.

[13] United Nations, 'Human Development Index,' United Nations. Accessed: 4th Jul. 2022. [Online]. Available: https://hdr.undp.org/data-center/human-development-index.

[14] T. Greyling, 'A composite index of quality of life for the Gauteng city-region: A principal component analysis approach,' 31st Oct. 2017. DOI: 10/11312.

[15] A. Sen, 'Development as capability expansion,' *Human development and the international development strategy for the 1990s*, vol. 1, no. 1, 1990.

[16] J. E. Stiglitz, A. Sen and J.-P. Fitoussi, 'Report by the Commission on the Measurement of Economic Performance and Social Progress,' 1st Jan. 2009, p. 1. Accessed: 28th Mar. 2025. [Online]. Available: https://ui.adsabs.harvard.edu/abs/2009cmep.rept....1S.

[17] UN Habitat, *A Brief History of Quality of Life: From Aristotle to the Present, Evolution of a Concept at the Core of Human Development*, Aug. 2023. Accessed: 28th Mar. 2025. [Online]. Available: https://unhabitat.org/a-brief-history-of-quality-of-life-from-aristotle-to-the-present-evolution-of-a-concept-at-the-core.

[18] G. S. Becker, T. J. Philipson and R. R. Soares, 'The Quantity and Quality of Life and the Evolution of World Inequality,' *American Economic Review*, vol. 95, no. 1, pp. 277–291, Mar. 2005, ISSN: 0002-8282. DOI: 10.1257/0002828053828563.

[19] E. Stanton, 'The Human Development Index: A History,' Political Economy Research Institute, University of Massachusetts at Amherst, Working Paper, 2007. Accessed: 29th Mar. 2025. [Online]. Available: https://econpapers.repec.org/paper/umaperiwp/wp127.htm.

[20] OECD, *The Path to Becoming a Data-Driven Public Sector* (OECD Digital Government Studies). Paris: OECD Publishing, 28th Nov. 2019. Accessed: 16th May 2022. [Online]. Available: https://doi.org/10.1787/059814a7-en..

[21] Statistics South Africa (Stats SA). 'Census | Statistics South Africa,' Statistics South Africa, Accessed: 12th Jul. 2022. [Online]. Available: https://www.statssa.gov.za/?page_id=3836.

[22] 'The DHS Program - Quality information to plan, monitor and improve population, health, and nutrition programs,' Accessed: 28th Mar. 2025. [Online]. Available: https://dhsprogram.com/.

[23] World Bank, *Living Standards Measurement Study Brochure*, Sep. 2020. [Online]. Available: https://documents1.worldbank.org/curated/en/708961597206589588/pdf/Living-Standards-Measurement-Study.pdf.

[24] D. E. Sahn and D. Stifel, 'Exploring Alternative Measures of Welfare in the Absence of Expenditure Data,' *Review of Income and Wealth*, vol. 49, no. 4, pp. 463–489, 2003, ISSN: 1475-4991. DOI: 10.1111/j.0034-6586.2003.00100.x.

[25] 'Gauteng City-Region Observatory,' GCRO, Accessed: 28th Mar. 2025. [Online]. Available: https://www.gcro.ac.za/.

[26] M. Orkin, 'Technical Review of the GCRO QoL Surveys: Synthesis Report,' Gauteng City-Region Observatory, Johannesburg, 2020, p. 65. [Online]. Available: https://gcro.ac.za/research/project/detail/quality-life-survey-10-year-review/.

[27] M. Burke, A. Driscoll, D. B. Lobell and S. Ermon, 'Using satellite imagery to understand and promote sustainable development,' *Science*, vol. 371, no. 6535, eabe8628, 19th Mar. 2021. DOI: 10.1126/science.abe8628.

[28] M. Jerven, *Poor Numbers: How We Are Misled by African Development Statistics and What to Do about It.* Cornell University Press, 2013, ISBN: 978-0-8014-5163-8. JSTOR: 10.7591/j.ctt1xx4sk. Accessed: 25th Oct. 2022. [Online]. Available: https://www.jstor.org/stable/10.7591/j.ctt1xx4sk.

[29] M. Jerven, 'How Much Will a Data Revolution in Development Cost?' *Forum for Development Studies*, vol. 44, no. 1, pp. 31–50, Mar. 2017, ISSN: 08039410.

[30] *GCRO Quality of Life 7 (2023/24) Survey – Fieldwork and Data Preparation RFP Procurement Document*, Gauteng City Region Observatory (GCRO), 7th May 2023.

[31] D. J. de Kadt and K. Crowe-Pettersson, 'GCRO's Quality of Life Survey: Survey Management Workshop Briefing Document,' Gauteng City-Region Observatory, Johannesburg, 2020. [Online]. Available: https://gcro.ac.za/research/project/detail/quality-life-survey-10-%20year-review/.

[32] 'MICS: Multiple Indicator Cluster Surveys programme,' Accessed: 1st Sep. 2025. [Online]. Available: https://mics.unicef.org/home.

[33] C. Hamann and D. J. de Kadt, 'GCRO Quality of Life Survey 6 (2020/21): Sample design,' Gauteng City-Region Observatory, Johannesburg, 2021.

[34] 'The challenges of survey research in high crime communities,' SaferSpaces, Accessed: 28th Mar. 2025. [Online]. Available: https://www.saferspaces.org.za/blog/entry/the-challenges-of-survey-research-in-high-crime-communities.

[35] The World Bank, *Data on Statistical Capacity*, 2nd Mar. 2021. DOI: 10.57966/0q5s-2368.

[36] H.-A. H. Dang, J. Pullinger, U. Serajuddin and B. Stacy, 'Statistical performance indicators and index—a new tool to measure country statistical capacity,' *Scientific Data*, vol. 10, no. 1, p. 146, 1 20th Mar. 2023, ISSN: 2052-4463. DOI: 10.1038/s41597-023-01971-0.

[37] The World Bank, *Statistical Performance Indicators*, version 8, 12th Nov. 2024.

[38] O. Hall, F. Dompae, I. Wahab and F. M. Dzanku, 'A review of machine learning and satellite imagery for poverty prediction: Implications for development research and applications,' *Journal of International Development*, vol. 35, no. 7, pp. 1753–1768, 2023, ISSN: 1099-1328. DOI: 10.1002/jid.3751.

[39] G. Li and Q. Weng, 'Measuring the quality of life in city of Indianapolis by integration of remote sensing and census data,' *International Journal of Remote Sensing*, vol. 28, no. 2, pp. 249–267, 20th Jan. 2007, ISSN: 0143-1161. DOI: 10.1080/01431160600735624.

[40] R. Marty and A. Duhaut, 'Global poverty estimation using private and public sector big data sources,' *Scientific Reports*, vol. 14, no. 1, p. 3160, 7th Feb. 2024, ISSN: 2045-2322. DOI: 10.1038/s41598-023-49564-6.

[41] K. He, X. Zhang, S. Ren and J. Sun. 'Deep Residual Learning for Image Recognition.' arXiv: 1512.03385 `[cs]`, Accessed: 2nd Apr. 2023. [Online]. Available: http://arxiv.org/abs/1512.03385, pre-published.

[42] A. Head, M. Manguin, N. Tran and J. E. Blumenstock, 'Can Human Development be Measured with Satellite Imagery?' In *Proceedings of the Ninth International Conference on Information and Communication Technologies and Development*, ser. ICTD '17, New York, NY, USA: Association for Computing Machinery, 16th Nov. 2017, pp. 1–11, ISBN: 978-1-4503-5277-2. DOI: 10.1145/3136560.3136576. Accessed: 26th Jul. 2022. [Online]. Available: https://doi.org/10.1145/3136560.3136576.

[43] A. Daoud, F. Jordán, M. Sharma, F. Johansson, D. Dubhashi, S. Paul and S. Banerjee, 'Using Satellite Images and Deep Learning to Measure Health and Living Standards in India,' *Social Indicators Research*, vol. 167, no. 1, pp. 475–505, 1st Jun. 2023, ISSN: 1573-0921. DOI: 10.1007/s11205-023-03112-x.

[44] E. Suel, S. Bhatt, M. Brauer, S. Flaxman and M. Ezzati, 'Multimodal deep learning from satellite and street-level imagery for measuring income, overcrowding, and environmental deprivation in urban areas,' *Remote Sensing of Environment*, vol. 257, p. 112 339, 1st May 2021, ISSN: 0034-4257. DOI: 10.1016/j.rse.2021.112339.

[45] C. D. Elvidge, K. E. Baugh, E. A. Kihn, H. W. Kroehl, E. R. Davis and C. W. Davis, 'Relation between satellite observed visible-near infrared emissions, population, economic activity and electric power consumption,' *International Journal of Remote Sensing*, vol. 18, no. 6, pp. 1373–1379, 1st Apr. 1997, ISSN: 0143-1161. DOI: 10.1080/014311697218485.

[46] X. Chen and W. D. Nordhaus, 'Using luminosity data as a proxy for economic statistics,' *Proceedings of the National Academy of Sciences of the United States of America*, vol. 108, no. 21, pp. 8589–8594, 24th May 2011, ISSN: 0027-8424. DOI: 10.1073/pnas.1017031108. PMID: 21576474.

[47] J. V. Henderson, A. Storeygard and D. N. Weil, 'Measuring Economic Growth from Outer Space,' *The American Economic Review*, vol. 102, no. 2, pp. 994–1028, Apr. 2012, ISSN: 00028282. DOI: 10.1257/aer.102.2.994.

[48] M. A. Wulder, T. R. Loveland, D. P. Roy, C. J. Crawford, J. G. Masek, C. E. Woodcock, R. G. Allen, M. C. Anderson, A. S. Belward, W. B. Cohen, J. Dwyer, A. Erb, F. Gao, P. Griffiths, D. Helder, T. Hermosilla, J. D. Hipple, P. Hostert, M. J. Hughes, J. Huntington, D. M. Johnson, R. Kennedy, A. Kilic, Z. Li, L. Lymburner, J. McCorkel, N. Pahlevan, T. A. Scambos, C. Schaaf, J. R. Schott, Y. Sheng, J. Storey, E. Vermote, J. Vogelmann, J. C. White, R. H. Wynne and Z. Zhu, 'Current status of Landsat program, science, and applications,' *Remote Sensing of Environment*, vol. 225, pp. 127–147, 1st May 2019, ISSN: 0034-4257. DOI: 10.1016/j.rse.2019.02.015.

[49] C. D. Elvidge, K. Baugh, M. Zhizhin, F. C. Hsu and T. Ghosh, 'VIIRS night-time lights,' *International Journal of Remote Sensing*, vol. 38, no. 21, pp. 5860–5879, 2nd Nov. 2017, ISSN: 0143-1161. DOI: 10.1080/01431161.2017.1342050.

[50] 'PlanetScope - Earth Online,' Accessed: 28th Aug. 2022. [Online]. Available: https://earth.esa.int/eogateway/missions/planetscope.

[51] S. Piaggesi, L. Gauvin, M. Tizzoni, C. Cattuto, N. Adler, S. Verhulst, A. Young, R. Price, L. Ferres and A. Panisson, 'Predicting City Poverty Using Satellite Imagery,' presented at the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, 2019, pp. 90–96. Accessed: 23rd Feb. 2025. [Online]. Available: https://openaccess.thecvf.com/content_CVPRW_2019/html/cv4gc/Piaggesi_Predicting_City_Poverty_Using_Satellite_Imagery_CVPRW_2019_paper.html.

[52] J. Machicao, A. Specht, D. Vellenich, L. Meneguzzi, R. David, S. Stall, K. Ferraz, L. Mabile, M. O'Brien and P. Corrêa, 'A Deep-Learning Method for the Prediction of Socio-Economic Indicators from Street-View Imagery

Using a Case Study from Brazil,' *Data Science Journal*, vol. 21, no. 1, p. 6, 1 11th Feb. 2022, ISSN: 1683-1470. DOI: 10.5334/dsj-2022-006.

[53] J. L. Abitbol and M. Karsai, 'Interpretable socioeconomic status inference from aerial imagery through urban patterns,' *Nature Machine Intelligence*, vol. 2, no. 11, pp. 684–692, Nov. 2020. DOI: 10.1038/s42256-020-00243-5.

[54] 'Sustainable Development Report 2022,' Accessed: 3rd Jun. 2023. [Online]. Available: https://dashboards.sdgindex.org/.

[55] 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, 19th Aug. 2016. DOI: 10.1126/science.aaf7894.

[56] C. Yeh, A. Perez, A. Driscoll, G. Azzari, Z. Tang, D. Lobell, S. Ermon and M. Burke, 'Using publicly available satellite imagery and deep learning to understand economic well-being in Africa,' *Nature Communications*, vol. 11, no. 1, p. 2583, 1 22nd May 2020, ISSN: 2041-1723. DOI: 10.1038/s41467-020-16185-w.

[57] K. Ayush, B. Uzkent, M. Burke, D. Lobell and S. Ermon. 'Generating Interpretable Poverty Maps using Object Detection in Satellite Images.' arXiv: 2002.01612 `[cs]`, Accessed: 4th Jul. 2022. [Online]. Available: http://arxiv.org/abs/2002.01612, pre-published.

[58] D. Everatt, 'Poverty and inequality in the Gauteng city-region,' in *Changing Space, Changing City*, ser. Johannesburg after Apartheid - Open Access Selection, P. Harrison, G. Gotz, A. Todes and C. Wray, Eds., Wits University Press, 2014, pp. 63–82, ISBN: 978-1-86814-765-6. DOI: 10.18772/22014107656.7. JSTOR: 10.18772/22014107656.7. Accessed: 29th Nov. 2022. [Online]. Available: https://www.jstor.org/stable/10.18772/22014107656.7.

[59] P. Nussbaumer, M. Bazilian and V. Modi, 'Measuring energy poverty: Focusing on what matters,' *Renewable and Sustainable Energy Reviews*, vol. 16, no. 1, pp. 231–243, 1st Jan. 2012, ISSN: 1364-0321. DOI: 10.1016/j.rser.2011.07.150.

[60] S. K. T. Baah, D. M. Jolliffe, D. G. Mahler, C. Lakner and A. Atamanov, 'Assessing the Impact of the 2017 PPPs on the International Poverty Line

and Global Poverty,' *Policy Research Working Paper*, vol. 9941, © World Bank 2022. HDL: 10986/37061.

[61] A. Bruederle and R. Hodler. 'Nighttime Lights as a Proxy for Human Development at the Local Level.' Social Science Research Network: 3013006, Accessed: 23rd Feb. 2025. [Online]. Available: https://papers.ssrn.com/abstract=3013006, pre-published.

[62] C. D. Elvidge, P. C. Sutton, T. Ghosh, B. T. Tuttle, K. E. Baugh, B. Bhaduri and E. Bright, 'A global poverty map derived from satellite data,' *Computers & Geosciences*, vol. 35, no. 8, pp. 1652–1660, 1st Aug. 2009, ISSN: 0098-3004. DOI: 10.1016/j.cageo.2009.01.009.

[63] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg and L. Fei-Fei, 'ImageNet Large Scale Visual Recognition Challenge,' *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 1st Dec. 2015, ISSN: 1573-1405. DOI: 10.1007/s11263-015-0816-y.

[64] R. Engstrom, J. Hersh and D. Newhouse, 'Poverty from Space: Using High Resolution Satellite Imagery for Estimating Economic Well-being,' *The World Bank Economic Review*, vol. 36, no. 2, pp. 382–412, 2nd May 2022, ISSN: 0258-6770. DOI: 10.1093/wber/lhab015.

[65] J. Blumenstock, G. Cadamuro and R. On, 'Predicting poverty and wealth from mobile phone metadata,' *Science*, vol. 350, no. 6264, pp. 1073–1076, 27th Nov. 2015. DOI: 10.1126/science.aac4420.

[66] J. E. Steele, P. R. Sundsøy, C. Pezzulo, V. A. Alegana, T. J. Bird, J. Blumenstock, J. Bjelland, K. Engø-Monsen, Y.-A. de Montjoye, A. M. Iqbal, K. N. Hadiuzzaman, X. Lu, E. Wetter, A. J. Tatem and L. Bengtsson, 'Mapping poverty using mobile phone and satellite data,' *Journal of The Royal Society Interface*, vol. 14, no. 127, p. 20 160 690, 28th Feb. 2017. DOI: 10.1098/rsif.2016.0690.

[67] E. Sheehan, C. Meng, M. Tan, B. Uzkent, N. Jean, M. Burke, D. Lobell and S. Ermon, 'Predicting Economic Development using Geolocated Wikipedia Articles,' in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, ser. KDD '19, New York, NY, USA:

Association for Computing Machinery, 25th Jul. 2019, pp. 2698–2706, ISBN: 978-1-4503-6201-6. DOI: 10.1145/3292500.3330784. Accessed: 28th Jul. 2022. [Online]. Available: https://doi.org/10.1145/3292500.3330784.

[68] G. Chi, H. Fang, S. Chatterjee and J. E. Blumenstock, 'Microestimates of wealth for all low- and middle-income countries,' *Proceedings of the National Academy of Sciences*, vol. 119, no. 3, e2113658119, 18th Jan. 2022. DOI: 10.1073/pnas.2113658119.

[69] L. Espın-Noboa, J. Kertész and M. Karsai, 'Challenges of inferring high-resolution poverty maps with multimodal data,' 2022.

[70] 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: 1711.06323 `[cs, stat]`, Accessed: 25th Oct. 2022. [Online]. Available: http://arxiv.org/abs/1711.06323, pre-published.

[71] K. Lee and J. Braithwaite, 'High-resolution poverty maps in Sub-Saharan Africa,' *World Development*, vol. 159, p. 106 028, 1st Nov. 2022, ISSN: 0305-750X. DOI: 10.1016/j.worlddev.2022.106028.

[72] E. Rolf, J. Proctor, T. Carleton, I. Bolliger, V. Shankar, M. Ishihara, B. Recht and S. Hsiang, 'A generalizable and accessible approach to machine learning with global satellite imagery,' *Nature Communications*, vol. 12, no. 1, p. 4392, 1 20th Jul. 2021, ISSN: 2041-1723. DOI: 10.1038/s41467-021-24638-z.

[73] L. Sherman, J. Proctor, H. Druckenmiller, H. Tapia and S. Hsiang. 'Global High-Resolution Estimates of the United Nations Human Development Index Using Satellite Imagery and Machine-Learning,' Accessed: 6th Jul. 2023. [Online]. Available: https://papers.ssrn.com/abstract=4393385, pre-published.

[74] E. Suel, J. W. Polak, J. E. Bennett and M. Ezzati, 'Measuring social, environmental and health inequalities using deep learning and street imagery,' *Scientific Reports*, vol. 9, no. 1, p. 6229, 1 18th Apr. 2019, ISSN: 2045-2322. DOI: 10.1038/s41598-019-42036-w.

[75] E. S. Steyn, *South Africa Street History Mapping*. [Online]. Available: https://github.com/Emily-RoseSteyn/south-africa-street-history-mapping.

[76] E. S. Steyn, *Measuring Quality of Life from Satellite Images in Gauteng, South Africa*. [Online]. Available: https://github.com/Emily-RoseSteyn/measuring-quality-of-life-gauteng.

[77] K. Beck, M. Beedle, A. van Bennekum, A. Cockburn, W. Cunningham, M. Fowler, J. Grenning, J. Highsmith, A. Hunt, R. Jeffries, J. Kern, B. Marick, R. C. Martin, S. Mellor, K. Schwaber, J. Sutherland and D. Thomas, *The Agile Manifesto*, 2001. Accessed: 4th Apr. 2024. [Online]. Available: http://agilemanifesto.org/.

[78] J. Thornhill, 'A space revolution: Do tiny satellites threaten our privacy?' *FT.com*, 15th Feb. 2018.

[79] C. Beam. 'Soon, satellites will be able to watch you everywhere all the time,' MIT Technology Review, Accessed: 29th Mar. 2025. [Online]. Available: https://www.technologyreview.com/2019/06/26/102931/satellites-threaten-privacy/.

[80] C. Santos and L. Rapp, 'Satellite Imagery, Very High-Resolution and Processing-Intensive Image Analysis: Potential Risks Under the GDPR,' *Air and Space Law*, vol. 44, no. 3, 1st Jun. 2019, ISSN: 0927-3379.

[81] M. M. Coffer, 'Balancing Privacy Rights and the Production of High-Quality Satellite Imagery,' *Environmental Science & Technology*, vol. 54, no. 11, pp. 6453–6455, 2nd Jun. 2020, ISSN: 0013-936X. DOI: 10.1021/acs.est.0c02365.

[82] Gauteng City-Region Observatory (GCRO), *Quality of Life Survey V 2017-2018 [dataset]. Round 5*, version 2, Johannesburg: GCRO [producer], 2022. Cape Town: DataFirst [distributor], 2022. DOI: 10.25828/8yf7-9261. Accessed: 25th May 2023. [Online]. Available: https://www.datafirst.uct.ac.za/dataportal/index.php/catalog/766/.

[83] Gauteng City-Region Observatory (GCRO), *Quality of Life Survey 2020-2021 [dataset]. Round 6*, version 1, Johannesburg: GCRO [producer], 2021. Cape Town: DataFirst [distributor], 2022. DOI: 10.25828/WEMZ-VF31. Accessed: 25th May 2023. [Online]. Available: https://www.datafirst.uct.ac.za/dataportal/index.php/catalog/874/.

[84] Planet Labs PBC, *Planet Application Program Interface: In Space for Life on Earth*, Planet, 2023. Accessed: 28th Aug. 2022. [Online]. Available: https://api.planet.com.

[85] Municipal Demarcation Board, *MDB Wards 2020*, Shapefile, 11th Dec. 2020. [Online]. Available: https://dataportal-mdb-sa.opendata.arcgis.com/datasets/e0223a825ea2481fa72220ad3204276b.

[86] Municipal Demarcation Board, *MDB Wards 2016*, Shapefile, 7th Jun. 2017. [Online]. Available: https://dataportal-mdb-sa.opendata.arcgis.com/datasets/cfddb54aab5f4d62b2144d80d49b3fdb/about.

[87] The pandas development team, *Pandas-dev/pandas: Pandas*, version latest, Zenodo, Feb. 2020. DOI: 10.5281/zenodo.3509134. [Online]. Available: https://doi.org/10.5281/zenodo.3509134.

[88] J. V. den Bossche, K. Jordahl, M. Fleischmann, M. Richards, J. McBride, J. Wasserman, A. G. Badaracco, A. D. Snow, B. Ward, J. Tratner, J. Gerard, M. Perry, cjqf, G. A. Hjelle, M. Taves, E. ter Hoeven, M. Cochran, R. Bell, rraymondgh, M. Bartos, P. Roggemans, L. Culbertson, G. Caria, N. Y. Tan, N. Eubank, sangarshanan, J. Flavin, S. Rey and J. Gardiner, *Geopandas/geopandas: V1.0.1*, version v1.0.1, Zenodo, 2nd Jul. 2024. DOI: 10.5281/zenodo.12625316. Accessed: 28th Mar. 2025. [Online]. Available: https://zenodo.org/records/12625316.

[89] S. Gillies and and team, *Rasterio: Geospatial raster I/O for Python programmers*, Mapbox, 2013–. [Online]. Available: https://github.com/rasterio/rasterio.

[90] R. Sefala, T. Gebru, L. Mfupe, N. Moorosi and R. Klein, 'Constructing a Visual Dataset to Study the Effects of Spatial Apartheid in South Africa,' presented at the Thirty-Fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2), 19th Nov. 2021. Accessed: 25th Apr. 2023. [Online]. Available: https://openreview.net/forum?id=WV0waZz9dTF.

[91] T. Developers, *TensorFlow*, version v2.19.0, Zenodo, 12th Mar. 2025. DOI: 10.5281/zenodo.15009305. Accessed: 28th Mar. 2025. [Online]. Available: https://zenodo.org/records/15009305.

[92] Python Software Foundation, *Python Language Reference*, version 3.10. [Online]. Available: http://www.python.org.

[93] T. M. D. Team, *Matplotlib: Visualization with Python*, version v3.10.1, Zenodo, 27th Feb. 2025. DOI: 10.5281/zenodo.14940554. Accessed: 28th Mar. 2025. [Online]. Available: https://zenodo.org/records/14940554.

[94] 'DVC: Data Version Control - Git for Data & Models,' DOI: 10.5281/zenodo.8421832.

[95] S. Eustace and The Poetry contributors, *Poetry: Python packaging and dependency management made easy.* [Online]. Available: https://github.com/python-poetry/poetry.

[96] E. Aiken, E. Rolf and J. Blumenstock. 'Fairness and representation in satellite-based poverty maps: Evidence of urban-rural disparities and their impacts on downstream policy.' arXiv: 2305.01783 `[cs]`, Accessed: 28th Sep. 2024. [Online]. Available: http://arxiv.org/abs/2305.01783, pre-published.

[97] D. Lam, R. Kuzma, K. McGee, S. Dooley, M. Laielli, M. Klaric, Y. Bulatov and B. McCord. 'xView: Objects in Context in Overhead Imagery.' arXiv: 1802.07856 `[cs]`, Accessed: 29th Nov. 2022. [Online]. Available: http://arxiv.org/abs/1802.07856, pre-published.

A Source Code

A.1. Introduction

The source code for the dissertation titled: "Measuring quality of life in Gauteng, South Africa from satellite images using machine learning" is available on Github [1]. The repository contains all scripts used for data processing, model training, evaluation, and visualisation.

Installation and dependency information is provided in the repository README file. The code is implemented using Python 3.10 [2] with key dependencies including TensorFlow 2.18 [3], GeoPandas [4].

The version of code used to generate the results presented in this dissertation corresponds to release v1.0.0 of the repository.

References

[1] E. S. Steyn, *Measuring Quality of Life from Satellite Images in Gauteng, South Africa*. [Online]. Available: https://github.com/Emily-RoseSteyn/measuring-quality-of-life-gauteng.

[2] Python Software Foundation, *Python Language Reference*, version 3.10. [Online]. Available: http://www.python.org.

[3] T. Developers, *TensorFlow*, version v2.19.0, Zenodo, 12th Mar. 2025. DOI: 10.5281/zenodo.15009305. Accessed: 28th Mar. 2025. [Online]. Available: https://zenodo.org/records/15009305.

[4] J. V. den Bossche, K. Jordahl, M. Fleischmann, M. Richards, J. McBride, J. Wasserman, A. G. Badaracco, A. D. Snow, B. Ward, J. Tratner, J. Gerard, M. Perry, cjqf, G. A. Hjelle, M. Taves, E. ter Hoeven, M. Cochran, R. Bell, rraymondgh, M. Bartos, P. Roggemans, L. Culbertson, G. Caria, N. Y. Tan, N. Eubank, sangarshanan, J. Flavin, S. Rey and J. Gardiner, *Geopandas/geopandas: V1.0.1*, version v1.0.1, Zenodo, 2nd Jul. 2024. DOI: 10.5281/zenodo.12625316. Accessed: 28th Mar. 2025. [Online]. Available: https://zenodo.org/records/12625316.

B Ethics Clearance

B.1. Introduction

Research ethics is a primary consideration of academic endeavours at the University of the Witwatersrand. The dissertation titled: "Measuring quality of life in Gauteng, South Africa from satellite images using machine learning" has been conducted in accordance with the University's ethical guidelines. The remainder of this appendix includes the University's ethics risk assessment table and the ethics waiver letter for the conducted research. These artifacts verify that this research has been conducted in accordance with the University of the Witwatersrand's ethical standards and requirements.

HREC (Non-Medical) Risk level categories definitions (January 2020)

This table identifies broad categories of risk. Schools/Departments can provide specific examples of these categories that are specific to that particular discipline, or the types of data collection methods or participant groups that are most common in that discipline. Please note that any study involving minors cannot be considered by Schools irrespective of the risk level.

Risk category	Definition	Example	Notes
No risk	**No contact with human participants**	Document analysis or literature review Studies based on theoretical or secondary analysis alone Use of non-human, quantitative datasets (e.g. economic data)	These studies do not require ethics clearance A waiver may be given, however, if required by a university faculty or external body
		Use of previously-collected human datasets (where permission from previous participants have been explicitly granted) Use of anonymized and aggregated human datasets (e.g. census data)	These studies may require ethics clearance, dependent on the type of study and faculty requirements A waiver may be given, however, if required by a university faculty or external body Applications deemed No Risk can be considered at School level
Minimal risk	Where the likelihood and magnitude of possible harm are no greater than those imposed by daily life in a stable society, or routine educational or psychological tests	Questions about people's everyday lives, activities and opinions rather than detailed biographical information No sensitive questions or topics Review of privileged information (e.g. documentation not publically available)	Applications deemed Minimal Risk can be considered at School level
Low risk	Where the only foreseeable risks is that of discomfort, or where there may be some sensitivity involved in terms of the questions asked	Questions about people's everyday lives, activities and opinions – may include biographical information and some potentially sensitive questions and/or topics May include some vulnerable contexts	Applications deemed Low Risk can be considered at School level
Medium risk	Where there is a likely risk of some harm for participants and/or the researcher, but where appropriate steps can be taken to mitigate or reduce risk	Sensitive topics and/or questions that may have potential for trauma and emotional distress May include vulnerable categories or marginalized groups, may include some types of low-level illegal activities, such as artisanal mining Research locality itself may contain potential risks to the participants and/or researcher There is a clear justification to undertake the research using this participant group and/or using the proposed instruments, despite the potential risks	Applications deemed Medium Risk cannot be considered at School level and must be referred to the main committee Support/counselling services must be provided for participants, if appropriate A distress protocol should be given, if appropriate
High risk	Where there is a real and foreseeable risk of harm which may lead to serious	Highly sensitive topics, e.g. experiences of violence, rape, illegal activities Vulnerable or marginalized groups, or where multiple vulnerabilities exist Research involving deception of the participants	Applications deemed High Risk cannot be considered at School level and must be referred to the main committee

UNIVERSITY OF THE WITWATERSRAND, JOHANNESBURG

HUMAN RESEARCH ETHICS COMMITTEE (NON-MEDICAL)

Registration number: REC-101114-044

30 March 2023

Re: Mrs Emily-Rose Simoes Steyn (0405400P)

Waiver letter number: HRECNMW23/04/03

To whom it may concern,

Mrs Steyn is a registered student at the School of Electrical and Information Engineering at the University of the Witwatersrand, Johannesburg. This letter is to confirm that, at the time of writing, Mrs. Steyn does not need ethical clearance for her Masters study entitled *'Measuring Quality of Life in Gauteng Using Satellite Imagery and Machine Learning'.* This decision has been reached based upon a description of the project supplied by Mrs. Steyn to the University Human Research Ethics Committee (Non-Medical), which has been evaluated by the Chairs and Deputy Chairs. If, however, if Mrs. Steyn changes the methods of data collection and analysis for this study, this decision may no longer be valid. If such changes take place, this should be communicated to the University Human Research Ethics Committee (Non-Medical) as soon as possible. This waiver letter is valid until 29 March 2026.

Please feel free to contact me should you require any further information.

Thank you.

Yours sincerely,
S Schoeman

Shaun Schoeman (Administrative Officer)
Solomon Mahlangu House, 10th Floor, Room 10004, Jorissen Street, Braamfontein, Johannesburg
Private Bag 3, Wits 2050
T + 27(0)11 717 1408 | **E** Shaun.Schoeman@wits.ac.za | hrec-medical.researchoffice@wits.ac.za
www.wits.ac.za/research/about-our-research/ethics-and-research-integrity/