Data DENT

Data for Decisions in Nutrition

TECHNICAL NOTE

Estimating nutrition intervention coverage for small areas using Bayesian spatial modeling



Project Note

DataDENT (Data for Decisions in Nutrition, www.datadent.org) aims to transform the availability and use of nutrition data by addressing gaps in nutrition measurement and advocating for stronger nutrition data systems. This work was carried out by the following DataDENT partners: Johns Hopkins Bloomberg School of Public Health (Hallie Eilerts-Spinelli, Nadia Akseer, Abdoulaye Maiga) with collaboration from International Food Policy Research Institute (IFPRI). This work is part of a broader set of activities focused on improving availability, quality and analysis of nutrition intervention coverage data. The work was funded by the Gates Foundation. The findings and conclusions contained within are those of the authors and do not necessarily reflect positions or policies of the Gates Foundation.

Suggested citation: DataDENT. 2026. Technical Note: Estimating nutrition intervention coverage for small areas using Bayesian spatial modeling. Johns Hopkins Bloomberg School of Public Health.

Introduction

As countries work to target nutrition and health interventions to populations with the greatest need, they need reliable subnational estimates of intervention coverage. Most national surveys produce estimates at national or first administrative levels (e.g. region, state) and may not provide sufficiently precise estimates for smaller geographic units (e.g. district, local area). This limitation constrains the use of data by governments and partners to guide implementation and progress tracking.

This technical note presents a flexible step-wise approach for using Bayesian spatial modeling to generate more reliable district- and local-level estimates of intervention coverage from household surveys. It accompanies the DataDENT technical note *Measuring who is being reached with multisector nutrition interventions: a methodol*ogical[1] framework for co-coverage & composite coverage analysis.

The primary audience for this technical note is data analysts and researchers. A strong foundation[1] in frequentist and Bayesian statistics is necessary to understand and apply the recommended methods.

[1]Access at https://datadent.org/resource/coverage-analysis-framework/

Key definitions

Establishing common definitions of key concepts is important for understanding the subsequent sections of this technical note.

Nutrition interventions

Services or programs implemented at scale to improve growth, health, development, and survival outcomes. Multisector nutrition interventions may be delivered through health systems, community platforms, schools, social protection programs, or other contexts. Common health sector nutrition interventions include micronutrient supplementation during pregnancy, growth monitoring and promotion, infant and young child feeding counseling, screening and treatment of acute malnutrition. Food fortification, WASH interventions, cash transfers, and school meals are examples of nutrition interventions implemented by other sectors.

Small area

In survey statistics, a small area refers to a geographic domain or population subgroup for which available data are insufficient to produce reliable estimates using conventional direct estimation methods. Small areas are not necessarily small in terms of population size; rather, they are areas where the number of observations is too limited to generate precise estimates. (e.g. districts, municipalities, subdistricts, catchment areas, demographic subgroup).

Intervention coverage

In the context of nutrition, intervention coverage refers to the proportion of individuals, households, or communities eligible for a specific nutrition service or interventions who receive the service or intervention. Coverage should always be expressed as a percentage or a fraction. Intervention coverage indicators are used to monitor progress, assess equity, and guide resource allocation and program planning.

Bayesian spatial modeling

Statistical approach that improves estimation by combining Bayesian inference with information about the geographic relationships between areas.

In Bayesian framework unknown quantities are treated as random variables and represented by probability distributions. Information from observed data is combined with prior assumptions or knowledge to generate posterior distributions that summarize the most likely values of model parameters and their uncertainty.

Spatial modeling extends this framework by accounting for the tendency of neighboring geographic areas to share similar characteristics. Through spatial smoothing, information is borrowed from nearby areas and relevant covariates to improve estimates in locations where data are sparse or unavailable.

Why use Bayesian small area estimation?

Sample sizes for national population-based household surveys including Demographic and Health Surveys (DHS), Multiple Indicator Cluster Surveys (MICS) and national nutrition surveys are often too small to generate precise estimates for districts or other local administrative units. Some interventions, such as treatment of childhood diarrhoea with oral rehydration solution (ORS) and zinc only apply to children who recently experienced diarrhoea, may have very few observations in some districts. This results in unstable estimates or no estimates at all. Even interventions that target large population groups, such as preventative vitamin A supplementation in children or iron-containing supplements during pregnancy, can have district coverage estimates with wide uncertainty intervals that limit their usefulness.

Bayesian small area estimation (SAE) addresses these limitations by borrowing strength from neighbouring geographic areas and relevant covariates to produce more stable and reliable local estimates. The approach can generate estimates for districts with sparse or even no observations while preserving underlying geographic patterns and providing a rigorous quantification of uncertainty.

Although Bayesian SAE has broad applications across health, nutrition, and other development sectors, this technical note focuses specifically on estimating the coverage of nutrition interventions from nationally-representative household survey data.

Selecting a Bayesian method for SAE of interventon coverage indicators

A range of Bayesian methods have been developed for SAE, including unit-level and area-level models, as well as approaches implemented using Markov Chain Monte Carlo (MCMC) or Integrated Nested Laplace Approximation (INLA). The choice of method depends on the characteristics of the available data, the outcome of interest, and the objectives of the analysis.

In this technical note, we apply Gao and Wakefield's spatial variance-smoothing area level model (Gao & Wakefield, 2023). This model accounts for heterogeneity in area-level variance and captures the dependence between mean and variance quantities for more accurate uncertainty propagation. The model is particularly well-suited for when the primary quantity of interest is a proportion generated from a weighted sum of binary responses (i.e., yes/no), and thus the mean and sampling variance are often strongly correlated (i.e., the variance changes as the mean changes). Evaluations of the Gao and Wakefield model indicate that it provides improved inference over mean smoothing models and may yield a more accurate representation of estimator uncertainty.

The model can be applied at the area-level or the unit-level. For large household survey data, an area-level model is better suited than a unit-level model because 1) area model better accounts for survey design effects from clustering and weighting, and 2) unit-level models require information on population composition for covariates to make post-stratification adjustments of modeled estimates; this need for population-level information is an additional constraint for low- and middle-income settings.

We considered the fully Bayesian nature of this approach (using MCMC methods) an additional advantage over methods relying on INLA, which have been widely implemented for SAE in the statistical software package, SUMMER (Li et al., 2025). INLA was developed as an alternative for fitting latent Gaussian models (Rue et al., 2009), achieving an accurate approximation of the posterior distribution for quantities of interest through deterministic methods of numerical integration (Martino & Riebler, 2019). INLA has been shown to produce similar results to fully Bayesian models in SAE for epidemiological quantities. However, there is evidence that INLA results in less reliable estimates when data is sparse or the outcome of interest is very rare (De Smedt et al., 2015). Additionally, because MCMC methods yield full posterior samples, they allow for estimation of joint distributions and complex functions of the model parameters, providing a more flexible and coherent treatment of uncertainty.

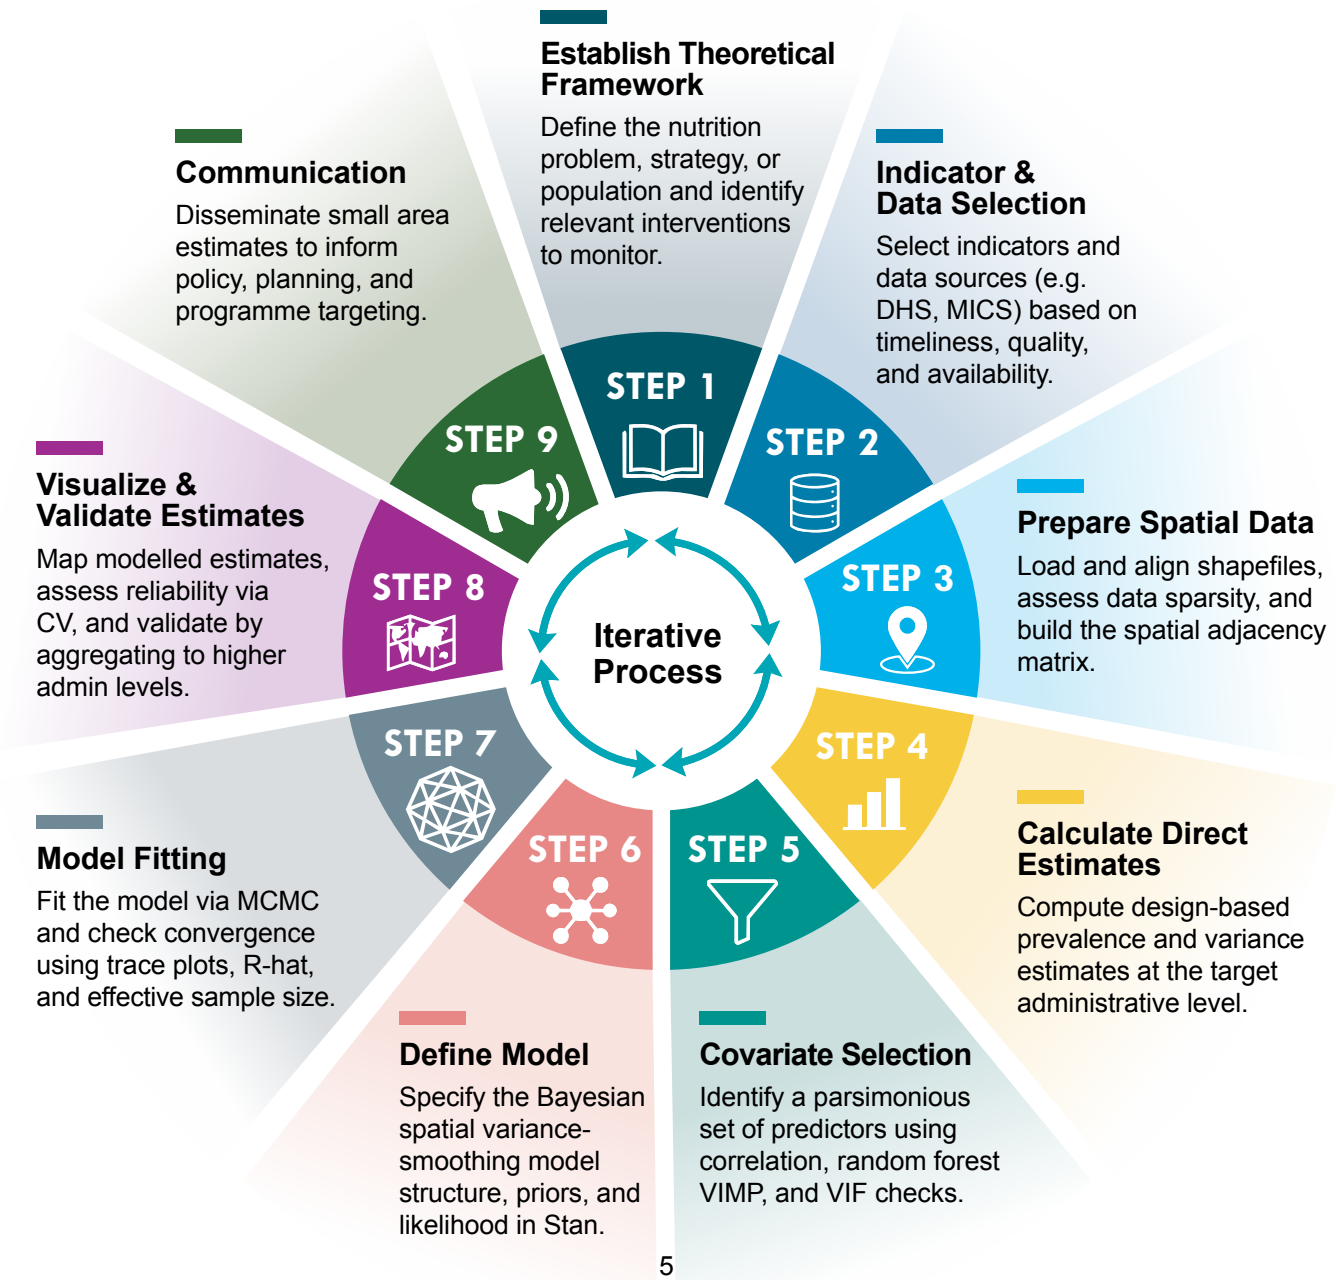
Overview of the step-wise approach

The Bayesian SAE approach presented in this technical note consists of nine sequential steps **(Figure 1)**. The process begins by establishing a conceptual framework and selecting appropriate indicators and data sources, followed by preparation of the spatial data and calculation of direct survey estimates. A Bayesian spatial model is then specified and fitted using carefully selected covariates, after which the resulting estimates are visualized, validated, and communicated for policy and programme decision-making.

Although presented sequentially, the framework is iterative, with model specification, covariate selection, and validation refined as needed to achieve reliable and robust SAE.

In this technical note, we demonstrate the framework using nutrition and health intervention coverage indicators from the Bangladesh DHS 2017 and 2022. Household survey microdata and georeferenced cluster coordinates were obtained from the DHS Program website (https://dhsprogram.com/data/available-datasets.cfm) and used throughout the worked examples.

Figure 1: Overview of stepwise approach to Bayesian SAE



Guiding principles



Use high-quality, representative survey data and account for the complex survey design when calculating direct estimates.



Select parsimonious and theoretically relevant covariates that improve model performance without overfitting.



Incorporate spatial dependence so neighbouring areas appropriately inform estimates where survey data are sparse.



Assess model convergence and estimate reliability using appropriate Bayesian diagnostics and validation procedures.



Communicate both estimates and their uncertainty, ensuring that maps and tables distinguish reliable estimates from those with high uncertainty.



Interpret modelled estimates alongside direct survey estimates, recognizing that small area estimates complement rather than replace design-based survey estimates.

STEP 1: ESTABLISH THE PURPOSE / THEORETICAL FRAMEWORK



The first step in generating small area estimates of intervention coverage from household surveys is to identify the purpose of the estimates and, if needed, establish the theoretical framework to guide the analysis (e.g. maternal-child continuum of care). Relevant interventions can be identified by reviewing policies, strategies, programme documents, and monitoring frameworks as well as stakeholder engagement.

Establishing a clear purpose and theoretical framework ensures that the small area estimates are policy-relevant and actionable. The small area estimates can be used to answer policy or programme questions. They can also be used as inputs for analyses including co-coverage or composite coverage.[2] The purpose will determine the geographic unit for the estimates (e.g. district, municipality, health facility catchment area).

STEP 2: INDICATOR AND DATA SELECTION



The next step is to identify intervention coverage indicators and nationally representative household survey data sources. Data are typically obtained from DHS, MICS, national nutrition surveys, or other surveys that collect intervention coverage data.

Survey datasets must be georeferenced to support SAE. Most recent DHS surveys include anonymized geographic coordinates for sampled clusters, enabling observations to be linked to administrative areas and incorporated into spatial models.

[2]DataDENT 2026. Measuring who is being reached with multisector nutrition interventions: a methodological framework for co-coverage & composite coverage analysis. Access at https://datadent.org/resource/coverage-analysis-framework/.

The suitability of each dataset should also be assessed based on its timeliness, data quality, representativeness, sample size, and geographic coverage.

Examples in this technical note use datasets from two rounds of the Bangladesh DHS – 2017 and 2022. Datasets are geocoded with GPS readings for most clusters accurate to less than 15 meters. Sample size in 2017 was 20,250 households and 30,018 households in 2022. Key populations include households, women aged 15-49 years and children aged 0-59 months. Per design, the surveys provide representative estimates at the national level, for urban and rural areas, and across eight administrative divisions (Barishal, Chattogram, Dhaka, Khulna, Mymensingh, Rajshahi, Rangpur, and Sylhet).

STEP 3: PREPARE SPATIAL DATA

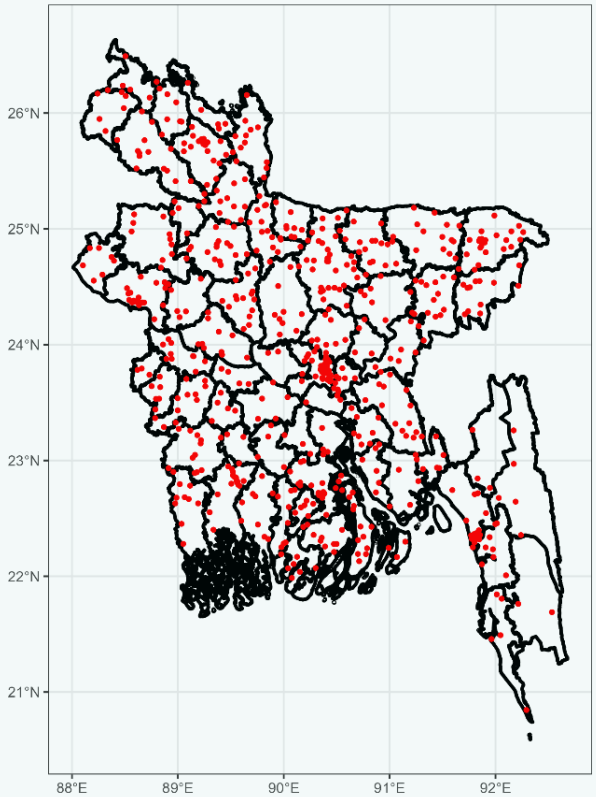
To generate small area estimates, shapefiles must be obtained for the administrative areas at which estimates are being produced. A shapefile stores geographic location data and associated attributes. Despite its name, a shapefile usually consists of several files bundled together in a folder. These include the core file (.shp) that stores geometric shapes and coordinates, the shape index file (.shx) that allows the software to read the former more easily, the database file (.dbf) that contains descriptive attributes about each shape such as the name, and the projection file(.prj) that defines the coordinate reference system.

Shapefiles can be loaded into Rstudio using the sf package (https://r-spatial.github.io/sf/). This is required for both the administrative boundary shapefile, which defines the areas at which estimates are desired, and the point-level shapefile containing the coordinates of the geolocated indicator data. It is important to ensure that both shapefiles are operating under the same Coordinate Reference System (CRS): WGS 84 is a common choice, but any CRS can be used as long as it is applied consistently to all shapefiles.

Prior to modelling, it is helpful to assess the distribution of raw data points across administrative areas. Assessment for data sparsity helps identify areas where estimates may be less reliable due to limited observations. First, tabulate the number of observations per administrative area. Then, visualize the spatial distribution of data clusters by overlaying point locations onto the shapefile. If multiple indicators are being modelled, data sparsity should be examined separately for each indicator; availability of data points may vary for indicators in the same dataset.

For Bangladesh we obtained shapefiles at the subnational administrative 2 level (i.e. “district”) from the Humanitarian Data Exchange (HDX, 2025). The plot **(Figure 2)** shows the administrative 2 level boundaries overlaid with survey cluster locations for the 2022 Bangladesh DHS.

Figure 2: Administrative 2 level boundaries overlaid with survey cluster locations (Bangladesh DHS 2022)

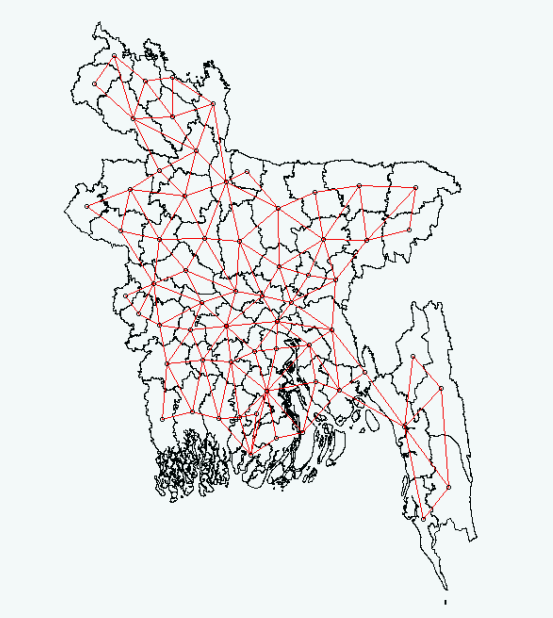


A spatial adjacency matrix must be prepared as an input to the SAE model. To do this, define which districts are considered neighbors. The spdep package (https://cran.r-project.org/web/packages/spdep/index.html) can be used to construct a neighbors list from the polygon boundaries in the shapefile. The Queen contiguity criterion, under which any two areas sharing at least one boundary point are considered neighbors, is generally the most appropriate choice for modelling spatially continuous outcomes such as nutrition intervention coverage.

Once the neighbors list is defined, it can be converted into a binary adjacency matrix with values of 1 or 0 indicating where two areas are neighbors or not, respectively. The matrix is then scaled for use as the spatial structure component of the small area model.

In **Figure 3** we plotted the neighbor links, as defined using Queen contiguity, for administrative 2 level districts in Bangladesh.

Figure 3: Neighbor links for administrative 2 level districts (Bangladesh DHS 2022)



STEP 4: CALCULATE DIRECT ESTIMATES

The primary inputs to the spatial variance-smoothing area level model are the direct estimates of prevalence and variance for the intervention coverage indicator at the selected administrative level of interest. These estimates must be calculated prior to modelling, accounting for survey design features including stratification, clustering, and individual sampling weights. The direct estimates will also serve as a benchmark against which modelled estimates are compared in Step 8.

The DHS Program GitHub provides code to construct binary indicator variables from the survey microdata (DHS Program, 2025). Once the indicator is defined, prevalence and variance can be estimated at the desired administrative level using a package such as *srvyr* (https://cran.r-project.org/web/packages/srvyr/index.html), which allows for design-based estimation, incorporating the stratification, clustering, and weighting when the survey design object is correctly specified.

Where data are sparse, some small areas may contain no observations at all. In other cases, the small area may contain sufficient observations to calculate prevalence but not variance. This occurs when either (i) there is only one cluster within the intersection of the small area and the larger sampling stratum, or (ii) multiple clusters are present in the small area-stratum, but all share identical weighted prevalences. To address these data gaps, the microdata can be augmented to enable calculation of design-based variance in any area where observations are available.

To augment the microdata, calculate the weighted average of cluster-level prevalences and variances within the stratum. Introduce this quantity as a “phantom cluster” in the small area where variance calculation was previously impossible. This approach supplements the small area data with the weighted mean of the surrounding stratum, providing the variance calculation with sufficient information to proceed. The result is a stabilized estimate for the small area that remains consistent with the broader spatial pattern, rather than an artificially inflated or deflated value.

Confidence intervals for the direct estimates can be constructed from the design-based variance estimates by taking their square root to obtain standard errors and applying a z-score approach to derive 95% intervals. This assumes the estimates are approximately normally distributed. These intervals will later serve as a benchmark for comparing the uncertainty of the modelled estimates.

STEP 5: COVARIATE SELECTION

The goal of covariate selection is to identify a parsimonious set of covariates that collectively explain the most variation in the indicator, without overfitting the model. A candidate list of covariates should first be assembled based on data availability and a priori knowledge of the factors expected to influence the outcomes. This list is then systematically narrowed through a series of assessments.

Candidate covariates for the model could be selected based on several considerations. Ideally the selected covariates are commonly collected so, for a given indicator, the model can be applied consistently across different surveys and contexts. Second, covariate selection should reflect the underlying pathways through which factors influence intervention coverage and nutrition outcomes. Rather than selecting covariates on an indicator-by-indicator basis, a broad set of covariates from multiple domains is preferable. Those operating more distally (e.g. socioeconomic, environmental, and demographic factors) tend to be the most predictive, as they capture structural conditions that shape both care-seeking behavior and uptake of nutrition interventions.

Covariates should be calculated at the small area level. For DHS data, the code available at the DHS Program GitHub (https://github.com/dhsprogram) allows for the calculation of a variety of covariates. Covariates from any source can be incorporated, provided they are available at the same administrative level used for estimation.

To narrow down the candidate covariates, first, Pearson's correlation coefficient (r) is calculated between all pairs of candidate covariates. Where two covariates are highly collinear ($r > 0.60$), the variable importance measure (VIMP) from a random forest model is used to determine which to retain. VIMP is derived by fitting a random forest to each indicator using all candidate covariates, permuting each covariate in turn, and measuring the resulting change in out-of-sample prediction error. The covariate with higher importance is retained from each collinear pair.

Second, a stepwise VIMP procedure is applied to the remaining covariates. At each step, the least important covariate is dropped until a single covariate remains, producing an importance ranking of covariates. Rankings are then averaged across all indicators to identify an overall ranking of covariate importance.

Finally, variance inflation factors (VIF) are calculated for the retained covariates for each indicator. Where a VIF exceeds 1.25, the stepwise VIMP analysis is repeated excluding that covariate.

STEP 6: DEFINE MODEL

The spatial variance-smoothing area level model can be fit in R using Stan, a probabilistic programming language for Bayesian statistical inference. Stan can interface with R through the rstan package, which allows users to write and compile Stan models directly within an R session (Guo et al., 2015). Gao and Wakefield provide a Stan script which defines the spatial variance-smoothing area level model (https://github.com/peteragao/VSALM/blob/main/inst/stan/spatial_joint_smooth_logit.stan). The script specifies the model structure, parameters, and priors. It can be adapted for a specific analysis by modifying the hyperparameter values or extending the model structure as needed. The script is organized into the standard Stan block structure, with each block serving a distinct role.

Key Stan blocks include:

Functions defines log-probability density functions used elsewhere in the model. Here, a penalized complexity (PC) prior is defined. PC priors are designed to shrink toward a simpler base model, avoiding overfitting while remaining weakly informative

Data declares all inputs that must be passed to the model. This includes the number of small areas, the number of small areas with valid estimates, spatial adjacency information, the direct survey estimates and their variance estimates.

Parameters declares the unknown quantities to be estimated by the model. These include a global intercept, fixed covariate effects, unstructured area-level random effects, a spatial random effects vector, and a mixing parameter controlling the balance between structured and unstructured spatial variation.

Transformed parameters constructs derived quantities from the parameters. This includes the linear predictor that is assembled from the intercept, random effects, and any covariates.

Model specifies the full joint log-posterior by declaring priors and the likelihood. Priors are placed on all parameters: the intercept and covariate effects receive weakly informative normal priors; the spatial effects use an intrinsic conditional autoregressive (ICAR) prior; the mixing parameter phi receives a Beta (0.5, 0.5) prior; and the precision parameters receive the PC priors defined in the functions block.

STEP 7: MODEL FITTING



Once the model is defined in the Stan script, it can be fit to the data. This requires four inputs prepared in earlier steps:1) small area boundaries (shapefile), 2) spatial adjacency matrix, 3) direct estimates of prevalence and variance, and 4) covariates.

Before fitting, several MCMC sampling parameters must be set. The number of warm-up (burn-in) iterations, sampling iterations, and chains should be specified. Values of 1000, 1000, and 4 respectively are generally adequate for this model and provide a reasonable balance between computation time and convergence.

The model is compiled in Rstudio using the **cmdstan_model()** function **(Figure 4)**, which indexes the Stan script in which the model is defined. Model inputs are assembled into a named list that maps directly onto the variable names declared in the data block of the Stan script.

Figure 4: Rstudio cmdstan_model function example

```
# Compile model
mod <- cmdstan_model(stanfile)
data_list <- list(
D = D,
X = X_full,
N = nrow(ind_dist),
NS = nrow(ind_dist_complete),
adm2_index = ind_dist_complete$district_id,
p_hat = ind_dist_complete$dir,
v_hat = ind_dist_complete$dir_var,
d = ind_dist_complete$degf,
k = ind_dist_complete$obs_un,
N_edges = length(prep$n1),
node1 = prep$n1,
node2 = prep$n2,
scaling_factor = prep$scaling_factor
)
```

Models can be fit in a cascade approach, starting with a baseline model that contains no covariates and only an intercept. The next model can be fit with the most important covariate identified in **Step 5**. A subsequent model can be fit with the two most important covariates, and so on, until a *full model* is fit with all desired covariates.

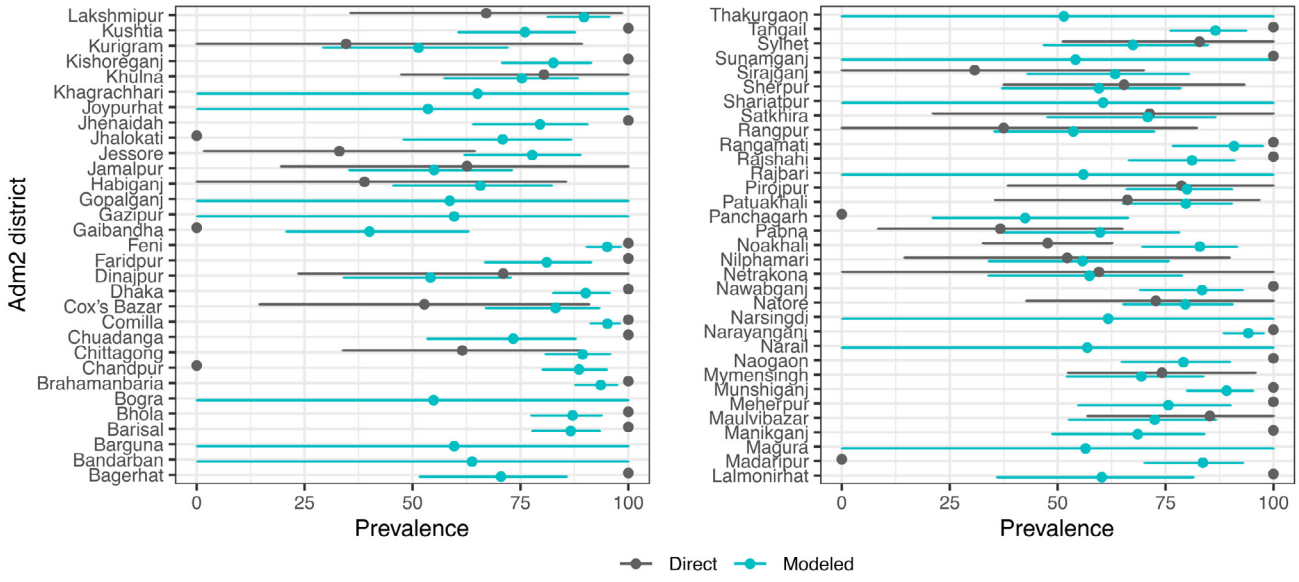
For each model, it is necessary to perform checks to ensure that MCMC simulations have converged and that the estimates parameters are reliable. First, visually inspected trace plots of selected parameters to confirm that the chains mixed thoroughly and that there were no sampling issues. Second, examine the potential scale reduction factor (denoted as $\hat{R}$) which compares variation within and between chains. Values of $\hat{R}$ close to 1 indicate that the chains have mixed well. Third, check the effective sample size (ESS), which measures how many independent samples the chains effectively provide after accounting for autocorrelation. Following the guidance of *rstan* developers, aim for bulk-ESS and tail-ESS values of at least 100 per chain, which helps ensure that parameter estimates and uncertainty intervals were stable (Guo et al., 2015). Adjust the number of chains, warm-up period, and total draws to meet these criteria.

STEP 8: VISUALIZE ESTIMATES AND VALIDATE ESTIMATES

After the models are fit and convergence is confirmed, posterior draws for the modeled area-level prevalence estimates can be extracted for summarization and mapping. Extract the mean and variance from the posterior draws for the estimated prevalence. Take the quantiles to generate 95% credible intervals. Plot these alongside direct estimates with design-based variance to see how the direct and modeled estimates compare.

Figure 5 below displays direct and modeled estimates with 95% uncertainty bounds for the prevalence of children under 24 months of age in Bangladesh who had diarrhea in the preceding two weeks and received oral rehydration solution as treatment. This was an indicator with sparse data in the 2022 DHS. There were 13 districts with no observations (e.g. Bandarban, Barguna, Bogra, Gazipur) and design-based variance could not be estimated for 28 districts with sparse data (e.g. Bagerhat, Barisal, Bhola, Brahamanbaria). When design-based variance could be directly estimated, uncertainty intervals were often extremely wide, with roughly 80% of indicators having a range of more than 50 percentage points between the direct estimate lower and upper bounds.

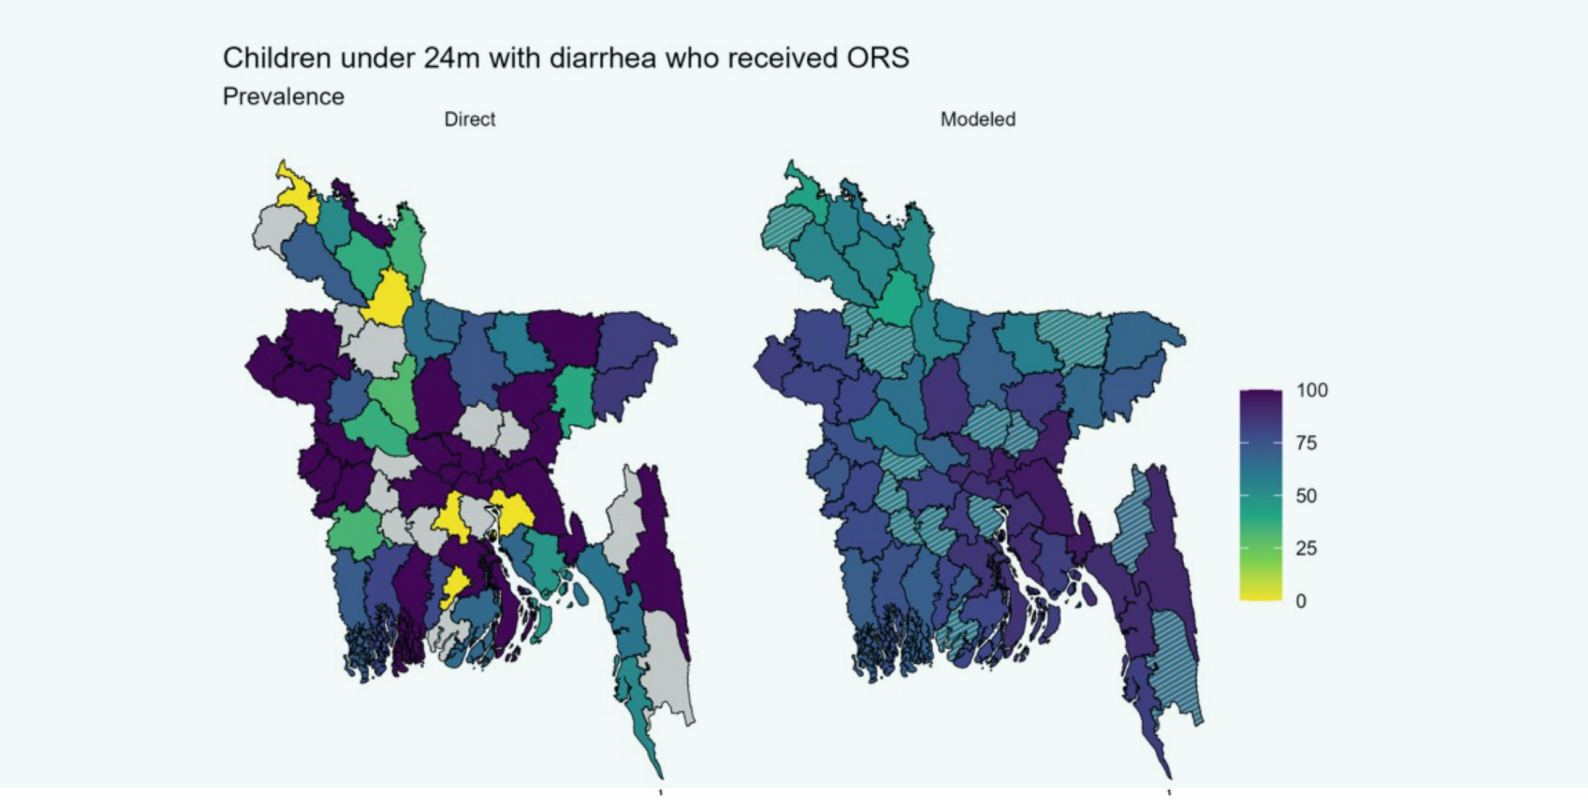
Figure 5: Direct vs Bayesian small area estimates for prevalence of children under 24 months of age who had diarrhea in the preceding two weeks and received oral rehydration solution as treatment (Bangladesh DHS 2022)



Compared with the direct estimates, the modeled prevalence estimates were less extreme, with values that were initially low tending to increase, while higher values decreased. For districts with no observations, the model generated estimates by borrowing strength from neighboring areas. However, some of these estimates remained unreliable, with very wide credible intervals. In cases where direct estimates had uncertainty bounds, the modeled uncertainty intervals were generally narrower, reflecting the stabilizing influence of the spatial smoothing and covariate information.

The direct and modeled prevalence estimates can also be visualized in map form to illustrate their spatial distribution **(Figure 6)**. The below figure plots direct and modeled estimates of prevalence side by side. Districts with no observations are shaded grey on the left panel displaying the direct estimates. On the right panel displaying modeled estimates, districts with a coefficient of variation (CV) larger than 30% are crosshatched with grey.

Figure 6: Mapped direct vs Bayesian small area estimates for prevalence of children under 24 months of age who had diarrhea in the preceding two weeks and received oral rehydration solution as treatment (Bangladesh DHS 2022)



The CV is one way of assessing the reliability of the modeled estimates. The CV can be calculated by deriving the standard error from the modeled variance and then dividing by the prevalence. This provides a measure of the size of the uncertainty relative to the size of the estimate itself. Following conventional guidance, a threshold of 30% can be used to classify estimates as reliable or unreliable.

Often the same districts that had no observations were deemed to have unreliable modeled estimates. Otherwise, the effects of spatial smoothing are evident, with extreme values in the direct estimates tending to shift towards the values of their neighbors in the modeled estimates. Despite these changes, the overall spatial pattern of low- and high-prevalence areas remains largely consistent.

The smoothed variance estimates from the model can also be visualized alongside the direct estimates of variance **(Figure 7)**. Comparing these maps provides insight into how the Bayesian model stabilizes uncertainty across small areas. Direct variance estimates are often highly variable, particularly in districts with small sample sizes, where they may be inflated or impossible to estimate. In contrast, the spatial variance-smoothing model borrows information from neighboring areas to produce more stable and realistic variance estimates while preserving genuine geographic differences in uncertainty. Visualizing both sets of estimates allows users to identify areas where uncertainty has been substantially reduced through modelling, as well as districts where uncertainty remains high due to limited data. These comparisons provide an important assessment of the reliability of the modelled estimates and help guide interpretation of the resulting small area maps.

Figure 7: Mapped smoothed variance for direct vs Bayesian small area estimates for prevalence of children under 24 months of age who had diarrhea in the preceding two weeks and received oral rehydration solution as treatment (Bangladesh DHS 2022)

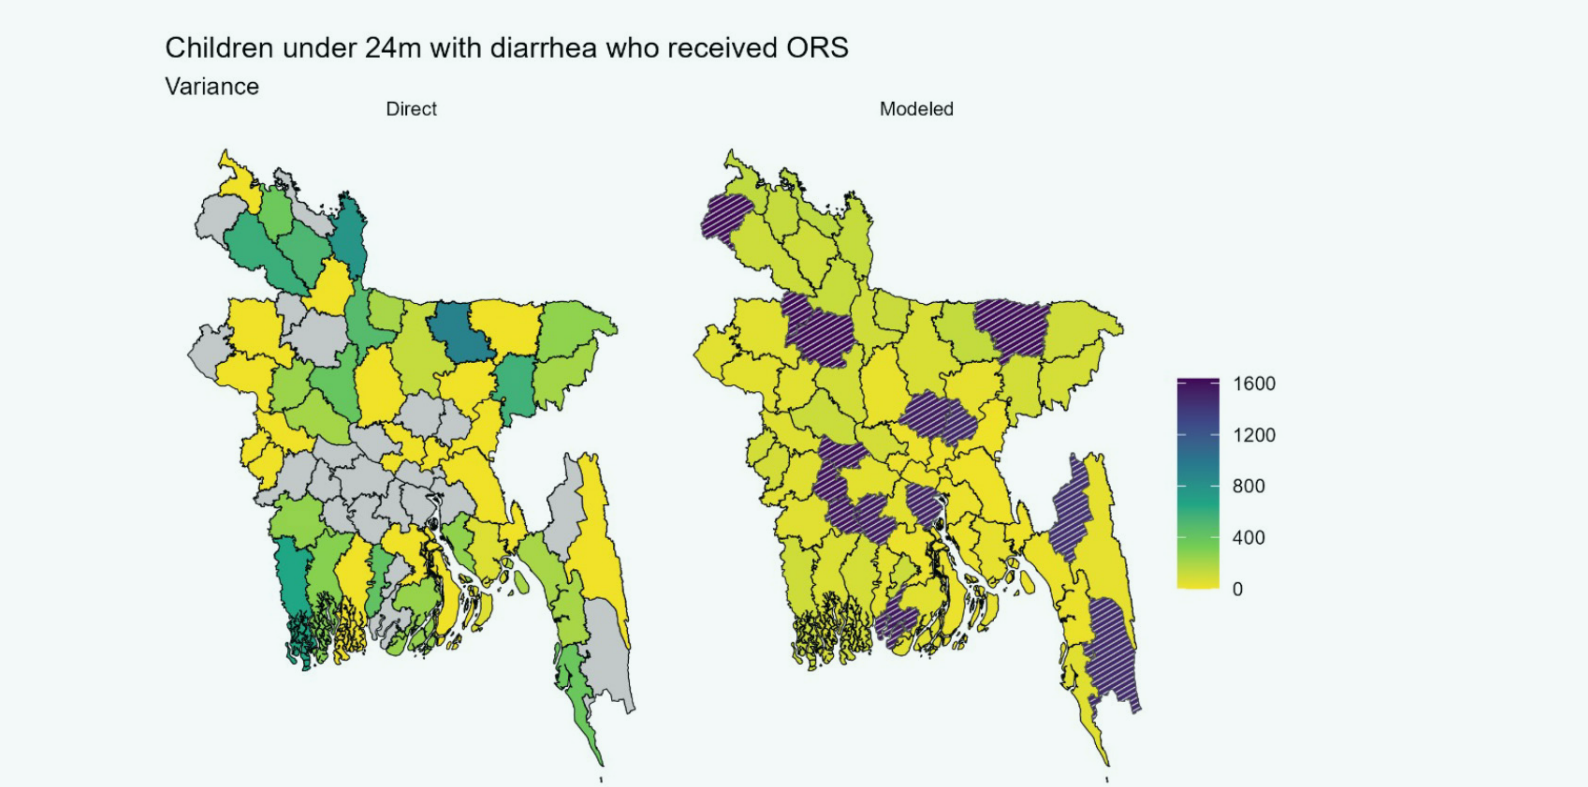


Figure 8 displays direct and modeled estimates with 95% uncertainty bounds for the prevalence of children aged 6-23 months who had received vitamin A supplements. Unlike the ORS example, this indicator was supported by substantially more survey observations, with all districts containing data and only one district (Bandarban) lacking sufficient information to estimate the design-based variance directly.

Figure 8: Direct vs Bayesian small area estimates for vitamin A supplementation to children 6-23 months old, Bangladesh DHS 2022

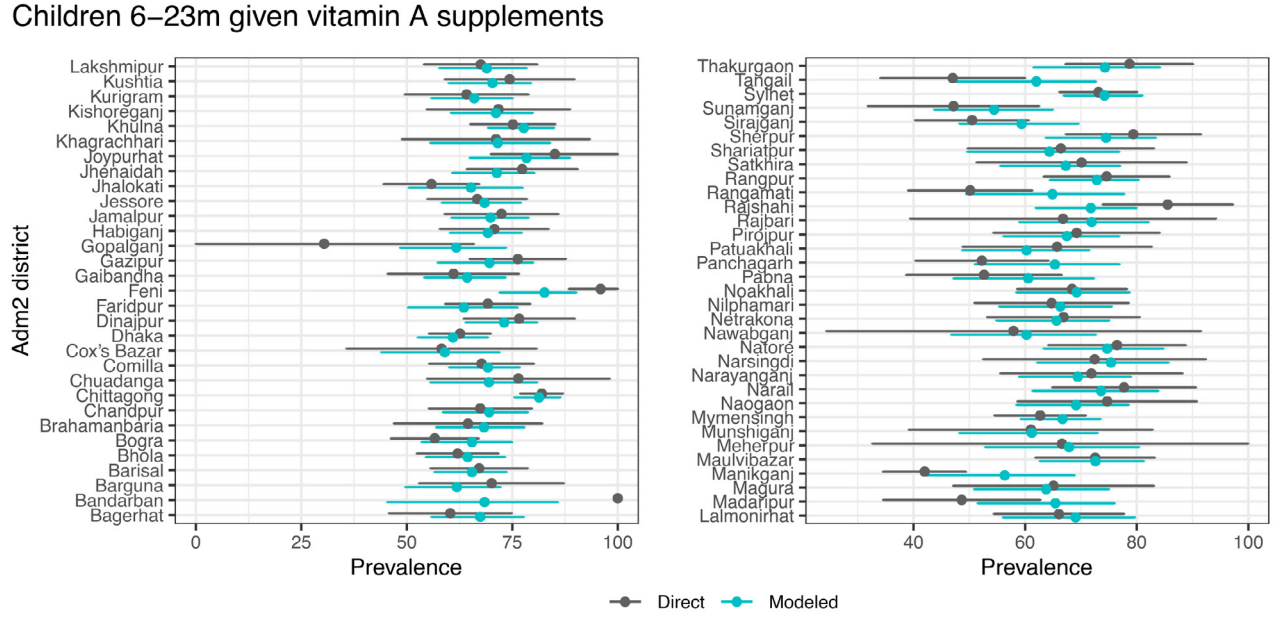


Figure 9 presents the mapped direct and modelled estimates of vitamin A supplementation. All modelled prevalence estimates met the reliability threshold of a CV below 30%. The effects of spatial smoothing remain evident, with districts that had unusually low (e.g. Gopalganj) or high (e.g. Bandarban and Feni) direct estimates shifting towards values that are more consistent with neighboring districts. However, because the underlying survey data were less sparse, the adjustments were generally modest, and the overall geographic pattern of vitamin A supplementation remained largely unchanged. This example illustrates that Bayesian SAE not only improves estimation for sparse indicators but also stabilizes estimates for more common interventions while preserving meaningful spatial variation.

The corresponding variance maps **(Figure 10)** demonstrate the stabilizing effect of the spatial variance-smoothing model on uncertainty estimates. Compared with the direct variance estimates, which are often highly variable or unavailable in districts with sparse data, the modelled variance estimates are smoother and more spatially coherent while still reflecting areas of relatively high and low uncertainty. For this indicator, where survey data were more abundant than for ORS treatment, the differences between the direct and modelled variance estimates were generally modest. Nevertheless, the Bayesian model reduced erratic fluctuations in variance caused by small sample sizes and produced more stable uncertainty estimates, providing greater confidence in the resulting district-level prevalence estimates.

Figure 9: Mapped direct vs Bayesian small area estimates for vitamin A supplementation to children 6-23 months old (Bangladesh DHS 2022)

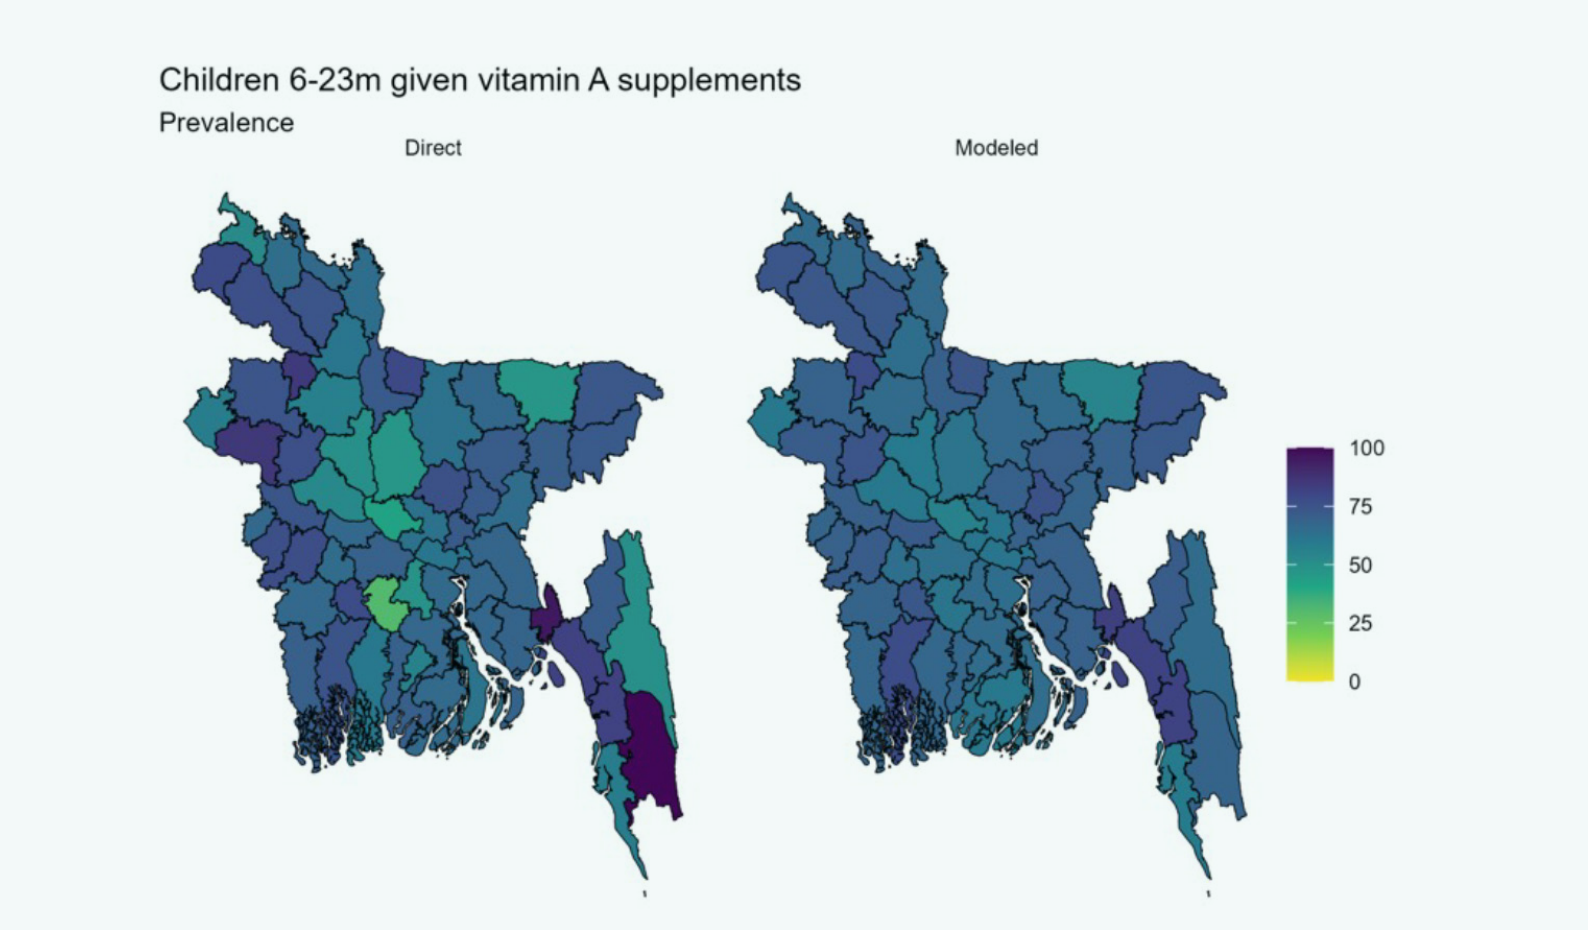
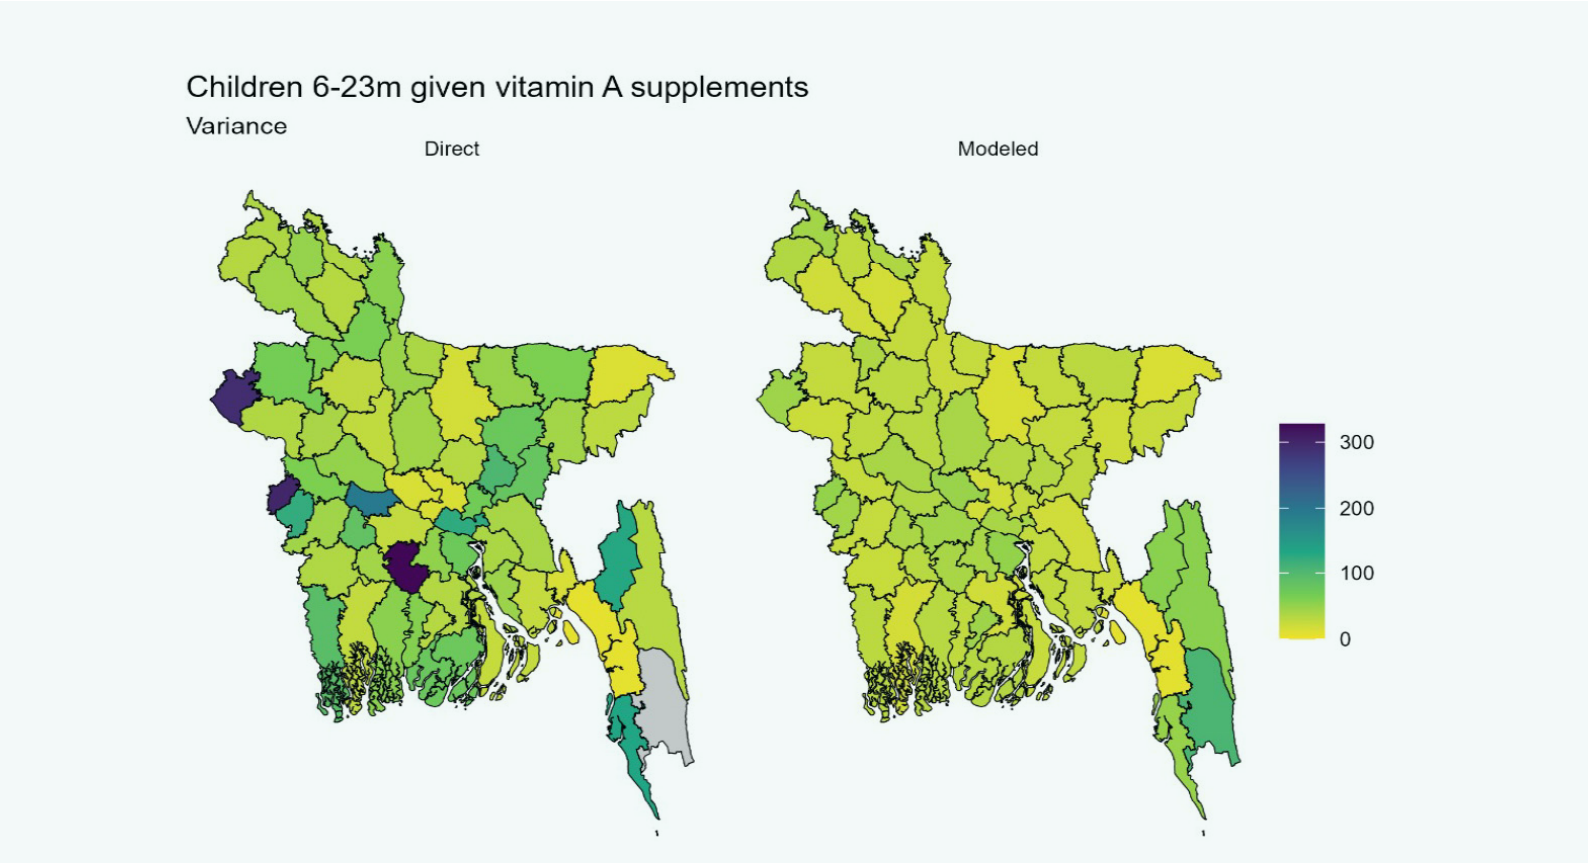


Figure 10: Mapped smoothed variance for direct vs Bayesian small area estimates for vitamin A supplementation to children 6-23 months old (Bangladesh DHS 2022)



Sensitivity analyses and validation of the models is critical. The sensitivity of choice of priors should be assessed by fitting the model under a range of prior specifications and comparing the resulting estimates. Ideally, estimates will be robust to reasonable variation in prior choice, with no substantial changes in posterior prevalence or uncertainty intervals.

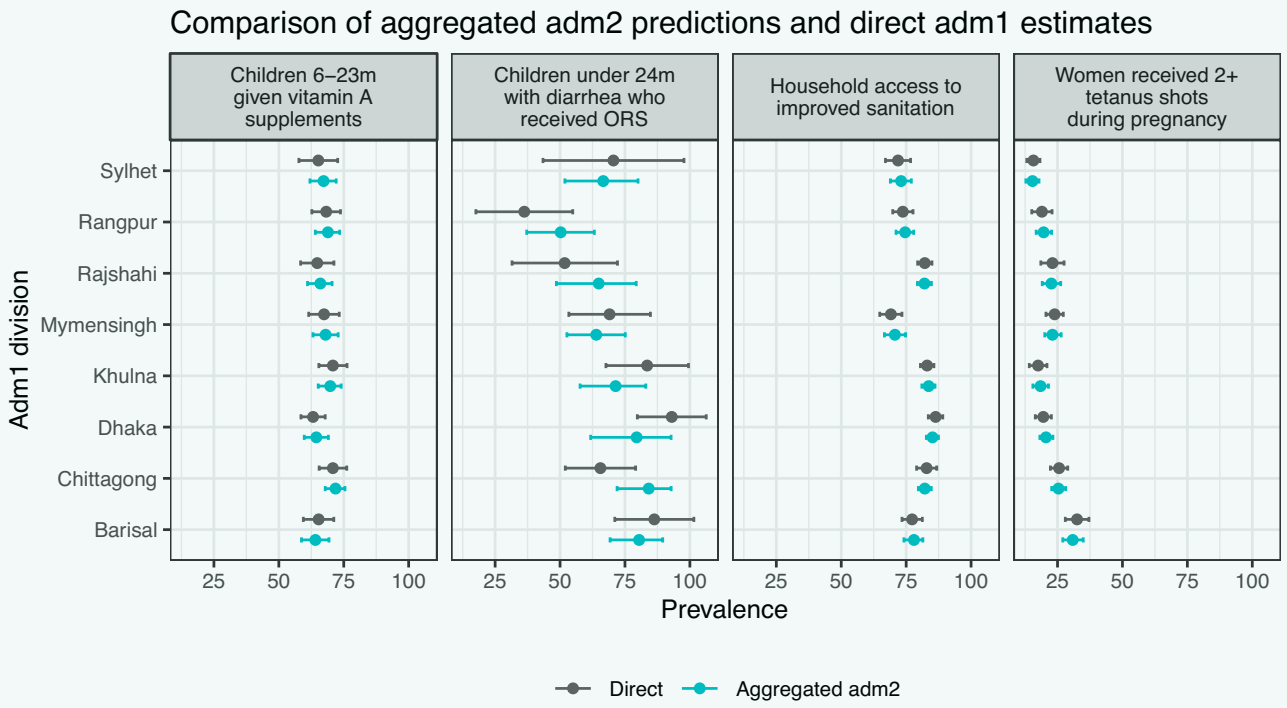
Model validation

Since no ground truth exists for small area prevalence, model validation relies on indirect approaches. As a first step, estimates from the spatial variance-smoothing area level model can be compared with those from the Bayesian hierarchical area-level smoothing model implemented in the R SUMMER package (Li et al., 2025) which uses the INLA method (Rue et al., 2009). The two models should produce broadly similar prevalence estimates. Uncertainty intervals from the spatial variance-smoothing model may be wider, which is expected: by explicitly accounting for heterogeneity in the design-based variances used as model inputs, it provides a more realistic characterization of uncertainty.

A further validation approach is to aggregate the modelled small area estimates upward to the administrative level at which reliable direct estimates are available. For DHS surveys, this is typically the administrative level 1 and national levels. Aggregation is performed by taking a weighted average of the small area estimates, where each area is weighted by the sum of survey weights for all observations within it. The aggregated modelled estimates can then be compared with direct estimates computed from the survey data, with validation assessed by examining whether the uncertainty intervals overlap for each indicator and area unit.

Figure 11 below presents an example of this comparison for a subset of indicators from the 2022 Bangladesh DHS. Across all four indicators, the aggregated small area estimates closely align with the direct division-level estimates, with substantial overlap in the corresponding uncertainty intervals. This agreement provides evidence that the Bayesian SAE model preserves the overall distribution of intervention coverage while producing stable estimates at finer geographic scales. Minor differences between the direct and aggregated estimates are expected due to spatial smoothing but are generally small, supporting the validity of the modelling approach.

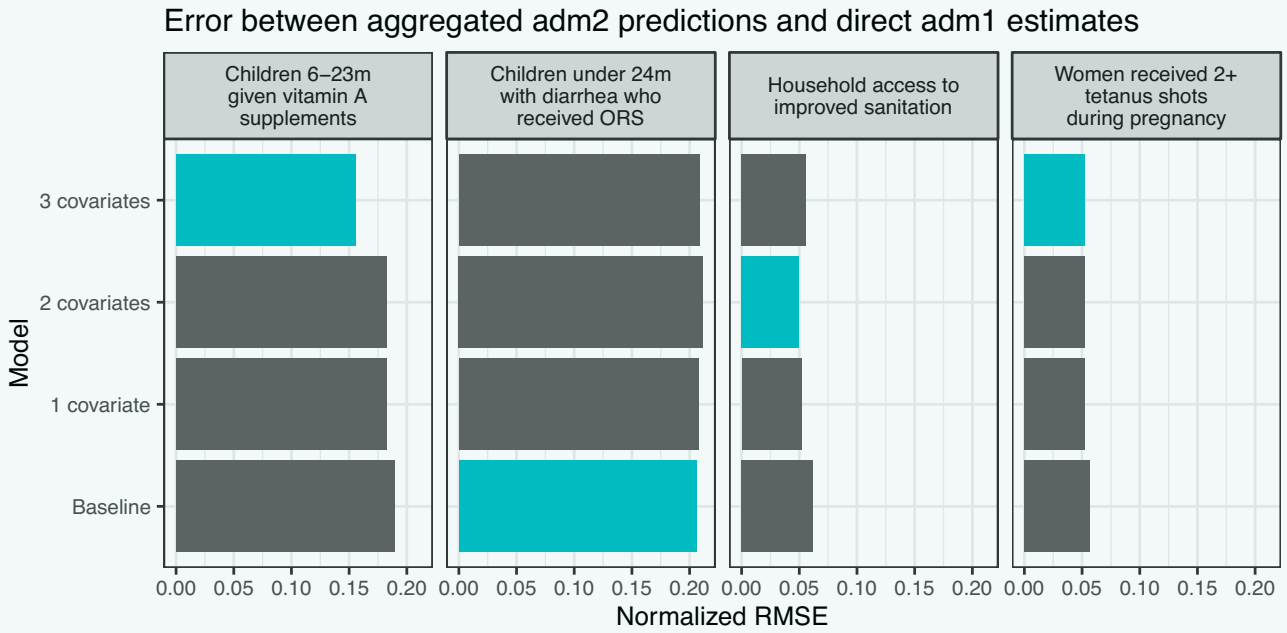
Figure 11: Comparison and validation of coverage estimates (Bangladesh DHS 2022)



This aggregation comparison can also be used to evaluate alternative model specifications such as a baseline model with no covariates and models incorporating one, two, or three covariates. The normalised root mean squared error (nRMSE) between aggregated modelled and direct estimates is calculated for each specification, and the model with the lowest nRMSE is preferred. Rather than selecting a different model specification for each indicator individually, it is advisable to average nRMSE values across indicators within substantive groups (e.g. antenatal care, delivery and postnatal care, child nutrition interventions) and identify the specification that performs most consistently across each group. This ensures that the final estimates for composite coverage are drawn from the model that performed consistently well.

Figure 12 compares the predictive performance of alternative model specifications using the normalized root mean squared error (nRMSE) between aggregated district-level estimates and direct division-level estimates. Lower nRMSE values indicate better agreement with the direct survey estimates. The optimal model specification varied by indicator, with the three-covariate model performing best for vitamin A supplementation and maternal tetanus vaccination, the baseline model performing best for ORS treatment, and the two-covariate model performing best for improved sanitation. These findings highlight the importance of evaluating alternative model specifications rather than assuming that more complex models always provide superior performance.

Figure 12: Comparison of error (nRMSE) between models and indicators (Bangladesh DHS 2022)



In practice, model selection should be based on predictive accuracy across related indicators while balancing model complexity and interpretability.

STEP 9: COMMUNICATION

The final step is to communicate the small area estimates in a manner that supports their use by the target audience. Results should be presented using clear and interpretable visualizations, tables, and infographics that highlight geographic variation in intervention coverage and the associated uncertainty. Maps of district-level intervention coverage can help identify underserved areas, while comparisons across indicators can reveal geographic patterns and opportunities for integrated programme delivery. The DataDENT technical note on Building Data Literacy: Intervention Coverage includes additional guidance on visualizing and communicating intervention coverage data.[3]

Communicating results in a transparent and accessible manner maximizes the value of SAE for informing policy, planning, and strengthening programmes. Clearly document in the main text or footnotes that small area estimates are modeled and whenever appropriate, also include measures of reliability, such as credible intervals or coefficients of variation. Communications to policy maker and programme audiences will emphasize key messages and recommendations rather than the SAE methods. However, more detailed information about the SAE methods should be accessible.

[3]Access technical note at https://datadent.org/resource/data-literacy-coverage/.

Strengths & Limitations

A key strength of the framework is its use of the spatial variance-smoothing area-level model, which explicitly accounts for heterogeneity in sampling variance across geographic areas. The approach is flexible and can be adapted to a wide range of intervention coverage indicators collected using household surveys. The framework also incorporates several procedures to improve the reliability of the resulting estimates, including assessment of data sparsity, stabilization of design-based variance estimates using phantom clusters, evaluation of estimate reliability using coefficients of variation, and validation against direct estimates at higher administrative levels.

The primary limitations of this approach relate to its computational requirements and the technical expertise needed for implementation. Fitting Bayesian spatial models using MCMC methods is more computationally intensive than approximate methods and requires familiarity with Bayesian modelling, the Stan programming language, and the interpretation of model diagnostics such as trace plots, R-hat, and effective sample size.

The quality of the resulting estimates also depends on the availability of high-quality survey data and relevant covariates. Although Bayesian spatial models can generate estimates for areas with sparse or no observations, uncertainty remains greater in these locations and estimates should be interpreted alongside their corresponding credible intervals or coefficients of variation.

Finally, alternative software implementations, such as the SUMMER package using INLA, provide more accessible out-of-the-box functionality and substantially faster computation. However, the fully Bayesian spatial variance-smoothing approach presented in this technical note offers greater flexibility and more explicitly models heterogeneity in sampling variance, making it particularly well suited to estimating intervention coverage from complex household survey data.

Conclusion

This technical note presents practical guidelines for generating small area estimates of nutrition intervention coverage using Bayesian spatial modeling. By combining design-based survey estimates with spatial information and relevant covariates, the approach produces more stable district-level estimates than direct survey estimation alone while providing a rigorous quantification of uncertainty. These reliable small area estimates can be used to support co-coverage analysis and more broadly evidence-based policy development, program planning, resource allocation, and monitoring of health and nutrition interventions.

References

Besag, J., York, J., & Mollié, A. (1991). Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics*, *43*(1), 1–20. https://doi.org/10.1007/BF00116466

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., & Riddell, A. (2017). Stan: A Probabilistic Programming Language. Journal of Statistical Software, 76, 1–32. https://doi.org/10.18637/jss.v076.i01

Datta, G., & Ghosh, M. (1991). Bayesian Prediction in Linear Models: Applications to Small Area Estimation. *The Annals of Statistics*, *19*. https://doi.org/10.1214/aos/1176348369

De Smedt, T., Simons, K., Van Nieuwenhuyse, A., & Molenberghs, G. (2015). Comparing MCMC and INLA for disease mapping with Bayesian hierarchical models. *Archives of Public Health*, *73*(S1), O2. https://doi.org/10.1186/2049-3258-73-S1-O2

DHS Program. (2025). *DHS-Indicators-R* [R]. The Demographic and Health Surveys (DHS) Program. https://github.com/DHSProgram/DHS-Indicators-R (Original work published 2021)

Fay, R. E., & 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*, *85*, 398–409. https://doi.org/10.1080/01621459.1979.10482505

Gao, P. A., & Wakefield, J. (2023). A Spatial Variance-Smoothing Area Level Model for Small Area Estimation of Demographic Rates. *International Statistical Review*, *91*(3), 493–510. https://doi.org/10.1111/insr.12556

Ghosh, M., & Rao, J. N. K. (1994). Small Area Estimation: An Appraisal. *Statistical Science*, *9*(1). https://doi.org/10.1214/ss/1177010647

Guo, J., Gabry, J., Goodrich, B., Johnson, A., Weber, S., & Badr, H. S. (2015). *rstan: R Interface to Stan* (p. 2.32.7) [Dataset]. https://doi.org/10.32614/CRAN.package.rstan

HDX. (2025). *Bangladesh—Subnational Administrative Boundaries, Humanitarian Dataset*. Retrieved November 10, 2025, from https://data.humdata.org/dataset/cod-ab-bgd

Li, Z. R., Martin, B. D., Dong, T. Q., Fuglstad, G.-A., Paige, J., Riebler, A., Clark, S., & Wakefield, J. (2025). *Space-Time Smoothing of Survey Outcomes using the R Package SUMMER* (No. arXiv:2007.05117). arXiv. https://doi.org/10.48550/arXiv.2007.05117

Martino, S., & Riebler, A. (2019). *Integrated Nested Laplace Approximations (INLA)* (No. arXiv:1907.01248). arXiv. https://doi.org/10.48550/arXiv.1907.01248

Neal, R. M. (2011). Handbook of Markov Chain Monte Carlo. In S. Brooks, A. Gelman, G. Jones, & X.-L. Meng (Eds.), *MCMC using Hamiltonian dynamics*. Chapman & Hall/CRC Press. https://doi.org/10.1201/b10905

Rao, J. N. K., & Molina, I. (2025). *Small Area Estimation*. John Wiley & Sons, Inc. https://onlinelibrary.wiley.com/doi/book/10.1002/9781118735855

Riebler, A., Sørbye, S. H., Simpson, D., & Rue, H. (2016). An intuitive Bayesian spatial model for disease mapping that accounts for scaling. *Statistical Methods in Medical Research*, *25*(4), 1145–1165. https://doi.org/10.1177/0962280216660421

Rue, H., Martino, S., & Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 71(2), 319–392. https://doi.org/10.1111/j.1467-9868.2008.00700.x

Sugasawa, S., & Kubokawa, T. (2020). Small area estimation with mixed models: A review. *Japanese Journal of Statistics and Data Science*, *3*(2), 693–720. https://doi.org/10.1007/s42081-020-00076-x


Data
DENT
Data for Decisions in Nutrition