sae4health: A Guided Tour of the DHS / MIS Shiny App

Subnational estimates of health indicators,
with Kenya as a running example

The sae4health team

CEMA workshop

Small Area Estimation for Improving Health in LMICs

What we will cover today

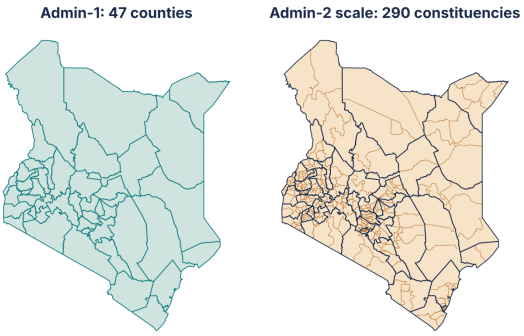
1. Why small area estimation (SAE)?
2. Meet sae4health and its ecosystem
3. A guided tour of the app, panel by panel
4. Data sparsity checks and model choice
5. The models behind the buttons
6. Reading the visualizations (a Kenya walk-through)
7. Reports and the Tool Kit
8. DHS vs MIS surveys
9. Shapefiles and boundaries
10. Hands-on exercise, troubleshooting, what's next

Running example: the **2022 Kenya DHS** (preloaded in the web app) and the **2020 Kenya Malaria Indicator Survey**. Indicator used throughout: **MCV1**, Percentage of children age 12–23 months who had received MCV 1 (Measles-Containing Vaccine).

Why small area estimation?

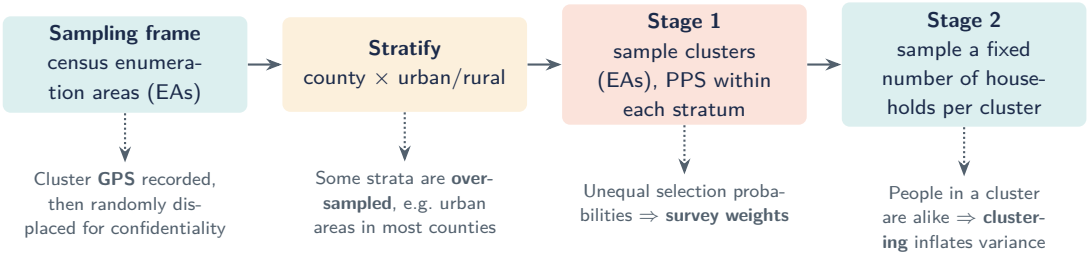
National surveys, local decisions

- Since the 2010 Constitution, Kenya's **47 counties** run health services and write their own County Integrated Development Plans.
- Planning needs indicators for **counties and sub-counties**, but household surveys are designed mainly for national and county-level reporting.
- The DHS Program computes national estimates of **2,500+ indicators** per survey. Most are binary (yes/no) outcomes.
- **The gap:** at finer levels, sample sizes shrink and simple survey-weighted estimates become noisy or impossible.



Croft et al. (2023), cited in Dong et al. (2025), arXiv:2504.16435. Boundaries: Kenya counties and IEBC constituencies.

How a DHS sample is drawn: stratified two-stage cluster design



Any valid estimate must respect this design: weights, strata and clusters. sae4health does this for you, but you should know *why* it matters when you read the results.

Kenya DHS 2022 at a glance

1,692	92	25	42,022	98%
clusters sampled	strata: 47 counties × urban/rural	households per cluster	households selected	household response rate

- Frame: Kenya Household Master Sample Frame (2019 census EAs).
- 1,026 rural and 666 urban clusters; **Nairobi and Mombasa are entirely urban**, which is why there are 92 strata rather than 94.
- Power allocation across counties, so county estimates are comparable for *some* indicators.

Designed for

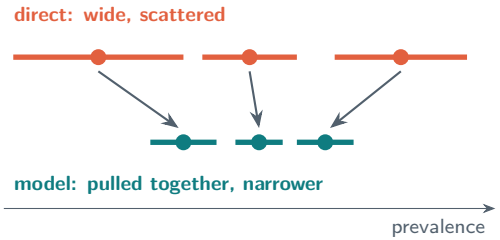
National, urban/rural, and (for some indicators) **county**-level estimates.

Not designed for sub-county estimates. That is where SAE comes in.

KNBS & ICF (2023), KDHS 2022 report, Appendix A; KNBS / World Bank microdata catalogue.

The core idea of SAE: borrowing strength

- **Direct estimate:** use only the data inside the area. It is unbiased, but has huge variance when n is small.
- **Model-based (SAE):** link areas together with a statistical model, so small areas *borrow strength* from
  - their **neighbours** (spatial smoothing)
  - the **national** or county level (shrinkage)
  - **covariates** (e.g. travel time to a facility)
- The trade-off: a little bias for a large drop in variance.



Areas with *less* data get pulled *more*. Areas with plenty of data keep estimates close to their direct value.

Meet sae4health

sae4health in one slide

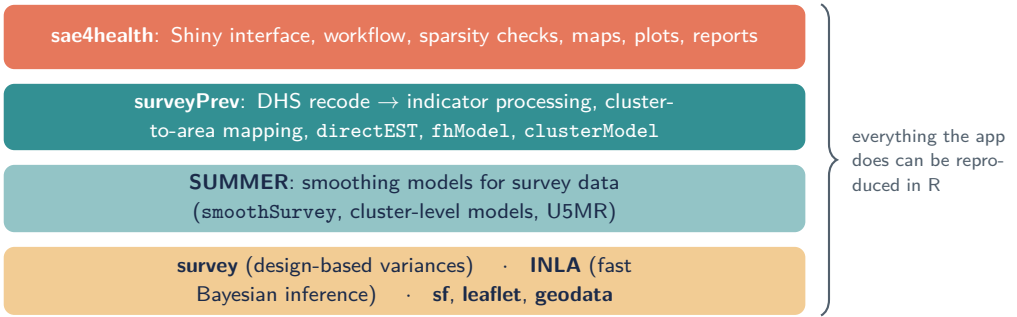
- An **R Shiny app**, distributed as an R package on CRAN (`sae4health`), for subnational prevalence mapping from household surveys.
- **No coding needed:** pick a country, survey and indicator, run checks and models, then explore maps, plots, tables and a PDF report.
- Three estimation approaches: **direct**, **area-level** and **unit-level**, with Bayesian inference via **INLA**.
- Built by Yunhan Wu, Victoria Knutson, Qianyu Dong, Jieyi Xu, Zehang Richard Li and Jon Wakefield (UW and UC Santa Cruz).

430+ binary indicators	**150+** DHS/MIS surveys
60 countries	**0** lines of code needed

Web app: https://rsc.stat.washington.edu/sae4health/ Docs: https://sae4health.stat.uw.edu/

Wu et al. (2025), arXiv:2505.01467; CRAN package sae4health 1.2.3.

The software ecosystem behind the app



- The app's **Model Log** prints the equivalent **surveyPrev** calls, so you can move from clicking to scripting.
- The indicator definitions come from `surveyPrev::indicatorList`, which is organised by DHS final-report chapter.

Two ways to run the app

Web version

- Open a browser: nothing to install.
- Most surveys are **preloaded on the server**, including **Kenya DHS 2022**, so there is no data upload.
- GADM boundaries are preloaded.
- For confidentiality, the raw data cannot be downloaded or fully displayed.

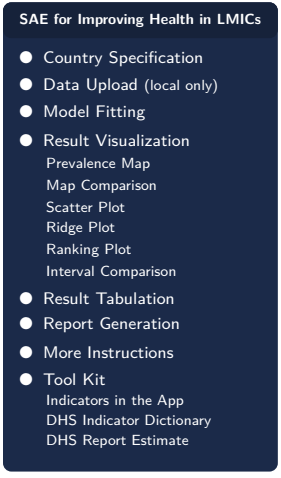
Local version (R package)

- Install R, INLA, SUMMER, surveyPrev and sae4health.
- **Upload your own DHS .zip**, which needs approved DHS data access.
- You can preview and download the prepared analysis data (.csv).
- Runs offline once installed. You can also deploy your own variants.

```
install.packages("INLA", dep = TRUE,
  repos = c(getOption("repos"), INLA = "https://inla.r-inla-download.org/R/testing"))
devtools::install_github(c("richardli/SUMMER", "richardli/surveyPrev", "wu-thomas/sae4health"))
library(sae4health); run_app()
```

A guided tour of the app

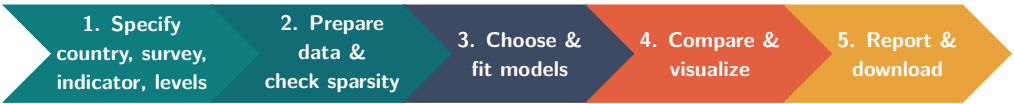
App layout: the sidebar is your roadmap



- The menu runs **top to bottom in workflow order**: specify, load data, fit, look, tabulate, report.
- **More Instructions** is the landing page, with expandable step-by-step boxes and links to videos.
- The **Tool Kit** holds reference material: the indicator list, the DHS dictionary, and official DHS estimates.
- In the web version the Data Upload tab is hidden; data load from the server instead.

> Changing the country, survey or indicator **resets** fitted models and the model log, so finish one analysis before switching.

The workflow in five steps



Steps 1–2: understand the data
Which areas have clusters? Does the app's national estimate match the DHS report?

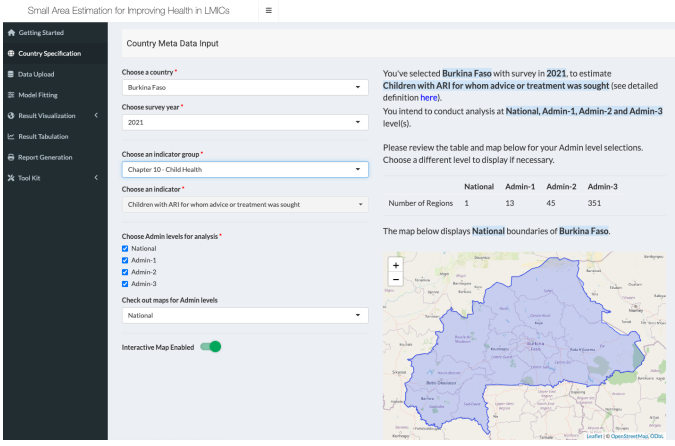
Step 3: pick an approach
Follow the recommendations and the sparsity check, and fit several models where allowed.

Steps 4–5: judge and share
Compare models and uncertainty, choose what you trust, then export a report.

The app follows the four-stage *prevalence mapping workflow* of Dong et al. (2025): understand the design and data, analyse with a range of models, evaluate and compare, then summarise and visualise.

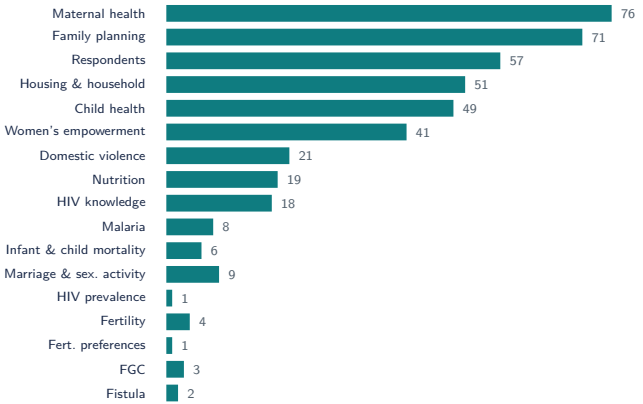
Step 1 — Country Specification

1. **Choose a country:** Kenya.
2. **Choose a survey year.** The labels show the survey type, e.g. 2022 (`DHS`) or 2020 (`MIS`).
3. **Choose an indicator group**, i.e. a DHS report chapter.
4. **Choose an indicator** (the list is searchable).
5. **Tick the Admin levels** to analyse: National, Admin-1, Admin-2.
6. **Check the maps and region counts** for each level.



Screenshot of the Country Specification panel (Burkina Faso shown). The panel reads "You've selected ... to estimate ... at National, Admin-1, Admin-2 level(s)."

Choosing an indicator: organised by DHS report chapter



- The menu is built from `surveyPrev::indicatorList`. The current GitHub version has **437 entries across 17 chapters** (above). The published paper quotes "150+" for the CRAN release.
- Which indicators you see depends on what the survey collected. Each indicator needs particular **recode files** (IR, PR, KR, BR, HR, MR, …).
- Each indicator has a DHS-style ID, e.g. `CH_VACC_C_MSL` for MCV1.

Counts computed from surveyPrev/data/indicatorList.rda (GitHub, Nov 2026).

Kenya-relevant indicators to try

ID	Description	Recode	Chapter
CH_VACC_C_MSL	Children 12–23 months who had received MCV1 (measles-containing vaccine)	KR	Child health
RH_ANCN_W_N4P	Antenatal visits for pregnancy: 4+ visits (2 years before the survey)	IR	Maternal health
RH_DELP_C_DHF	Place of delivery: health facility (5 years before the survey)	BR	Maternal health
CH_VACS_C_BAS	Fully vaccinated (8 basic antigens)	KR	Child health
CH_DIAT_C_ORS	Treatment of diarrhoea: oral rehydration solution	KR	Child health
CN_NUTS_C_HA2	Stunting (HAZ < -2), children under five	PR	Nutrition
FP_CUSA_W_MOD	Current use of any modern contraceptive method (all women)	IR	Family planning
ML_NETP_H_IT2	Households with $\geq$1 ITN per two persons	HR	Malaria
ML_PMAL_C_RDT	Malaria prevalence according to RDT	PR	Malaria

Click the indicator name on the first panel to see its **full definition**: numerator, denominator, reference period.

Admin levels: what “Admin-1” and “Admin-2” mean in Kenya

Admin-1: 47 counties



Admin-2 scale: 290 constituencies



- **Admin-1 = 47 counties.** These are also the DHS stratification units: data are designed to be adequate here.
- **Admin-2:** about **300** sub-county units in GADM. The map shows the 290 IEBC constituencies as a stand-in.
- The app shows a **table of region counts** and an interactive map for each selected level. Check them before you go on.

> Admin-2 units are not survey strata: expect sparse data and plan to use models.

Step 2 — Getting the data in

Web version

- Data is pre-loaded in the app
- The app builds the **analysis dataset** automatically once country/survey/indicator is chosen.

Local version

1. Register at dhsprogram.com; describe your project.
2. Request the country/year datasets, **including GPS data**. Approval usually takes about 48 hours.
3. Download **Stata (.dta)** recodes *and* the geographic data in **one .zip**.
4. Upload the .zip and submit.

Which files? It depends on the indicator

IR	Individual (women's) Recode
PR	Household Member Recode
KR	Children's Recode
BR	Births Recode
HR	Household Recode
MR	Men's Recode
AR	HIV Test Results Recode
GE	Geographic data (cluster GPS)

Verify first: does the app reproduce the DHS report?

- Before modelling, the app asks you to check its **national direct estimate** against the official **DHS/MIS final report** value. See **Tool Kit → DHS Report Estimate**.
- Official values come from the **DHS API**. If the API has no value, compare by hand with the report table.
- For indicators with several versions (by age group, sex, …), the app defaults to the total, unstratified version.

> A mismatch means the indicator coding may not fit this survey. Do not model until it is resolved; report it to the developers.

Kenya DHS 2022 — MCV1

From the app	**88.7%**
KDHS 2022 final report	**89%**

Step 3 — Model Fitting: the recommendations

The app opens with a recommendation box:

National	Survey-weighted **direct estimates**
Admin-1	**Area-level** (Fay–Herriot) model
Finer levels	**Unit-level** (cluster-level) model

- The panel is a **grid of methods** × **Admin levels**. Tick the combinations you want.
- Fitting *several* models and comparing them is encouraged.
- The *Model Details* tab documents each method.

	Direct	Area	Unit
National	✓	–	–
Admin-1	✓	✓	✓
Admin-2	▲	▲	✓

Example of a selection grid after the sparsity check (illustrative).

The Data Sparsity Check

Click **Data Sparsity Check**. For every selected method × level, the app counts:

- **No-data areas:** no sampled cluster.
- **Low-information areas:** a direct estimate exists but is 0 or 1, or its variance is missing, zero, or not finite.

Method	Flag raised when ...
Direct	no-data + low-info $> 25\%$ of areas $\Rightarrow$ ▲
Area-level	no-data + low-info $> 25\%$ $\Rightarrow$ not fitted
Unit-level	no-data $> 25\%$ $\Rightarrow$ ▲ (over-smoothing risk)

The three outcomes

✓ **Passed:** ready to fit.

▲ **Concern:** fitting is possible but not advised. Interpret with caution.

✗ **Failed:** high risk of failure or bias. Not fitted.

When the area-level model runs with some low-information areas ($\leq 25\%$), those areas are **dropped from the likelihood** and their values are filled in by the spatial random effects.

Rules from `screen_svy_model()` in the sae4health source code and Wu et al. (2025).

Sparsity check

	Areas	No data	Low info
Counties	47	0	0
Constituencies	290	13	31

Share flagged at the constituency level: **15%** for the direct and area-level checks, and **4%** for the unit-level check. So all three pass the 25% rule, but only just for an indicator like MCV1.

Rarer indicators fail sooner. With fewer events per cluster, more areas have estimates of exactly 0% or 100%, which means low information. Malaria RDT positivity or IMR at Admin-2 can easily cross 25%.

If Admin-2 fails: move up to county level, use the unit-level model, or report county-level results with a note.

Running models and tracking status

Once every selected model has been checked:

- **Run models that passed check**: the conservative, recommended option.
- **Run all selected models**: also fits the ▲ models. Take extra care when interpreting them.

A status table then shows, for each method × level:

- ✓ fitted successfully
- — not chosen for fitting
- ✗ failed the sparsity check, so not fitted

Fitting uses INLA, so most models take **seconds**. Unit-level models at Admin-2 take longest.

If a model errors out with a confusing message, check the **Model Log** tab. It shows the exact R call and the console output, so you can send it to the developers.

Advanced Options: nested models and covariates

Nested model (unit-level, Admin-2)

- Adds a **fixed effect for each Admin-1 area** (county), plus spatial smoothing *within* counties.
- Respects county-level differences that the survey was designed to capture, so it does not blur across county borders.
- Off by default. Recommended by Dong et al. (2025) when modelling at Admin-2.

Area-level covariates

1. Choose the Admin level.
2. **Download Covariate Template**: a .csv with one row per area.
3. Add columns (e.g. travel time to a facility, night-time lights, population).
4. Upload the file. The app reports which covariates were loaded.

Do not edit the area-name column.

The Model Log: from clicks to code

- Every sparsity check and model fit is logged as **reproducible R code** plus the console output.
- You can copy, download or clear the log.
- It is a great way to learn **surveyPrev**, and to document exactly what was run for a report.
- The log resets when you change country, survey or indicator.

```r
# Build cluster + admin info for the selected levels
cluster.info <- surveyPrev::clusterInfo(
  geo       = cluster.geo,
  poly.adm1 = gadm.list[['Admin-1']],
  poly.adm2 = gadm.list[['Admin-2']])
admin.info <- surveyPrev::adminInfo(
  poly.adm = gadm.list[['Admin-1']],
  admin    = 1,
  by.adm   = paste0('NAME_', 1))

# --- Area-level Model at Admin-1 (KEN, 2022, CH_VACC_C_MSL) ---
res <- surveyPrev::fhModel(
  data         = analysis.dat,
  cluster.info = cluster.info,
  admin.info   = admin.info,
  admin        = 1,
  aggregation  = TRUE,
  CI           = 0.95)
summary(res$res.admin)
```

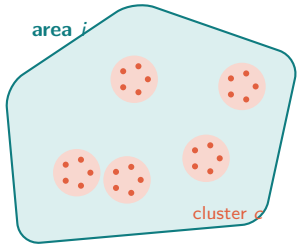
Follows the app's log template; arguments change with your selections.

The models behind the buttons

Notation

- Areas $i = 1, \ldots, m$ (e.g. 47 counties).
- Clusters c, with cluster c in area $i[c]$.
- Individuals j, with binary outcome $y_j \in \{0, 1\}$ and design weight w_j.
- Target: the **prevalence** p_i, the proportion with the outcome among the eligible population of area i.

> Example: $y_j = 1$ if child j (aged 12–23 months) had received MCV1, the first dose of measles-containing vaccine. Then p_i is MCV1 coverage in county i.



each dot is an individual j with (y_j, w_j)

Method 1 — Direct (survey-weighted) estimates

Weighted (Hájek) estimator within area i:

$$\hat{p}_i^{\,W} = \frac{\sum_{j \in S_i} w_j \, y_j}{\sum_{j \in S_i} w_j}$$

- The design-based variance D_i comes from the **survey** package, accounting for strata and clusters.
- Intervals are built on the **logit scale** and transformed back:

$$V_i = \frac{D_i}{\left(\hat{p}_i^{\,W}\right)^2 \left(1 - \hat{p}_i^{\,W}\right)^2}$$

$$\text{95\% CI: } \operatorname{expit}\left(\operatorname{logit} \hat{p}_i^{\,W} \pm 1.96 \sqrt{V_i}\right)$$

Pros

Design-consistent, transparent, and matches the DHS report at national level.

Cons

Very noisy with few clusters. Undefined with no data. Has no valid variance when $\hat{p} = 0$ or 1.

Use for: **national** and, if precise enough, **county** level.

Method 2 — Area-level (Fay–Herriot) model

Model the direct estimates themselves:

$$\operatorname{logit}\big(\hat{p}_i^{\,W}\big) \mid \lambda_i \sim \mathrm{N}\big(\lambda_i,\ V_i\big) \qquad \text{sampling model}$$

$$\lambda_i = \alpha + \mathbf{x}_i^{\top}\boldsymbol{\beta} + b_i \qquad \text{linking model}$$

- $p_i = \operatorname{expit}(\lambda_i)$ is the latent true prevalence.
- V_i is the design-based variance, treated as **known**.
- b_i are **BYM2** spatial random effects (next slide), and $\mathbf{x}_i$ optional area-level covariates.

Why it is a favourite

Because it starts from weighted estimates and their variances, it **acknowledges the survey design**. It is often the most reliable choice at the stratification level (counties).

Watch out

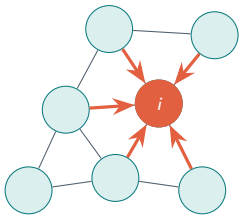
It needs a usable direct estimate in most areas. At Admin-2 the variances V_i can be unstable or near zero.

Spatial smoothing with BYM2

$$b_i \;=\; \sigma\left(\sqrt{1-\phi}\,e_i \;+\; \sqrt{\phi}\,s_i^{\star}\right)$$

- $e_i \sim \mathrm{N}(0,1)$: **unstructured**, area-specific noise.
- $s^{\star}$: **scaled ICAR**. Each area is pulled toward the average of its **neighbours** (areas that share a border).
- σ: total standard deviation of the area effects.
- $\phi \in [0,1]$: share of that variation that is **spatially structured**.

PC priors (default): $\Pr(\sigma > 1) = 0.01$ and $\Pr(\phi > 0.5) = 2/3$. Intercept: $\alpha \sim \mathrm{N}(0, 1000)$.



Under ICAR, area *i*'s effect is centred on the mean of its neighbours' effects

Riebler et al. (2016); Simpson et al. (2017). Islands with no neighbours need special handling in the adjacency graph.

Method 3 — Unit-level (cluster-level) model

Model the counts in each cluster:

$$Y_c \mid p_c, d \sim \mathrm{BetaBinomial}(n_c,\ p_c,\ d)$$
$$\mathrm{logit}(p_c) = \alpha + \mathbf{x}_{i[c]}^{\top}\boldsymbol{\beta} + b_{i[c]}$$

- Y_c is the number of "yes" answers among the n_c eligible people in cluster c.
- $d \in (0, 1)$ is the **overdispersion** (intra-cluster correlation). Prior: $\mathrm{logit}(d) \sim \mathrm{N}(0, 0.4)$.
- Area prevalence: $p_i = \mathrm{expit}(\alpha + \mathbf{x}_i^{\top}\boldsymbol{\beta} + b_i)$.

Why use it

It works with **very sparse data**: it does not need a direct estimate in each area, and areas with 1–2 clusters still contribute.

Current limitation

The app's version is **unstratified**: it ignores the survey design (weights and the urban/rural strata). If urban areas are over-sampled *and* the indicator differs by urban/rural, estimates can be **biased**.

Nested model for Admin-2

$$\text{logit}(p_c) = \underbrace{\alpha_{a[c]}}_{\text{county fixed effect}} + \underbrace{b_{i[c]}}_{\text{Admin-2 smoothing}}$$

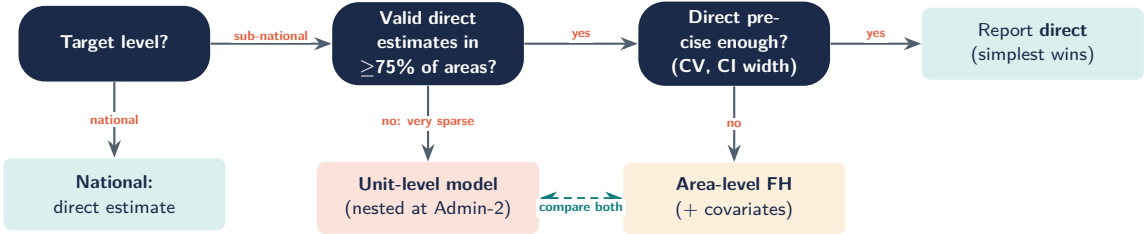
- $a[c]$ is the **county** (Admin-1 area) that contains cluster c.
- Each county gets its own level. The data support this, because counties are the strata.
- Constituencies are then smoothed **within** counties.
- Available as an *Advanced Option* for unit-level models at Admin-2 and below.

KDHS 2022 MCV1: formal comparisons (likelihood ratio test, WAIC) favoured the **nested** model over the non-nested and interaction alternatives.

Over-smoothing check: the between-Admin-2 variance was 0.0161 under the nested model vs 0.0122 when county-level FH results were simply copied down to Admin-2. So the Admin-2 model does keep real sub-county variation.

Dong et al. (2025), arXiv:2504.16435, Kenya case study.

Which model, when? A decision guide



The guide follows Dong et al. (2025): *default to simpler models if they give sufficiently precise estimates*; use FH when direct variances are usable and cluster-level models when data are very sparse. Fit several models and compare them before deciding.

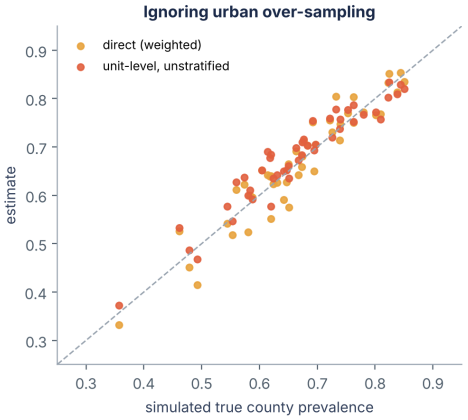
What shrinkage looks like



- Each dot is a constituency. The y-axis shows how far the model moved the estimate away from the direct one.
- Areas with **few clusters and large direct SEs** (dark dots) move the most.
- Areas with many clusters barely move.

Shrinkage is a *feature*, but too much of it hides real differences. Check with scatter plots and variance comparisons, as in the nested-model check above.

A known pitfall: ignoring urban over-sampling



- In KDHS 2022, urban clusters are over-sampled in most counties.
- An **unstratified** unit-level model treats the sample as if it reflected the population. So it **over-weights urban children** and drifts toward urban coverage levels.
- For MCV1, the average difference between unstratified and stratified estimates is [+1.4] percentage points.

The fix: stratified models with separate urban/rural levels, aggregated with population urban fractions. This is on the roadmap for the app and already available in **surveyPrev**.

Reading the results: a Kenya walk-through

About the example results

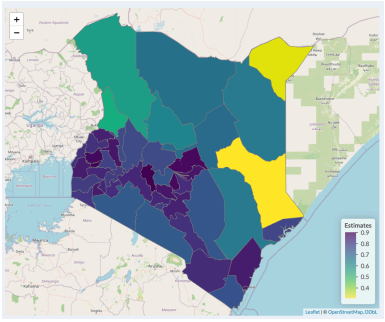
The maps and plots in this section come from the app, run on the **Kenya DHS 2022** with indicator `CH_VACC_C_MSL` (MCV1).

- Survey design: 1,692 clusters, 92 strata, about 36 clusters per county, with urban over-sampling.
- National MCV1 coverage: about **89%** (KDHS 2022 final report).
- Direct, area-level (BYM2) and unit-level estimates were fitted at county level. At constituency level, the area-level model could not be fitted, so only direct and unit-level estimates are shown.

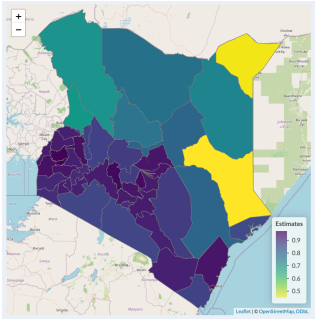
During the hands-on session, reproduce these figures yourself with the app and check that you get the same numbers.

Prevalence Map: county level, three methods

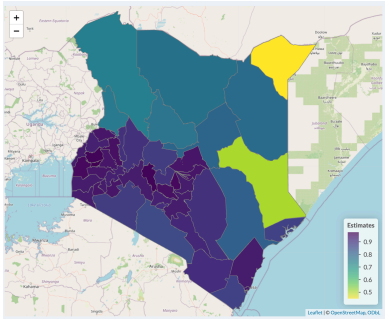
Direct



Area-level



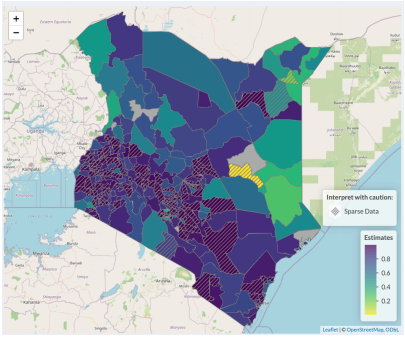
Unit-level



- At county level the three methods **agree closely**. Counties are the strata, so data are adequate.
- The direct estimates already look usable at this level.
- Controls: *Admin level*, *Map region* (zoom to one county's sub-areas), *Method* and *Statistic*.
- Toggle between the interactive (leaflet) and static maps, and download either.

Prevalence Map: constituency level

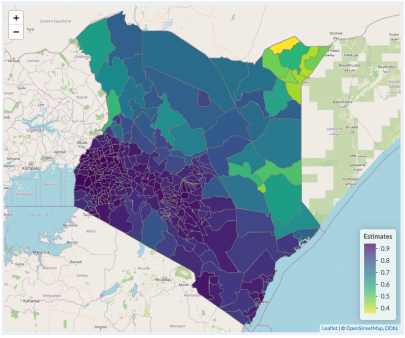
Direct



Area-level

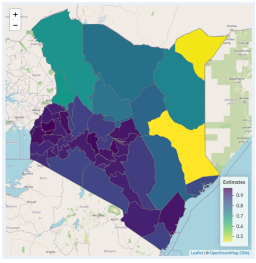
The area-level model **could not be fitted** at constituency level for MCV1.

Unit-level

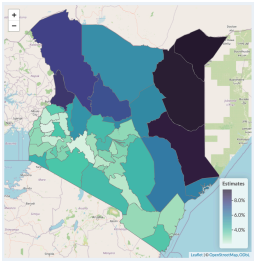


- The **direct** map is noisy. Neighbouring areas jump between extremes, and grey areas have no data.
- For MCV1, [%] of Admin-2 direct intervals are wider than the widest county interval ([width]).

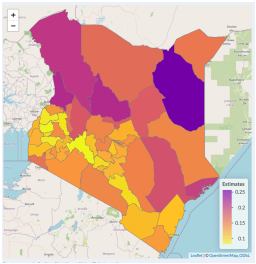
The four key statistics



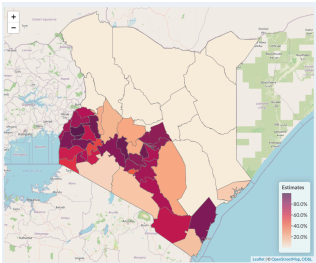
Mean
posterior mean of the prevalence (the point estimate)



CV $= 100 \times \text{SD}/\text{mean}$
relative precision. Lower is better.



95% interval width
absolute uncertainty in prevalence units



Exceedance
$\Pr(p_i > p_0 \mid \text{data})$ for a target p_0 you choose with a slider

Area-level model at county level, threshold p_0 =0.95, e.g. 95%.

Interpreting CV and interval width

- **CV is relative:** the same SD gives a larger CV when prevalence is low. Rare indicators always look worse.
- **Interval width is absolute:** it is easier to explain to policy makers ("we are 95% sure MCV1 coverage is between 82% and 93%").
- A common rule of thumb comes from Statistics Canada: $\mathrm{CV} \lesssim 16.6\%$ is usable without restriction, and larger CVs need caution or suppression.
- Look at **both** maps *next to* the mean map. The Map Comparison tab does exactly this.

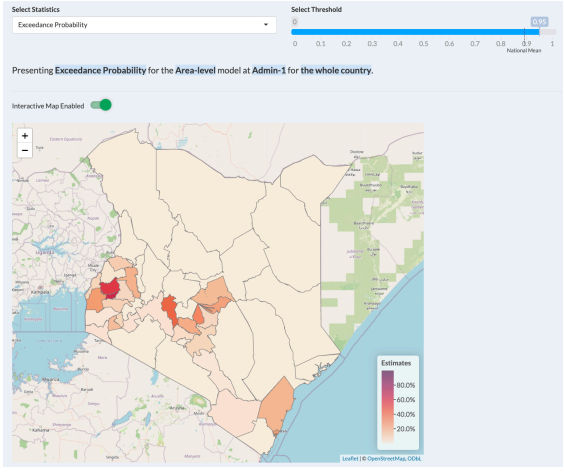
$\mathrm{CV} \leq 16.6\%$	16.6–33.3%	$> 33.3\%$
use as is	use with caution	too unreliable

These cut-offs are conventions, not laws. The app shows the numbers and leaves the threshold to you and your stakeholders.

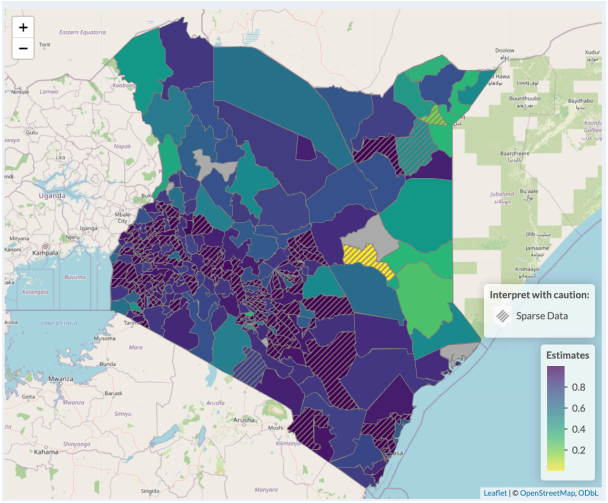
Statistics Canada guidance as cited in Dong et al. (2025).

Exceedance probabilities for policy targets

- For MCV1, the **95%** coverage benchmark for measles elimination is a natural thresholds p_0.
- $\Pr(p_i > p_0) \approx 1$ means the area is almost certainly **above** target.
- $\Pr(p_i > p_0) \approx 0$ means it is almost certainly **below** target, a priority area.
- Values near 0.5 mean the data cannot tell. **Do not rank these areas.**



Grey and hatched areas: how the maps flag sparse data



- **Grey** ("No Data"): no clusters, so no direct estimate.
- **Hatched** ("Sparse Data"): an estimate exists, but its uncertainty is not available, typically $\hat{p} = 0\%$ or 100%.
- Hover in the interactive map to see warnings for each area.
- Model-based maps fill these areas in, but their *uncertainty* will be large. Always show it alongside.

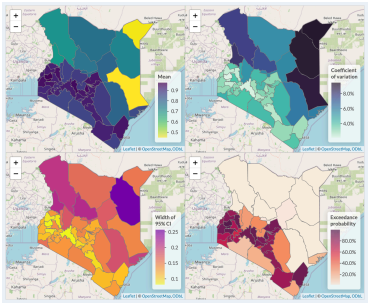
Map Comparison

Two modes:

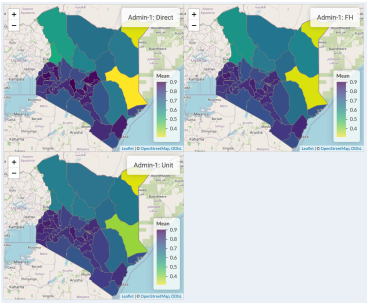
1. **Several statistics, one model:** mean, CV, CI width and exceedance side by side.
2. **One statistic, several models:** e.g. the mean from direct vs area-level vs unit-level, on a common colour scale.

> Mode 2 is the fastest way to spot **over-smoothing** (a model map that looks too flat) or **noise** (a direct map that looks like a patchwork).

1. Statistics (area-level model)



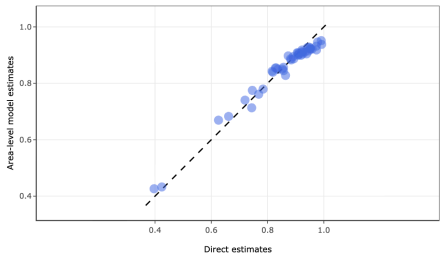
2. Mean from several models



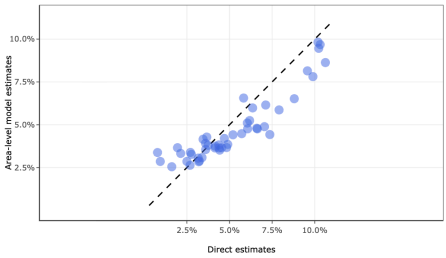
County level.

Scatter Plot: comparing two methods area by area

Estimates



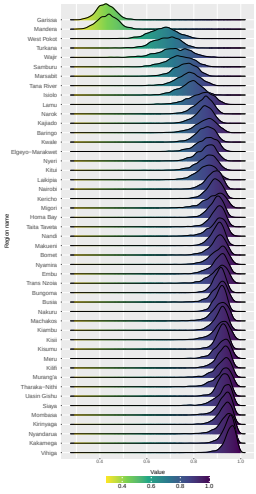
Coefficient of variation



- Choose a method for the x-axis and the y-axis, and a statistic.
- **Estimates:** points near the diagonal mean the methods agree. A flatter cloud than the diagonal means shrinkage toward the mean.
- **CV:** points **below** the diagonal mean the model is more precise than the direct estimate.
- An interactive (plotly) version lets you hover to identify areas.

Direct (x-axis) vs area-level model (y-axis), county level.

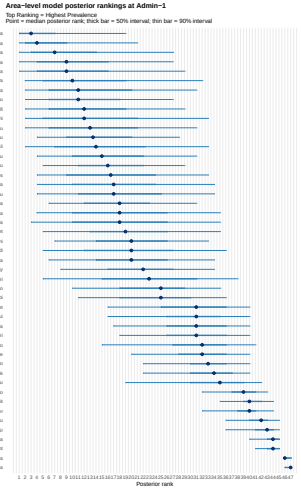
Ridge Plot: full posterior distributions

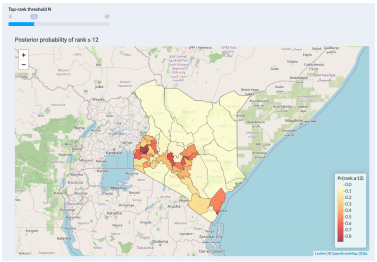


Area-level model, 47 counties.

- Shows the **whole posterior distribution**, not just a point and an interval.
- Options: all Admin-1 areas, the Admin-2 areas within a chosen county, or the k **highest and lowest** areas.
- **Overlapping ridges** mean the difference between two areas may not be meaningful.
- For MCV1, most counties fall between about [30 %] and [100 %]. Where they overlap heavily, a strict ranking would be misleading.

Ranking Plot

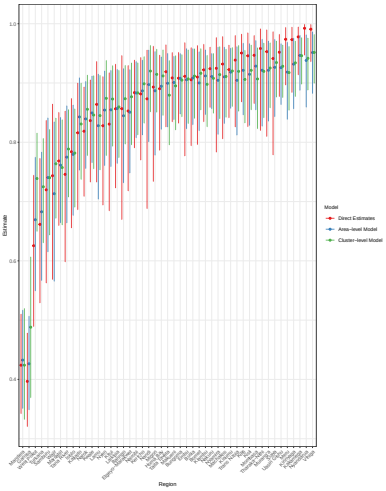




- **Caterpillar plot:** median posterior rank (point), 50% interval (thick bar), 90% interval (thin bar). Rank 1 = highest prevalence.
- **Mean posterior rank map**, and **Pr(rank $\leq N$)**: the probability of being in the top N, with N set by a slider.
- Only the extremes are clearly ranked. The middle of the table is a statistical tie.

Area-level model, 47 counties.

Interval Comparison



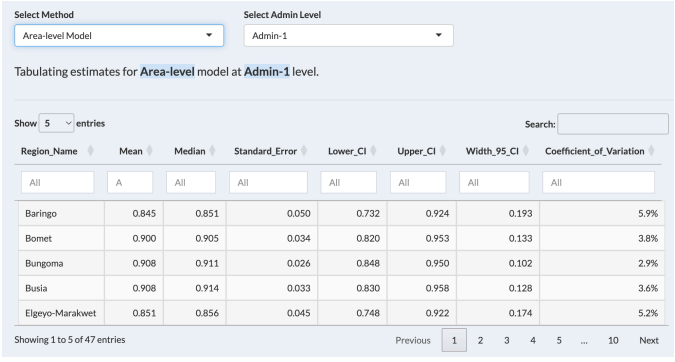
All 47 counties, county level, sorted by estimate.

- Point estimates and 95% intervals for every fitted model, area by area: direct, area-level and unit-level ("Cluster-level Model" in the legend).
- Model intervals are typically **narrower** than direct ones, and their point estimates are pulled toward the middle.
- Check whether the unit-level estimates sit systematically above or below the others. A consistent shift points to the urban over-sampling bias from the models section.

Result Tabulation

- Choose a method and Admin level to get a paginated table, one row per area.
- Columns include the **mean**, **SD/SE**, **coefficient of variation**, and **95% interval** (lower and upper).
- **Download** saves the table as **CSV** for Excel, R, Stata or QGIS.

> Join the CSV to your county shapefile in QGIS to make custom maps for county health departments.



Reports and the Tool Kit

Report Generation: one click to a PDF

The report collects your whole analysis:

1. **Summary info:** country, survey, indicator, admin levels, and number of regions.
2. **Indicator details:** DHS ID, report chapter, full definition.
3. **Consistency check:** the app's national estimate vs the DHS/MIS final report.
4. **Sample information:** clusters, sample size, events, and regions without data, overall and by level.
5. **Key-statistics maps** (mean, CV, CI width, exceedance) for each fitted model and level.

Languages

The report is available in **English, French, Spanish and Portuguese**. The language is pre-set to match the country's DHS report (English for Kenya) and can be changed.

Fit *all* the models you want in the report *before* you click **Download the Report**.

The Tool Kit

Indicators in the App

A searchable table of every supported indicator, with its definition and the recode files it needs.

DHS Indicator Dictionary

Official DHS definitions, measurement type and denominator for **4,433** DHS indicators, including ones the app does not model.

DHS Report Estimate

Official national estimates from the DHS API, used for the consistency check. The title reads "Comparing with DHS/MIS Final Report Estimates".

Before presenting any subnational map, check that the indicator is defined the way your audience expects: age range, reference period, denominator. Definitions differ between, e.g., "births in the last 2 years" and "last 5 years".

DHS vs MIS surveys

Malaria Indicator Surveys in the app

- MIS are smaller, malaria-focused household surveys run with DHS Program support. They use the same recode file structure.
- In the app, MIS surveys appear in the survey-year menu labelled `(MIS)`, e.g. `2020 (MIS)` for Kenya.
- MIS collect only a subset of DHS modules, so the indicator menu is **restricted to 58 indicators**. These were checked against DHS-published values for **KE2020MIS**.

The 58 MIS indicators cover

- Housing, water and sanitation, wealth
- Respondent education, literacy, media, internet
- Fertility; ANC visits and provider
- Fever and treatment-seeking in children
- Anaemia in children
- Malaria: ITN ownership and use, RDT positivity, fever treatment

`MIS_supported_indicators` in the sae4health source (fct_survey_type_helpers.R).

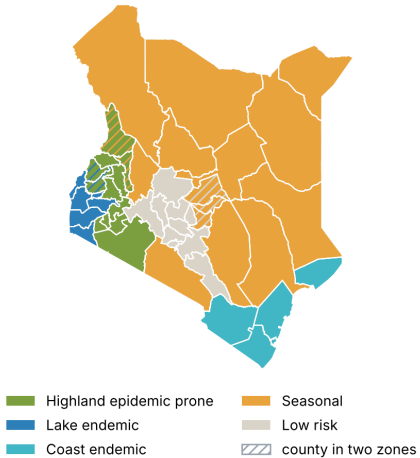
DHS vs MIS: what changes for SAE

	Kenya DHS 2022	Kenya MIS 2020
Purpose	Broad demographic and health	Malaria interventions, anaemia, malaria parasitaemia
Clusters	1,692	**301** (134 urban, 167 rural)
Households / cluster	25	30
Strata	92 (county × urban/rural)	92 (county × urban/rural), NASSEP V frame
Designed for	National, urban/rural, county (some indicators)	National, urban/rural, **5 malaria-endemicity zones**
Clusters per county	about 36	about **6** on average

With about 6 clusters per county, **county-level direct estimates from an MIS are usually too noisy**. Lean on area- or unit-level models; Admin-2 is generally out of reach.

KMIS 2020 and KDHS 2022 sample design documentation (IHSN / World Bank microdata catalogues).

Kenya MIS 2020: designed for zones, not counties



- The five malaria-endemicity zones cover the whole country. Each county falls in one or two zones.
- Seven counties span two zones (hatched): Kakamega, Bungoma, West Pokot, Elgeyo-Marakwet, Meru, Tharaka-Nithi and Embu.
- The app's Admin levels follow **administrative** boundaries, not zones. So zone-level reporting is *not* a built-in Admin level.

> Malaria is highly spatial, so smoothing helps. Always show the uncertainty maps with MIS-based county results.

Zone membership: KMIS 2020 sampling documentation.

Shapefiles and boundaries

Where the boundaries come from

GADM (default)

GADM version 4.1, retrieved with the R package `geodata`. **Preloaded** in the web version, with an OpenStreetMap basemap underneath.

WHO GIS Hub

The `run_app(version = 'DHS-WHO')` variant asks you to upload Admin-1 and Admin-2 shapefiles (.zip) from the WHO GIS Hub, and uses no basemap.

Custom boundaries

For national or ministry boundaries (e.g. official sub-counties), contact the developers. The app needs consistent names and adjacency.

Why boundaries matter: they decide **which area each cluster belongs to**, the **neighbour graph** for spatial smoothing, and the **names** in every table and report.

What's next?

Recently added and coming next

Recently added

- MIS survey support (58 indicators)
- Ranking and interval-comparison plots
- Child mortality indicators (U5MR, IMR) with mortality-specific settings
- Reports in 4 languages
- Model Log with reproducible R code
- Nested models and area-level covariates

On the roadmap

- **Stratified** unit-level models (urban/rural), aggregated with population fractions
- MICS integration (a Nigeria MICS app already exists)
- Space-time models for composite indicators (U5MR, fertility)
- Continuous spatial (SPDE) models and unit-level covariates
- Cross-validation and model checking tools

Result-only websites such as https://sae4lmic.stat.uw.edu/ publish pre-computed, model-checked estimates for selected countries and indicators, for users who need numbers rather than tools.

Resources

- **Web app:** https://rsc.stat.washington.edu/sae4health/
- **Documentation, videos, methods, gallery:** https://sae4health.stat.uw.edu/
- **Source and issues:** https://github.com/wu-thomas/sae4health
- **surveyPrev:** https://github.com/richardli/surveyPrev
- **SUMMER:** https://richardli.github.io/SUMMER/
- **DHS data access:** https://dhsprogram.com/data/

Key references

- Wu, Dong, Xu, Li & Wakefield (2025). *sae4health: An R Shiny application for SAE in LMICs.* arXiv:2505.01467.
- Dong et al. (2025). *Toward a principled workflow for prevalence mapping using household survey data.* arXiv:2504.16435.
- Wakefield, Okonek & Pedersen (2020). Small area estimation for disease prevalence mapping. *Int. Stat. Review.*
- Riebler et al. (2016). An intuitive Bayesian spatial model for disease mapping (BYM2). *SMMR.*
- KNBS & ICF (2023). *Kenya DHS 2022.*

Asante sana!

Questions, then let's open the app.



https://sae4health.stat.uw.edu/