


Child Stunting in Tanzania's Councils

Small area estimates of stunting, severe stunting and height-for-age for 195 councils and 150 districts, 2022



Suggested citation. SpacePilots (2026). *Child Stunting in Tanzania's Councils: Small Area Estimates of Stunting, Severe Stunting and Height-for-Age, 2022*. SpacePilots Research Report. Copenhagen: SpacePilots.

Disclaimer. The estimates in this report are model-based and were produced by SpacePilots. They are not official statistics of the National Bureau of Statistics (NBS), the Office of the Chief Government Statistician Zanzibar (OCGS) or the Ministry of Health. Official region and national figures are those published in the TDHS-MIS 2022 final report. Every estimate should be read with its confidence interval and reliability class.

Data. Tanzania Demographic and Health Survey and Malaria Indicator Survey 2022 (The DHS Program); 2022 Population and Housing Census 10% sample and administrative boundaries (NBS); Google Open Buildings 2.5D Temporal, Landsat 8 NDVI and Google Satellite Embedding (Google Earth Engine). Maps show the 2022 census boundaries and do not imply any judgement on the legal status of any territory.

Contact. Questions about this report, the estimates or the underlying files can be sent to SpacePilots at info@spacepilots.io.

Preview edition. This is an abridged edition of the full report. It shows the approach, the main findings and the quality of the estimates. Council, district and ward estimates with their confidence intervals, the full set of maps and the technical appendices are in the full report and data package (Chapter 8).

© SpacePilots, Copenhagen, 2026.

Contents

Executive summary	4
Abbreviations	5
1 Introduction	6
1.1 Background	6
1.2 The problem: too few children per council	6
1.3 Previous small area estimates of malnutrition	6
1.4 Objectives	6
1.5 Structure of the report	6
2 Data	7
2.1 Data sources	7
2.2 Survey outcomes	8
2.3 Census	8
2.4 Harmonizing survey and census variables	8
2.5 Survey cluster locations	8
3 Methodology	9
3.1 Overview of the estimators	9
3.2 How the data are used	9
3.3 Direct estimates	9
3.4 The main model	10
3.5 From model to area estimates	10
3.6 Uncertainty	10
3.7 Alternative 1–4: area-level Fay–Herriot models	11
3.8 Alternatives 5–6: machine-learning models	11
3.9 Model selection	11
3.10 Benchmarking and indicators reported for regions only	11
4 Results	13
4.1 Stunting prevalence by council	13
4.2 Number of stunted children	14
4.3 Prevalence and numbers compared	14
4.4 Severe stunting	15
4.5 Mean height-for-age	15
4.6 Districts and regions	15
4.7 Ward estimates (experimental)	16
5 Quality and validation	18
5.1 Precision	18
5.2 Gain over direct survey estimates	18
5.3 Independent accuracy test	18
5.4 Consistency with official figures	19
6 Discussion and limitations	20
6.1 What drives the results	20
6.2 Why the household-level model outperforms	20
6.3 Limitations	20
7 Conclusions and recommendations	21
7.1 Conclusions	21
7.2 Recommendations for users	21
References	22
8 The full report and data package	23
8.1 Data package	23
8.2 What the full report adds	23

LIST OF FIGURES

Figure 1	Stunting prevalence among children under 5 by council, 2022, in four classes	13
Figure 2	Estimated number of stunted children under 5 by council, 2022, in four classes	14
Figure 3	Stunting prevalence and number of stunted children under 5 by council, 2022	15
Figure 4	Experimental stunting estimates by ward, 2022, in four classes	17
Figure 5	Split-sample validation error for council stunting, by method	19

LIST OF TABLES

Table 1	Summary of data sources	7
Table 2	The estimators compared	9
Table 3	Priority councils: the ten with the highest stunting prevalence and the ten with the most stunted children, 2022	14
Table 4	Region estimates: official TDHS-MIS figures and aggregated council estimates, 2022	16
Table 5	Precision of the estimates by level and indicator	18
Table 6	Split-sample validation results, council level (estimates fitted on half the survey)	19
Table 7	Contents of the data package	23

Executive summary

Stunting exceeds 30% in 72 of Tanzania's 195 councils, and prevalence varies more than fivefold: from 11% in Moshi Municipal Council to 58% in Kilolo District Council. Region figures hide much of this: councils in the same region typically differ by 10 percentage points.

2.7 million

children under 5 stunted in 2022, 0.8 million of them severely

11–58%

range of stunting prevalence across the 195 councils

171 of 195

council estimates meet the 20% CV reliability threshold

3.5×

more accurate than council estimates from the survey alone, in an independent test

Purpose. The 2022 Tanzania Demographic and Health Survey and Malaria Indicator Survey (TDHS-MIS) gives reliable stunting figures for the nation and its 31 regions, but planning and budgeting happen in councils. With about 25 measured children per council, direct survey estimates for councils are too imprecise to use (median CV 35%). This report fills that gap.

Approach. A statistical model links each surveyed child's height-for-age to household characteristics measured in the same way in the survey and the 2022 census, including the mother's education, plus council features from the census and satellite imagery. The model is then applied to all 940,619 children under 5 in the census 10% sample. The method was chosen from seven candidates, including two machine-learning approaches, because it was the most accurate in an independent test.

KEY FINDINGS

- **Highest prevalence.** 16 councils exceed 40% stunting, all but one (Ngorongoro DC) in Iringa, Njombe and Rukwa. Kilolo DC (58%), Makete DC (57%) and Mufindi DC (57%) are highest and also have the highest severe stunting, at 23–24%.
- **Largest numbers.** About 2.7 million children under 5 are stunted, 0.8 million of them severely. The 40 councils with the most stunted children hold 42% of the total. Geita DC (64,000), Sumbawanga DC (44,000), Kaliua DC (41,000) and Chato DC (40,000) have the largest numbers.
- **Prevalence and numbers point to different places.** The highest-prevalence fifth of councils holds only 30% of stunted children. Targeting by rate and targeting by numbers give different priority lists (Section 4.3).
- **Consistency with official figures.** Aggregated to region, the council estimates agree with the published TDHS-MIS region figures within sampling error in all 31 regions.

Reliability. 171 of 195 council stunting estimates meet the standard reliability threshold (CV below 20%); the other 24 are “use with caution” and none is unreliable. The median CV is 16% and the median 95% confidence interval is ± 8 points. In an independent test, the typical estimation error of the council estimates (the root mean squared error) was about 29% as large as that of council estimates taken from the survey alone, a roughly 3.5-fold reduction in estimation error. Moving from direct survey estimates to these small area estimates therefore makes council figures about three and a half times as accurate.

Recommendations

1. Use the council estimates to rank and target councils and to size programmes by the number of stunted children, always with their confidence intervals.
2. Use both prevalence and numbers when setting priorities: high-prevalence councils in the Southern Highlands and populous councils in the Lake Zone and the west.
3. Use ward figures for allocation only with great caution; they are experimental.

Abbreviations

Abbreviation	Meaning
CI	Confidence interval
CV	Coefficient of variation (standard error divided by the estimate)
DC, TC, MC, CC	District, Town, Municipal and City Council
EBP	Empirical best predictor
FH	Fay–Herriot (area-level) model
HAZ	Height-for-age z-score (WHO 2006 Child Growth Standards)
LGA	Local government authority
MSE, RMSE	Mean squared error, root mean squared error
ML	Machine learning
NBS	National Bureau of Statistics
OCGS	Office of the Chief Government Statistician, Zanzibar
PHC	Population and Housing Census
pp	Percentage points
REML	Restricted maximum likelihood
SAE	Small area estimation
SE	Standard error
TDHS-MIS	Tanzania Demographic and Health Survey and Malaria Indicator Survey

Names of the estimation methods

Name used	Method
Main model	Unit-level census model (EBP) of child height-for-age; source of all headline estimates (Sections 3.4–3.6)
FH-census, FH-geo, FH-census+geo, FH-hybrid	Area-level Fay–Herriot models with census, satellite, both, or both plus a machine-learning prediction as covariates (Section 3.7)
ML-unit	Child-level machine-learning model (gradient boosting with LightGBM, or regularized logistic or ridge regression) applied to census children (Section 3.8)
ML-grid	Cluster-level machine-learning model (gradient boosting with LightGBM, or ridge regression) predicted on a ~500 m satellite grid (Section 3.8)
Direct estimate	Survey-weighted estimate from the TDHS-MIS children in the area alone

1 Introduction

This report estimates child stunting for every council in Tanzania in 2022, the level at which nutrition programmes are planned and budgeted.

1.1 Background

Stunting, low height for age, is the main indicator of chronic undernutrition in early childhood. It reflects long-term deficits in diet, health and care, and is associated with lower schooling, earnings and adult health. A child is stunted when their height-for-age z-score is more than 2 standard deviations below the median of the WHO Child Growth Standards (WHO Multicentre Growth Reference Study Group 2006), and severely stunted when it is more than 3 below.

The 2022 TDHS-MIS found that 30% of children under 5 in Tanzania were stunted (Ministry of Health et al. 2023). The survey is designed to give reliable estimates for the nation and its 31 regions. It is not designed for councils, the local government authorities that plan, budget and deliver nutrition interventions.

1.2 The problem: too few children per council

The survey measured 5,516 children with valid height-for-age, an average of about 25 per council. Six councils had no sampled households at all. Direct survey estimates for councils are therefore very imprecise, with a median coefficient of variation of 35%. 88% of councils would fail the usual reliability threshold of 20%.

Small area estimation (SAE) addresses this by borrowing strength from auxiliary data that cover every area, here the 2022 Population and Housing Census and satellite imagery. The census and survey were carried out only months apart, which makes them well suited to being combined.

1.3 Previous small area estimates of malnutrition

Small area estimation was developed for poverty mapping (Elbers, Lanjouw and Lanjouw 2003; Molina and Rao 2010) and is now standard practice in national statistics offices and the World Bank (Rao and Molina 2015; Corral et al. 2022). Its first application to child nutrition in Africa was in Tanzania: Simler (2006) combined household survey and census data to map undernutrition by district, and concluded that nutrition mapping is feasible. Models of child anthropometry explain less of the variation than poverty models, but they produced a plausible spatial pattern. The approach has since been applied to stunting and underweight in Cambodia (Fujii 2010) and to all woredas of Ethiopia (Sohnesen et al. 2017). Area-level models have also been used, for example spatio-temporal Fay–Herriot models of zonal undernutrition in Ethiopia (Yilema et al. 2022), and geostatistical models have mapped child growth failure across Africa on a 5 × 5 km grid (Osgood-Zimmerman et al. 2018).

This report builds on that work in three ways. The survey and the census were carried out in the same year, whereas earlier census-based maps had to bridge gaps of several years (seven in Ethiopia); satellite data add information that neither source measures; and seven statistical and machine-learning estimators are compared on independent data before one is chosen.

1.4 Objectives

1. Estimate stunting prevalence, severe stunting prevalence, mean height-for-age and the number of stunted children for all 195 councils and 150 districts, with standard errors, confidence intervals and reliability classes.
2. Choose the estimation method by measured accuracy rather than by assumption, comparing model-based and machine-learning approaches on independent data.
3. Check that the estimates are consistent with the official TDHS-MIS region figures.
4. Assess whether finer geographies (wards, ~500 m grid cells) can be supported.

1.5 Structure of the report

Chapter 2 describes the data and Chapter 3 the methodology. Chapter 4 presents the results and Chapter 5 their quality and validation. Chapter 6 discusses limitations and Chapter 7 concludes. Chapter 8 describes the full report and data package, including the technical appendices that are not part of this preview edition.

2 Data

The estimates combine a national household survey, the 2022 census and satellite-derived data. All are publicly available or available on request from their producers, and all refer to 2022 or close to it: the survey (February–July 2022) and the census (August 2022) were carried out only months apart.

2.1 Data sources

Tanzania Demographic and Health Survey and Malaria Indicator Survey 2022 (TDHS-MIS). A nationally representative household survey conducted by the Ministries of Health of Tanzania Mainland and Zanzibar, NBS and OCGS, with technical assistance from ICF through The DHS Program (Ministry of Health et al. 2023). The survey measured the height and weight of children under 5 in a stratified two-stage cluster sample. The data were downloaded from [The DHS Program website](#) after registration and approval of a data request. Three files were used:

- the **household member (PR) recode file**, which holds the anthropometric measurements of children aged 0–59 months and the characteristics of their households;
- the **GPS data set (TZGE81FL)**, which gives the location of each survey cluster, displaced for confidentiality;
- the **geospatial covariates data set (TZGC81FL)**, which gives summaries of 135 environmental and demographic rasters at each cluster location, used here for comparison only.

The published region estimates used to check the results were taken from the DHS Program's indicator data API.

Population and Housing Census 2022, 10% sample (PHC). A 10 percent sample of all household and person records from the 2022 Population and Housing Census, conducted by NBS and OCGS in August 2022 (National Bureau of Statistics and Office of the Chief Government Statistician 2022). The data were downloaded from the [National Bureau of Statistics website](#). The household file and five person files were used: demographics, education, fertility (women aged 10 and over), national documents and orphanhood, and economic activity. Together they cover 1.42 million households, 6.2 million persons and 940,619 children aged 0–4.

PHC 2022 administrative boundaries. The post-census geodatabase published by NBS, with the boundaries of 31 regions, 150 districts, 195 councils and 4,339 wards used in the 2022 census. These define the estimation units and are used for all maps.

Satellite-derived data. Three products were exported from [Google Earth Engine](#):

- **Google Open Buildings 2.5D Temporal** (Google Research): annual estimates of building presence, counts and heights for 2018–2023, aggregated to a grid of about 500 m. Used for building density, height and built-up fraction, and to weight population and survey cluster locations towards settlements.
- **Landsat 8 vegetation index (NDVI), 2022:** an annual composite at 100 m, computed from Landsat 8 surface reflectance (USGS/NASA).
- **Google Satellite Embedding, 2024** (Google DeepMind): a 64-band annual representation of the land surface learned from multiple satellite sources, aggregated to 100 m. Its first 15 principal components (88% of the variance) were used as covariates.

Table 1. Summary of data sources

Data set	Producer	Obtained from	Content used	Role in the analysis
TDHS-MIS 2022, PR file	MoH, NBS, OCGS, ICF	The DHS Program (registration required)	5,516 children 0–59 months with valid height-for-age; 625 clusters	Outcomes; child and household covariates
TDHS-MIS 2022, GPS file (TZGE81FL)	The DHS Program	The DHS Program	Displaced cluster coordinates	Linking clusters to councils
TDHS-MIS 2022, covariates (TZGC81FL)	The DHS Program	The DHS Program	135 raster summaries at cluster points	Comparison only
PHC 2022, 10% sample	NBS, OCGS	NBS website	1.42 million households; 940,619 children aged 0–4; education and fertility of all members	Covariates; prediction population; under-5 counts
PHC 2022 boundaries	NBS	NBS	31 regions, 150 districts, 195 councils, 4,339 wards	Estimation units; maps
Open Buildings 2.5D Temporal, 2018–2023	Google Research	Google Earth Engine	Building density, height, built-up fraction (~500 m)	Covariates; settlement weights
Landsat 8 NDVI, 2022	USGS/NASA	Google Earth Engine	Vegetation index (100 m)	Covariate
Satellite Embedding, 2024	Google DeepMind	Google Earth Engine	64-band land-surface representation (100 m)	Covariates (15 principal components)

2.2 Survey outcomes

Outcomes follow the Guide to DHS Statistics (Croft et al. 2018): de facto children aged 0–59 months, measured against the WHO 2006 standards, with flagged measurements excluded. This processing reproduces the official figures: 30.0% stunting, 8.9% severe stunting, 12.1% underweight and 3.3% wasting nationally, and all 31 region stunting values to within 0.05 percentage points of the DHS Program's published values.

2.3 Census

The census 10% sample provides the population to which the model is applied: every child aged 0–4 with their household's characteristics. Household and person records were linked for 99.98% of households, and council and ward names were matched to the 2022 boundary files for over 99.9% of records.

2.4 Harmonizing survey and census variables

A unit-level model needs covariates measured identically in survey and census. 26 household variables were constructed from both questionnaires and tested for comparability against a pre-specified rule. 18 passed and enter the model at child level, covering the mother's and household head's education, housing materials, water, sanitation, cooking fuel, electricity, livestock and household composition. The others are used only as council averages, where definitional differences matter less. Urban residence is excluded at child level because the survey frame uses the 2012 urban definition (32% urban against 38% in the 2022 census).

Mother's education. The census does not identify each child's mother, so the mother's education is approximated at household level, in the same way in both sources. In the survey, where the actual mother is known, the proxy matches her education for 94–97% of children.

2.5 Survey cluster locations

To protect respondents' confidentiality, DHS moves the published location of each survey cluster at random, by up to 2 km in urban and 5 km in rural areas (10 km for 1% of rural clusters), but keeps each cluster in its original district (Burgert et al. 2013). Assignment of clusters to councils is therefore reliable, but assignment to wards is not.

The displacement also means that satellite data read at the published point may describe the wrong place, for example open bush a few kilometres from the village that was surveyed. To correct for this, 1,000 possible true locations were simulated for each cluster, following the DHS displacement rules and staying within the cluster's council. Each possible location was weighted by the number of buildings there, because the surveyed households are more likely to be where buildings are, and the satellite values were averaged over all of them. The result is the best estimate of the conditions at the cluster's true, unpublished location.

3 Methodology

Seven estimators were built and compared on independent data. The headline estimates come from a unit-level model of child height-for-age, fitted to the survey and applied to every child in the census. This empirical best predictor (EBP) approach is widely used for poverty mapping by the World Bank and national statistics offices (Elbers, Lanjouw and Lanjouw 2003; Molina and Rao 2010; Corral et al. 2022).

3.1 Overview of the estimators

Small area estimators differ in three respects: the level at which the model is fitted (child, survey cluster or council), the data used to fit it and to predict from it, and whether a council's own survey data enter its estimate. Table 2 summarizes the seven estimators compared in this report, alongside the direct survey estimate. All are fitted to the same 5,516 measured children and draw on the same census and satellite data; they differ in how they combine them.

Table 2. The estimators compared

Estimator	Model level	Fitted to	Predicts for	Council's own survey data	Uncertainty
Main model (EBP)	Child	5,516 survey children: age, sex, 18 household variables, region, 4 council context variables	All 940,619 census children 0–4, averaged by area	Through a council random effect, shrunk to the model	Semi-parametric bootstrap, 200 replicates
FH - census, FH - geo, FH - census+geo	Council	189 council direct estimates and council census and/or satellite covariates	195 councils	Weighted average of direct and model prediction	Parametric bootstrap, 500 replicates
FH - hybrid	Council	As above, plus a machine-learning prediction as a covariate	195 councils	As above	As above
ML - unit	Child	5,516 survey children: 100 child, household and council features; gradient boosting (LightGBM) or regularized logistic/ridge regression	All census children 0–4, averaged by area	No (prediction only)	Cluster bootstrap, 50 replicates
ML - grid	Survey cluster	625 cluster means: 79 satellite and council census features; gradient boosting (LightGBM) or ridge regression	1.63 million populated ~500 m grid cells	No (prediction only)	Cluster bootstrap, 50 replicates
Direct estimate	—	Survey children in the council only	189 councils with survey data	Only source	Design-based, smoothed

FH = Fay–Herriot area-level model; ML = machine learning. Sections 3.3–3.8 describe each estimator.

3.2 How the data are used

The survey supplies the outcome and the relationships the models learn; the census supplies the population to which those relationships are applied, and area-level information that covers every council; the satellite data add information on settlement, infrastructure and land cover that neither survey nor census measures. The full report sets out which pieces of each source enter which estimator.

Two design choices govern how survey and census are combined. First, a variable enters a child-level model only if it is measured the same way in both sources (Section 2.4); otherwise predictions for census children would be biased. Variables that fail this test can still be used as council averages, where definitional differences largely cancel. Second, council aggregates from the census are weighted by the number of children under 5 in each household, so they describe the conditions young children live in rather than the average household.

3.3 Direct estimates

Direct estimates use only the survey children in each area. They are survey-weighted means, with variances linearized by cluster within the 61 sampling strata (region by urban and rural). They are the benchmark that any model must beat and the input to the area-level models. With one to three clusters in most councils, direct variances are unstable and are zero when no sampled child (or every sampled child) is stunted. They were therefore smoothed with a generalized variance function: the variance of a council proportion is taken as the pooled design effect $\times p(1 - p) / n_{\text{eff}}$, where p is the region prevalence and n_{eff} the council's effective (Kish) sample size.

3.4 The main model

Specification. The height-for-age z-score of child i in survey cluster c of council d is modelled as

$$\text{HAZ}_{icd} = x'_{icd}\beta + u_d + v_c + e_{icd}, \quad u_d \sim N(0, \sigma_u^2), \quad v_c \sim N(0, \sigma_v^2)$$

where β are fixed coefficients, u_d is a council effect, v_c a cluster effect and e_{icd} the child-level error. This is the nested-error model of Battese, Harter and Fuller (1988), extended with a cluster level because survey children are sampled in clusters. The covariates x fall into four groups (the full report lists all coefficients):

- **Child.** Age in completed years (0–4, as separate categories) and sex.
- **Household (18 variables).** The mother's and household head's education, housing quality, water, sanitation, energy, livestock and household composition, each taken from the child's own household: from the survey when the model is fitted, from the census when it is applied.
- **Region.** Fixed effects for the 31 regions. They capture differences in diet, agro-ecology and disease that the census does not measure, and absorb the survey's stratification by region.
- **Council context.** Four standardized council variables, one from the census and three from the satellite data. They were selected by LASSO (Tibshirani 1996) in the area-level model with census and satellite covariates (Section 3.7), and let the main model use information that exists only for areas.

Estimation. The model was fitted on the 5,516 children by restricted maximum likelihood (REML), the standard method for estimating the variance components of mixed models; unlike ordinary maximum likelihood, it allows for the regression coefficients being estimated from the same data, which avoids underestimating the variances. The model is fitted unweighted: the DuMouchel–Duncan test (DuMouchel and Duncan 1983) finds the sampling weights uninformative once region and the other covariates are included ($p = 0.12$). The variance components are 0.0088 for councils, 0.054 for clusters and 1.44 for children. Councils therefore differ very little once household conditions, region and context are accounted for: the council effect has a standard deviation of 0.09 z, small compared with the coefficients of individual covariates. This is why estimates for councils without survey data are almost as reliable as those for sampled councils, and why the model relies mainly on the census information rather than on each council's few survey children: the estimated council effect gives the council's own survey residuals a weight of about 0.1 on average (the rest is shrunk to zero), whereas the area-level models give the direct estimate a weight of about 0.3. Rare missing covariate values were replaced by the mean in each source.

3.5 From model to area estimates

The fitted model is applied to every child aged 0–4 in the census 10% sample, between 712 and 19,451 children per council (median 4,172). For census child i in council d :

1. The linear predictor $x'_i\hat{\beta} + \hat{u}_d$ combines the child's own age, sex and household characteristics, the region and council context, and the estimated council effect $\hat{u}_d$. For the 189 councils with survey data, $\hat{u}_d$ is the best linear unbiased predictor from that council's survey residuals; for the six councils without survey data it is zero.
2. The probability of stunting is the probability that HAZ falls below -2 given this prediction:

$$\hat{p}_i = \sum_q w_q \hat{F}(-2 - x'_i\hat{\beta} - \hat{u}_d - z_q\sqrt{V_d})$$

where $\hat{F}$ is the empirical distribution of the model's residuals, V_d the remaining uncertainty about the council effect, and (z_q, w_q) five Gauss–Hermite quadrature points that integrate over it. Severe stunting uses -3 ; the expected HAZ is the linear predictor itself. A census child's cluster is unknown, so the cluster effect is part of the residual.

3. Area estimates are averages over the census children in the area:

$$\hat{\theta}_d = (1/N_d) \sum_{i \in d} \hat{p}_i$$

and the number of stunted children is $\hat{\theta}_d \times N_d \times 10$. The same averaging gives district, region and ward estimates.

Residual distribution. The residuals have heavier tails than a normal distribution (excess kurtosis 2.1), so the empirical distribution is used. Within the survey it reproduces observed stunting more closely: the mean gap between predicted and observed stunting is 2.1 points across regions and 2.3 points across deciles of predicted HAZ, against 2.8 and 3.7 points with normal errors, and national stunting is reproduced (29.9% against 30.0%), whereas normal errors over-predict it by 1.7 points.

3.6 Uncertainty

Standard errors come from a semi-parametric bootstrap with 200 replicates, following González-Manteiga et al. (2008) and Molina and Rao (2010). In each replicate:

1. new council effects are drawn from $N(0, \hat{\sigma}_u^2)$, and a synthetic census population is generated by adding them and resampled residuals to each census child's linear predictor; the area shares of children below -2 and -3 are that replicate's true values;

2. a synthetic survey sample is generated in the same way for the 5,516 survey children, with cluster effects;
3. the coefficients and council effects are re-estimated from the synthetic survey, and the estimates are recomputed for every area;
4. the squared difference between estimate and true value is recorded.

The mean squared error (MSE) of each area is the average over replicates. Variance components are held at their REML values, so the bootstrap does not capture their uncertainty, and it assumes the model is correct; the split-sample test in Section 5.3 measures accuracy without that assumption. 95% confidence intervals for prevalence are computed on the logit scale. Reliability classes follow common statistical-office practice: **reliable** (CV below 20%), **use with caution** (CV 20–33%) and **unreliable** (CV above 33%).

3.7 Alternative 1–4: area-level Fay–Herriot models

Model. Area-level models work with council direct estimates rather than individual children (Fay and Herriot 1979). For council d the direct estimate y_d is the true value plus sampling error, and the true value follows a regression on council covariates z_d :

$$y_d = \theta_d + e_d, \quad e_d \sim N(0, \psi_d); \quad \theta_d = z_d' \beta + u_d, \quad u_d \sim N(0, \sigma_u^2)$$

The estimate is a weighted average of the direct estimate and the regression prediction, $\tilde{\theta}_d = \gamma_d y_d + (1 - \gamma_d) z_d' \hat{\beta}$, with weight $\gamma_d = \hat{\sigma}_u^2 / (\hat{\sigma}_u^2 + \psi_d)$: the noisier the direct estimate, the more weight goes to the prediction. Councils without survey data receive the regression prediction.

Covariates. The four variants differ only in the pool of candidate covariates: the 57 census council variables (FH-census), the 23 satellite features (FH-geo), both (FH-census+geo), or both plus a machine-learning prediction of mean HAZ from the satellite data (FH-hybrid). The prediction is cross-fitted by region, so a council's own survey data never enter its covariate. From each pool, covariates were selected by LASSO (Tibshirani 1996) weighted by the precision of the direct estimates.

Satellite covariates explain more of the variation between councils than census averages (57–61% against 44%) and remove the spatial autocorrelation that census covariates leave in the residuals. The direct estimate receives a weight of 0.28–0.36 on average.

3.8 Alternatives 5–6: machine-learning models

The two machine-learning estimators replace the statistical model with flexible prediction. They use no random effects, no direct estimates and no distributional assumptions: each indicator is predicted directly, and a council's estimate depends on its survey data only through the training sample. Each is built with two kinds of learner, gradient boosting (LightGBM) and a regularized linear model (logistic or ridge regression), and the better one is chosen by cross-validation for each indicator.

ML-unit (child level) works like the main model but without its structure. It is trained on the 5,516 survey children with 100 features: age, sex, the 18 household variables and the 80 council census and satellite variables. The trained model is applied to every census child and the predictions are averaged by area. Stunting and severe stunting are predicted as probabilities (classification), mean HAZ as a value (regression). Children are weighted by their survey weights.

ML-grid (cluster level) learns from places rather than children. It is trained on the 625 survey clusters, with 79 satellite and council census features corrected for the displacement of cluster coordinates, and predicts all 1.63 million populated ~500 m grid cells, which are averaged with weights equal to their number of children under 5.

Model choice. Two kinds of learner were compared for each estimator and indicator: gradient boosting (LightGBM; Ke et al. 2017), an ensemble of small decision trees that captures non-linear effects and interactions, and a regularized linear model (penalized logistic regression for child-level outcomes, ridge regression otherwise), which is simpler and harder to overfit. The better one was chosen by council-grouped cross-validation, so that no council's own data are used to predict it. Gradient boosting won in most cases; ridge regression won where the signal is weakest.

Uncertainty. Standard errors come from a cluster bootstrap within regions; they measure estimation variance only, so the split-sample test is the guide to their accuracy.

3.9 Model selection

All seven estimators were compared in a split-sample test on independent data (Section 5.3): each method is refitted on a random half of the survey clusters and scored against direct estimates from the other half. The main model was the most accurate for stunting and no method was significantly better for the other indicators, so it supplies all headline estimates.

3.10 Benchmarking and indicators reported for regions only

The headline estimates are not rescaled to the TDHS-MIS region figures, because region gaps are within sampling noise; a benchmarked version is included in the data package.

Underweight and wasting are reported for regions only, as published in the TDHS-MIS (Section 4.6):

- **Underweight** (low weight for age) was modelled in the same way as stunting with all four area-level models. In none of them did any of the 80 census and satellite variables help predict the differences in underweight between councils: the covariate selection kept none. A council estimate would then rest almost entirely on the council's 25 or so measured children, and would be little better than the imprecise direct survey estimate.
- **Wasting** (low weight for height) is acute malnutrition and affects 3.3% of children nationally, so a council's measured children include on average less than one wasted child, too few to model. Wasting also changes with the season and with short-term shocks such as drought or disease outbreaks, and the survey was carried out between February and July 2022, so council differences in a single survey round would partly reflect when each area was visited.

4 Results

Stunting is highest in the Southern Highlands (Iringa, Njombe) and Rukwa, while the largest numbers of stunted children are in the populous councils of the Lake Zone and the west.

Preview edition. This chapter shows summary maps in four classes and the priority councils in rank order. The full report has detailed maps of every indicator, ranked tables with estimates and confidence intervals, and the data package gives every council, district and ward estimate (Chapter 8).

4.1 Stunting prevalence by council

Council stunting ranges from 10.7% (Moshi MC) to 57.8% (Kilolo DC). 72 councils exceed 30% and 16 exceed 40%, all of the latter in Iringa, Njombe and Rukwa except Ngorongoro DC in Arusha (40.3%) (Figure 1, Table 3). The lowest prevalence is in Moshi MC, Moshi DC, Korogwe TC, Ubungo MC and Mtwara Mikindani MC. Within a region, councils typically differ by about 10 percentage points.




Figure 1. Stunting exceeds 40% in 16 councils, all but one in the Southern Highlands and Rukwa. Stunting prevalence among children under 5 by council, 2022, in four classes.

Source: SpacePilots small area estimates (main model) from TDHS-MIS 2022 and PHC 2022. Grey lines are region boundaries. The data package gives each council's estimate with its standard error, 95% confidence interval, CV and reliability class.

Table 3. Priority councils: the ten with the highest stunting prevalence and the ten with the most stunted children, 2022

Rank	Highest prevalence	Most stunted children
1	Kilolo DC (Iringa)	Geita DC (Geita)
2	Makete DC (Njombe)	Sumbawanga DC (Rukwa)
3	Mufindi DC (Iringa)	Kaliua DC (Tabora)
4	Njombe DC (Njombe)	Chato DC (Geita)
5	Ludewa DC (Njombe)	Nkasi DC (Rukwa)
6	Mafinga TC (Iringa)	Muleba DC (Kagera)
7	Wanging'ombe DC (Njombe)	Temeke MC (Dar es Salaam)
8	Njombe TC (Njombe)	Biharamulo DC (Kagera)
9	Iringa DC (Iringa)	Igunga DC (Tabora)
10	Nkasi DC (Rukwa)	Uyui DC (Tabora)

Councils in rank order. Bold: in the top ten on both measures. Estimates, confidence intervals and the ranking of all 195 councils are in the data package.

4.2 Number of stunted children

An estimated 2.7 million children under 5 are stunted. Prevalence is a poor guide to where they live: the highest-prevalence fifth of councils holds only 30% of stunted children, while the 40 councils with the largest numbers hold 42% (Figure 2). Only one council, Nkasi DC, is among the ten highest on both measures (Table 3).




Figure 2. 42% of the 2.7 million stunted children live in 40 of the 195 councils. Estimated number of stunted children under 5 by council, 2022, in four classes.

Number = prevalence × children under 5 from the 2022 census (10% sample × 10). The data package gives each council's number with its confidence interval.

4.3 Prevalence and numbers compared

Figure 3 plots each council's stunting prevalence against its number of stunted children. The two are related, but less closely than one might expect. Councils with higher prevalence tend to have more stunted children (rank correlation 0.62), but the number depends mainly on how many young children a council has: it correlates 0.87 with the number of children under 5 and only 0.47 with prevalence. The grey lines join points with the same number of children under 5, so a council's position along a line shows its prevalence and the line shows its size.




Figure 3. High prevalence and large numbers of stunted children often occur in different councils. Stunting prevalence and number of stunted children under 5 by council, 2022.

Dotted lines: national prevalence from the council estimates (28.8%) and median number of stunted children per council. Grey lines: equal numbers of children under 5. Source: SpacePilots small area estimates (main model).

Dividing councils at the national prevalence (28.8%) and the median number of stunted children (about 12,100) gives four groups:

- **Higher prevalence, more stunted children** (67 councils, 56% of all stunted children). Mostly large rural councils in the Lake Zone (Geita, Mwanza, Kagera, Simiyu), the west (Kigoma, Tabora) and Rukwa, led by Geita DC, Sumbawanga DC, Kaliua DC and Chato DC. Iringa DC, Mufindi DC and Kilolo DC are the only councils with both prevalence near or above 50% and more than 20,000 stunted children. These councils are priorities on either criterion.
- **Higher prevalence, fewer stunted children** (20 councils, 7%). Small councils, many in the Southern Highlands. Makete DC (57%), Njombe DC (53%), Ludewa DC (52%) and Mafinga TC (51%) have some of the highest prevalence in the country but fewer than 12,000 stunted children each.
- **Lower prevalence, more stunted children** (30 councils, 20%). Four of the five Dar es Salaam councils, Dodoma CC and populous district councils in Mara, Tanga, Mbeya and Shinyanga. Temeke MC (20%) and Dar es Salaam CC (15%) each have about 30,000 stunted children, more than Kilolo DC, the council with the highest prevalence.
- **Lower prevalence, fewer stunted children** (78 councils, 17%).

The choice of targeting criterion therefore matters. The 40 councils with the highest prevalence hold 31% of stunted children, while the 40 with the most stunted children hold 42%. Only 18 councils appear on both lists. Programmes aiming to reduce the national number of stunted children need to reach the populous councils in the second and third groups, not only the high-prevalence councils of the Southern Highlands.

4.4 Severe stunting

Nationally 8.5% of children under 5 are severely stunted, about 804,000 children. Council estimates range from 2.6% to 24.2% and follow the same geography as stunting. The five highest are Kilolo DC, Makete DC, Mufindi DC, Njombe DC and Ludewa DC.

4.5 Mean height-for-age

The national mean HAZ is -1.34. Mean HAZ reflects the whole distribution of child growth, not only children below a cut-off. Council means range from -2.21 (Kilolo DC) to -0.56 (Moshi MC). In the five worst councils, all in Iringa and Njombe, the average child is more than two standard deviations below the WHO median.

4.6 Districts and regions

District estimates come from the same model and are as precise as council estimates. Kilolo, Makete and Mufindi districts are highest; Moshi, Ubungu and Hai are lowest.

Aggregated to region, the council estimates agree with the official TDHS-MIS figures within sampling error everywhere (Table 4). The two differ by 2.5 points on average and never by more than 1.4 standard errors of the survey figure. Nationally the council estimates total 28.8% against the survey's 30.0%, a gap of about one standard error.

Table 4. Region estimates: official TDHS-MIS figures and aggregated council estimates, 2022

Region	Stunting, TDHS-MIS (95% CI $\pm$)	Stunting, councils aggregated	Underweight, TDHS-MIS	Wasting, TDHS-MIS	Children under 5
Iringa	56.9% (± 12.2)	52.6%	10.7%	1.9%	157,320
Njombe	50.4% (± 9.4)	50.4%	12.2%	2.1%	115,580
Rukwa	49.8% (± 7.5)	47.3%	19.4%	8.3%	275,650
Geita	38.6% (± 13.2)	34.5%	10.3%	3.3%	534,010
Ruvuma	35.6% (± 9.3)	34.4%	12.2%	2.8%	263,780
Kagera	34.3% (± 12.4)	35.4%	15.4%	3.7%	488,370
Simiyu	33.2% (± 6.7)	29.5%	11.4%	1.7%	397,600
Tabora	33.1% (± 7.7)	29.3%	12.2%	3.0%	609,230
Katavi	32.2% (± 7.6)	33.6%	14.0%	1.8%	219,630
Manyara	32.0% (± 7.0)	30.9%	18.9%	4.7%	311,340
Songwe	31.9% (± 8.7)	27.9%	8.1%	0.0%	220,830
Mbeya	31.5% (± 10.9)	24.8%	1.4%	0.6%	337,830
Dodoma	30.7% (± 9.7)	30.8%	11.0%	1.4%	452,580
Arusha	30.7% (± 10.5)	26.6%	20.1%	4.5%	362,410
Morogoro	30.6% (± 6.7)	26.8%	10.0%	3.9%	454,300
Kaskazini Unga	29.0% (± 8.3)	23.3%	21.0%	10.0%	37,340
Mwanza	27.9% (± 6.4)	31.5%	15.1%	1.8%	588,130
Shinyanga	27.5% (± 6.4)	27.2%	8.6%	1.3%	377,680
Kigoma	27.1% (± 6.8)	31.7%	12.2%	2.5%	394,480
Singida	25.7% (± 7.7)	28.7%	11.1%	2.3%	337,980
Tanga	24.1% (± 6.3)	23.6%	11.2%	5.6%	374,070
Mara	23.4% (± 6.7)	25.0%	7.6%	3.4%	391,750
Mtwara	22.3% (± 8.1)	21.6%	6.4%	2.1%	199,410
Kaskazini Pemba	21.8% (± 6.5)	21.8%	15.5%	7.7%	44,300
Lindi	21.3% (± 12.0)	22.5%	10.0%	2.1%	148,510
Kilimanjaro	20.4% (± 9.0)	16.3%	16.4%	2.9%	225,410
Pwani	20.2% (± 13.1)	21.5%	12.1%	6.2%	268,370
Dar es Salaam	18.4% (± 6.9)	16.6%	9.7%	4.3%	623,450
Kusini Pemba	17.3% (± 7.8)	19.9%	11.0%	3.3%	42,850
Kusini Unga	16.7% (± 6.8)	15.9%	6.1%	3.5%	26,050
Mjini Magharibi	12.8% (± 6.5)	17.1%	15.2%	10.5%	125,920

Regions sorted by TDHS-MIS stunting prevalence. Underweight and wasting are not estimated below region level.

4.7 Ward estimates (experimental)

Experimental. Ward results show where to look more closely within a council. They should only be used for allocating resources with great caution.

The main model, averaged over each ward's census children, gives stunting estimates for 4,336 wards. They range from 8.8% to 63.2%. 352 wards exceed 40% stunting; 302 of them are in Iringa, Njombe and Rukwa, and the others are mainly in Ruvuma, Arusha and Kagera. Within a council, ward estimates typically span about 8 percentage points, and up to 15–16 points in councils such as Monduli DC, Ngorongoro DC and Morogoro MC. Figure 4 shows one region as a sample.

These estimates have important limits. GPS displacement means survey data cannot be assigned to wards, so ward variation within a council comes only from census household characteristics. There is no ward random effect, so a ward that differs from what its households predict cannot be detected. Ward CVs (median 20%) assume the model is correct and understate the true uncertainty.




Figure 4. Sample of the ward-level product: Arusha region. Experimental stunting estimates by ward, 2022, in four classes.

Ward variation comes only from census household covariates; there is no ward-level random effect. The data package covers all 4,336 wards.

5 Quality and validation

88% of council stunting estimates are reliable, and the main model was the most accurate of seven methods in an independent test, about 3.5 times as accurate as direct survey estimates.

5.1 Precision

Table 5. Precision of the estimates by level and indicator

Level	Indicator	Median CV	Median 95% CI width	Reliable	Use with caution	Unreliable
Council (195)	Stunting	16%	16.3 pp	171	24	0
Council (195)	Severe stunting	21%	6.4 pp	78	115	2
Council (195)	Mean HAZ	—	0.57 z	—	—	—
District (150)	Stunting	15%	15.7 pp	137	13	0
District (150)	Severe stunting	20%	6.2 pp	77	72	1
District (150)	Mean HAZ	—	0.53 z	—	—	—

Reliable: CV below 20%; use with caution: CV 20–33%; unreliable: CV above 33%. pp = percentage points.

No council or district stunting estimate is unreliable (CV above 33%); for severe stunting, two council estimates and one district estimate are. CVs are not reported for mean HAZ, because a z-score has no natural zero.

Councils marked “use with caution”. The 24 councils in this class all have low prevalence (11–23%); they include Moshi MC, Moshi DC, Hai DC, Rombo DC and Siha DC in Kilimanjaro, Ubungo MC, Tanga CC, Morogoro MC and several councils in Lindi and Mtwara. A CV is relative to the estimate, so low-prevalence councils reach 20% even though their absolute uncertainty (± 6 –10 points) is no larger than other councils'. Severe stunting is rarer, so more of its estimates fall in this class for the same reason.

Councils with no survey data. Korogwe TC, Pangani DC, Bagamoyo DC, Rufiji DC, Nanyamba TC and Nzega TC had no sampled households. Their estimates rest entirely on census and satellite information and have CVs of 16–27%.

5.2 Gain over direct survey estimates

The model reduces the median CV from 35% to 16%. Without it, 88% of councils would fail the reliability threshold. The model pulls imprecise direct estimates towards values supported by the census and satellite data, most strongly where the survey sample is smallest.

5.3 Independent accuracy test

Model-based CVs assume the model is correct, so accuracy was also measured directly. In each of 20 replicates, the survey clusters in each council with at least two clusters (168 councils) were split at random into two halves. Every method was refitted on one half and its council estimates compared with direct estimates from the other, independent half. The other half's sampling variance was subtracted from the squared differences to give each method's mean squared error (Figure 5, Table 6).




Figure 5. The main model's mean squared error is 92% lower than that of direct estimates. Split-sample validation error for council stunting, by method.

Bars: mean squared error relative to the direct estimate, 20 replicates; whiskers: 95% intervals.

Table 6. Split-sample validation results, council level (estimates fitted on half the survey)

Method	Stunting RMSE (pp)	MSE relative to direct	Paired t vs main model, stunting	Paired t vs main model, mean HAZ
Main model (unit-level census EBP)	4.1	0.08	—	—
ML -grid (satellite grid)	4.9	0.12	2.6	-0.8
FH -hybrid	5.7	0.16	3.9	-0.5
ML -unit (child level)	6.6	0.22	9.3	4.0
FH -geo	6.6	0.22	8.1	1.7
FH -census+geo	7.1	0.25	8.9	4.4
FH -census	7.3	0.27	10.5	6.5
Direct survey estimate	14.1	1.00	29.1	27.3

Positive t means a larger error than the main model; differences are paired within the 20 replicates.

The main model is significantly the most accurate for stunting. For mean HAZ it is statistically tied with ML-grid and FH-hybrid. Because it is best for the primary indicator and not significantly worse for any other, it is used for all three indicators. A mean squared error 8% of the direct estimate's corresponds to a root mean squared error 29% as large (4.1 against 14.1 percentage points), a 3.5-fold reduction. The full report gives the validation design and replicate-level results.

5.4 Consistency with official figures

Region aggregates of the council estimates are within 1.4 standard errors of the official TDHS-MIS figure in every region (Table 4). The spread of the 31 differences is smaller than sampling noise alone would produce, so there is no sign of systematic bias at region level.

6 Discussion and limitations

The estimates are precise enough to rank and target councils, but they inherit the limits of the predictors available and describe a single year.

6.1 What drives the results

The Southern Highlands, among the country's main food-producing areas, have the highest prevalence, so stunting does not simply follow food production. This is consistent with the official region figures and is reproduced by every model tested, including those based only on satellite data. Household characteristics explain part of the pattern: the model finds lower HAZ for children in households that cook with firewood, use surface water or are more crowded, and higher HAZ where the mother has secondary education or the household has electricity or owns cattle. Region effects capture much of what remains, such as diet and agro-ecological differences that the census does not measure.

6.2 Why the household-level model outperforms

The main model uses each census child's own household characteristics, so council estimates reflect each council's actual mix of households, not only council averages. It also keeps the council effect, so councils with survey data are not reduced to their covariate prediction. The satellite-based ML-grid model was the strongest alternative, which shows that imagery carries real signal for child growth. With 5,516 children, however, flexible machine-learning models learned less than a structured statistical model.

6.3 Limitations

- **Predictors.** Household characteristics explain little of an individual child's growth. Key determinants are not in the census: maternal height, birth weight, diet and illness. Maternal education is included, but through a household proxy because the census does not identify each child's mother; the proxy matches the actual mother for 97% of survey children. Given the household's housing, water and energy, education adds little further information. A council that is unusual for unmeasured reasons is pulled towards what its covariates and region predict.
- **Model-based uncertainty.** The standard errors, confidence intervals and CVs are calculated by simulating new data from the fitted model itself (Section 3.6). They show how uncertain the estimates are if the model is right, and they treat the model's variance parameters as known. If the model misses something, for example a council that is unusual in ways the census and satellite data do not capture, the actual error can be larger than the CV suggests. The split-sample test (Section 5.3) does not rely on the model: it compares the estimates with survey data that were not used to produce them, so it is the better guide to real accuracy.
- **Validation advantage.** The main model's contextual covariates and residual distribution were chosen on the full sample, which slightly favours it in the validation. Its lead over the runner-up ($t = 2.6$) is unlikely to be explained by this alone.
- **Boundaries.** The survey frame predates the 2022 boundaries. Four clusters were reassigned, and small misallocations near changed boundaries are possible.
- **Wards.** Survey data cannot be assigned to wards because of GPS displacement, so ward estimates are experimental (Section 4.7).

7 Conclusions and recommendations

Council-level stunting estimates for Tanzania can be produced with useful precision by combining the 2022 TDHS-MIS with the 2022 census and satellite data. The results show large differences between councils that region figures hide, and they identify two distinct sets of priority councils: those with the highest rates and those with the most stunted children.

7.1 Conclusions

1. Stunting varies more than fivefold across councils, from 11% to 58%. 72 councils exceed 30%.
2. The highest prevalence is concentrated in Iringa, Njombe and Rukwa; the largest numbers are in populous councils of the Lake Zone and the west, such as Geita, Sumbawanga, Kaliua and Chato.
3. 171 of 195 council stunting estimates are reliable, against 12% of councils with direct survey estimates alone.
4. The unit-level census model was the most accurate of seven methods in an independent test, with an estimation error about 3.5 times smaller than that of direct survey estimates.
5. The estimates agree with the official region figures within sampling error.

7.2 Recommendations for users

1. **Targeting.** Use both prevalence and numbers. High-prevalence councils need intensive programmes; high-burden councils account for most stunted children.
2. **Interpretation.** Always report confidence intervals and reliability classes alongside the estimates. Treat councils whose intervals overlap as not clearly different.
3. **Official figures.** Use the TDHS-MIS figures for national and region reporting, and these estimates for comparisons below region level.
4. **Wards.** Use ward figures for allocating resources only with great caution.

References

- Battese, G. E., Harter, R. M. and Fuller, W. A. (1988). An error-components model for prediction of county crop areas using survey and satellite data. *Journal of the American Statistical Association*, 83(401), 28–36.
- Burgert, C. R., Colston, J., Roy, T. and Zachary, B. (2013). *Geographic Displacement Procedure and Georeferenced Data Release Policy for the Demographic and Health Surveys*. DHS Spatial Analysis Reports No. 7. Calverton, MD: ICF International.
- Corral, P., Molina, I., Cojocar, A. and Segovia, S. (2022). *Guidelines to Small Area Estimation for Poverty Mapping*. Washington, DC: World Bank.
- Croft, T. N., Marshall, A. M. J., Allen, C. K. et al. (2018). *Guide to DHS Statistics*. Rockville, MD: ICF.
- DuMouchel, W. H. and Duncan, G. J. (1983). Using sample survey weights in multiple regression analyses of stratified samples. *Journal of the American Statistical Association*, 78(383), 535–543.
- Elbers, C., Lanjouw, J. O. and Lanjouw, P. (2003). Micro-level estimation of poverty and inequality. *Econometrica*, 71(1), 355–364.
- Fay, R. E. and Herriot, R. A. (1979). Estimates of income for small places: an application of James–Stein procedures to census data. *Journal of the American Statistical Association*, 74(366), 269–277.
- Fujii, T. (2010). Micro-level estimation of child undernutrition indicators in Cambodia. *World Bank Economic Review*, 24(3), 520–553.
- González-Manteiga, W., Lombardía, M. J., Molina, I., Morales, D. and Santamaría, L. (2008). Bootstrap mean squared error of a small-area EBLUP. *Journal of Statistical Computation and Simulation*, 78(5), 443–462.
- Ke, G., Meng, Q., Finley, T. et al. (2017). LightGBM: a highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems*, 30.
- Ministry of Health (Tanzania Mainland), Ministry of Health (Zanzibar), National Bureau of Statistics, Office of the Chief Government Statistician and ICF (2023). *Tanzania Demographic and Health Survey and Malaria Indicator Survey 2022: Final Report*. Dodoma and Rockville, MD.
- Molina, I. and Rao, J. N. K. (2010). Small area estimation of poverty indicators. *Canadian Journal of Statistics*, 38(3), 369–385.
- National Bureau of Statistics and Office of the Chief Government Statistician (2022). *2022 Population and Housing Census*. Dodoma and Zanzibar.
- Osgood-Zimmerman, A., Millea, A. I., Stubbs, R. W. et al. (2018). Mapping child growth failure in Africa between 2000 and 2015. *Nature*, 555(7694), 41–47.
- Rao, J. N. K. and Molina, I. (2015). *Small Area Estimation*. 2nd edition. Hoboken, NJ: Wiley.
- Simler, K. R. (2006). *Nutrition Mapping in Tanzania: An Exploratory Analysis*. FCND Discussion Paper 204. Washington, DC: International Food Policy Research Institute.
- Sohnesen, T. P., Ambel, A. A., Fisker, P., Andrews, C. and Khan, Q. (2017). Small area estimation of child undernutrition in Ethiopian woredas. *PLoS ONE*, 12(4), e0175445.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58(1), 267–288.
- WHO Multicentre Growth Reference Study Group (2006). *WHO Child Growth Standards: Length/Height-for-Age, Weight-for-Age, Weight-for-Length, Weight-for-Height and Body Mass Index-for-Age*. Geneva: World Health Organization.
- Yilema, S. A., Shiferaw, Y. A., Zewotir, T. and Muluneh, E. K. (2022). Small area estimation of undernutrition under age five children based on spatial-temporal models. *The Scientific World Journal*, 2022, 6882047.

Data sources

TDHS-MIS 2022 recode, GPS and covariate files (The DHS Program); published region values from the DHS Program API; PHC 2022 10% sample and boundaries (NBS); Google Open Buildings 2.5D Temporal, Landsat NDVI and Google satellite embeddings (Google Earth Engine exports).

8 The full report and data package

This preview is an abridged edition. The full report and the data package give the estimates for every area, the complete set of maps and the technical documentation needed for review.

8.1 Data package

Table 7. Contents of the data package

Component	Coverage	Contents
Council estimates	195 councils	Stunting, severe stunting and mean height-for-age, each with standard error, 95% confidence interval, CV and reliability class; number of stunted and severely stunted children; children under 5; survey sample sizes and direct estimates
District estimates	150 districts	The same indicators and quality measures
Region estimates	31 regions	Official TDHS-MIS figures for stunting, severe stunting, underweight and wasting, council aggregates and consistency checks
Ward estimates (experimental)	4,336 wards	Stunting prevalence and number of stunted children
Alternative estimates	Councils and districts	A version benchmarked to the official region figures, and the six alternative estimators for sensitivity analysis
Maps and charts	All levels	High-resolution maps of every indicator and all charts in the full report

Tables are supplied as CSV files that join to the 2022 census boundaries.

8.2 What the full report adds

- **Results.** Detailed council maps (seven classes) of stunting, the number of stunted children, severe stunting, mean height-for-age and the reliability of each estimate; ranked tables with estimates, confidence intervals and numbers affected; the full ward map.
- **Methods.** How each data source enters each estimator; the full specification of the main model, with all coefficients and variance components; the selected covariates and configurations of the six alternatives.
- **Quality.** The comparison of model and direct estimates for every council with survey data.
- **Technical appendices.** Data preparation; the main model's technical specification; detailed results of the alternative models; the validation design and the contribution of the education data; agreement between models; benchmarking to region figures; exploratory ~500 m grid predictions; and a guide to the data files.

For the full report and data package, contact SpacePilots at info@spacepilots.io.