Bayes Bayesian Biostatistics Leiden, NL 2026

21–23 October 2026

Leiden – The Netherlands



Martina Del Gaudio, Pietro Fusco, Mohd Rashid Khan, Danila Azzolina

Biostatistics and Clinical Trial Methodology Unit, Clinical Research Center DEMeTra, Department of Translational Medical Sciences, University of Naples Federico II

E-mail: martina.delgaudio@unina.it

BIOSTATISTICA DEMeTra

DIPARTIMENTO DI SCIENZE MEDICHE TRASLAZIONALI

DIPARTIMENTO DI ECCELLENZA 2023 - 2027

UNIVERSITÀ DEGLI STUDI DI NAPOLI FEDERICO II

BAYESIAN FEDERATED INFERENCE FOR A CLUSTER-RANDOMIZED TRIAL

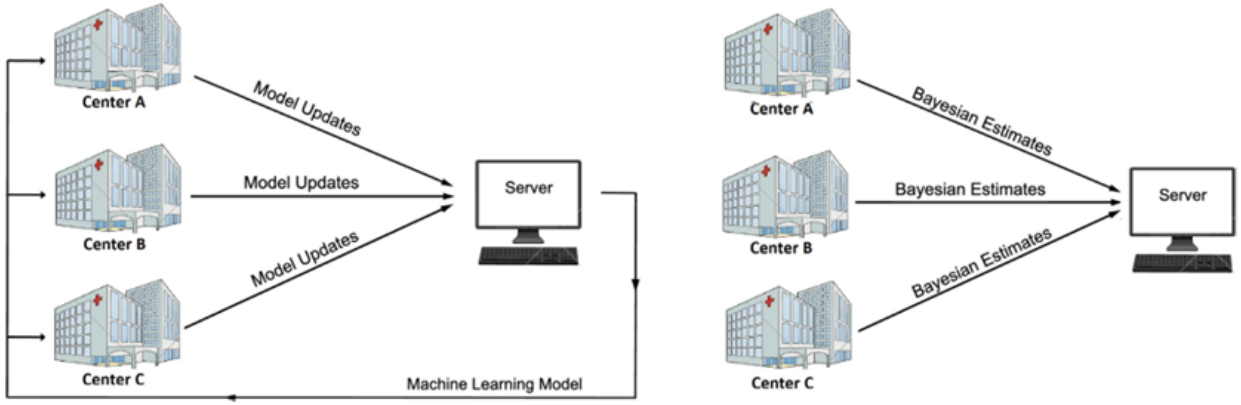
EIGHT-YEAR FOLLOW-UP OF A MATERNAL EDUCATION TRIAL

Why Bayesian models are more efficient for small, clustered, non-shareable data

The statistical problem

Data cannot always be pooled

Multicentre and cluster-randomized studies often have few patients per site and privacy or governance rules prevent building a single combined database, yet the pooled analysis is the benchmark.



Research question

Can we recover the pooled-data inference without pooling the data?

- **Efficiency through priors**

Gaussian priors act as ridge-type shrinkage: stable local estimates even when a site is small relative to the number of parameters.

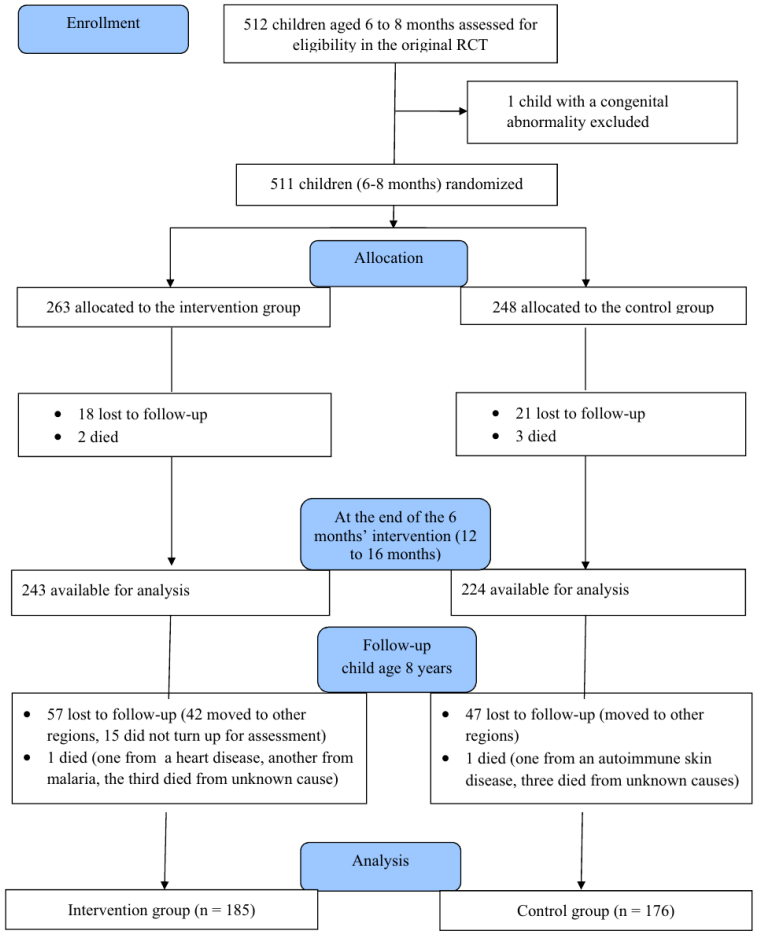
- **Full posterior, not just point estimates**

Uncertainty is propagated explicitly: credible intervals and direct probability statements on the effect.

- **One-shot communication**

Each site shares its local MAP and curvature once, instead of the many iterative cycles of federated learning.

CHNUDEV: a cluster-randomized trial with a natural multi-site structure



URE 1
SORT flowchart. The flowchart illustrates the inclusion into the original cluster-randomized trial and current follow-up study.

ing in, holding, and using information immediately and er some time. Higher visual processing scores indicated compared with control children. Furthermore, the int vention group had acquired more knowledge and th

Two-arm cluster-randomized controlled trial

CLUSTERS

10 clusters in 2 districts: Kabale (6) + Kisoro (4), SW Uganda

ENROLMENT 2013–14

511 mother–child pairs, children aged 6–8 months

ARMS

6-month maternal education package vs routine care

8-YEAR FOLLOW-UP

361 children (71%): intervention 185 · control 176

OUTCOMES

KABC-II FCI (continuous) + 5 subscales · 5 TOVA z-scores

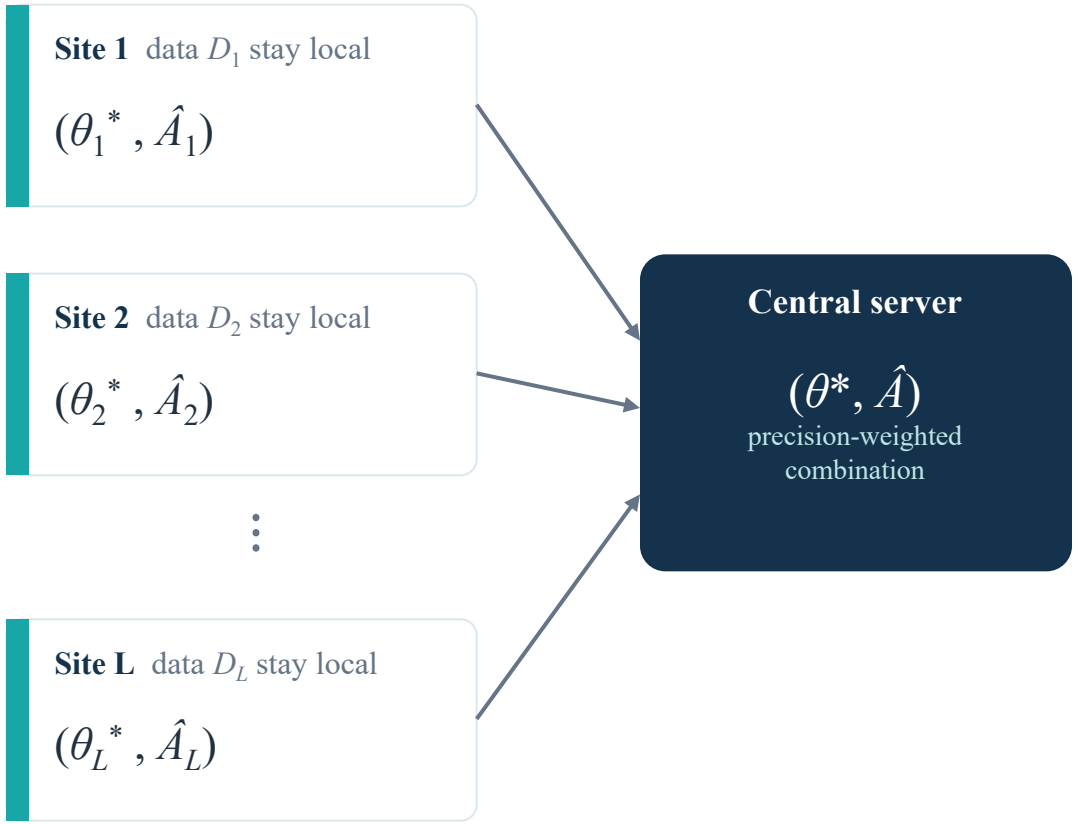
ORIGINAL ANALYSIS

Frequentist cluster-adjusted mixed models, BH-FDR

Data structure: children nested in 10 clusters → each cluster = one data site

Bayesian federated inference

One-shot aggregation of local posteriors (Jonker, Pazira & Coolen, Stat Med 2024)



Only summaries leave the site · single round · no individual data shared

01 Exact posterior decomposition

In a Bayesian analysis based on any data set $D = \{y_1, \ldots, y_n\}$ from a conditional density function $y \mid \boldsymbol{\theta} \to p(y \mid \boldsymbol{\theta})$ with prior density $\boldsymbol{\theta} \to p(\boldsymbol{\theta})$, the posterior density equals:

$$p(\boldsymbol{\theta}|\boldsymbol{D}) = \frac{p(\boldsymbol{D}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{Z(\boldsymbol{D})}, \qquad Z(\boldsymbol{D}) = \int p(\boldsymbol{D}|\boldsymbol{\theta})p(\boldsymbol{\theta})\mathrm{d}\boldsymbol{\theta}, \qquad p(\boldsymbol{D}|\boldsymbol{\theta}) = \prod_{i=1}^{n} p(y_i|\boldsymbol{\theta})$$

Where $Z(D)$ denotes the normalizing constant of the posterior distribution

02 Local posterior probabilities

$$p(\theta|\mathcal{D}_n) \approx q_n(\theta) \equiv \mathcal{N}(\theta|\mu_n, \Sigma_n)$$

We assume that each posterior $p(\theta \mid D_n)$ follows a multivariate Gaussian distribution $q_n(\theta)$

03 Closed-form aggregation (priors $N(0, \Lambda^{-1})$)

$$p(\boldsymbol{\theta}|\boldsymbol{D}) \approx \left(\frac{\det(\hat{\mathbf{A}})}{(2\pi)^d}\right)^{\frac{1}{2}} \exp\left\{-\frac{1}{2}(\boldsymbol{\theta}-\hat{\boldsymbol{\theta}})^t\hat{\mathbf{A}}(\boldsymbol{\theta}-\hat{\boldsymbol{\theta}})\right\}$$

Where $d = \dim(\boldsymbol{\theta})$ and $\det(\hat{A})$ denotes the determinant of the matrix $\hat{A}$

BFI applied to the trial

Local model

Logit link

FCI below average (yes/no)

Heterogeneity

site intercepts $\alpha\ell$ handled as nuisance parameters (Jonker Sec. 2.4)

Priors & exactness

$\beta, \gamma \sim N(0, \lambda^{-1} I)$

weakly informative, ridge-type shrinkage; flat prior on $\alpha\ell$

Gaussian outcomes

log-posterior exactly quadratic → BFI = pooled posterior

Cluster randomization

Arm constant within cluster

β not identifiable from a single subcounty

BFI pools information

curvatures $\hat{A}\ell$, not point estimates: every cluster still contributes

Design choice

site intercepts confounded with arm → sites spanning both arms, or model between-cluster variance

Evaluation: agreement of BFI with the pooled analysis of the same model — MSE of β, of posterior SDs and of predictions

Pooled benchmark at 8 years and what the Bayesian federated approach adds

+14 points

KABC-II FCI difference

95% CI 12–16 · ≈ 0.9 SD of the scale

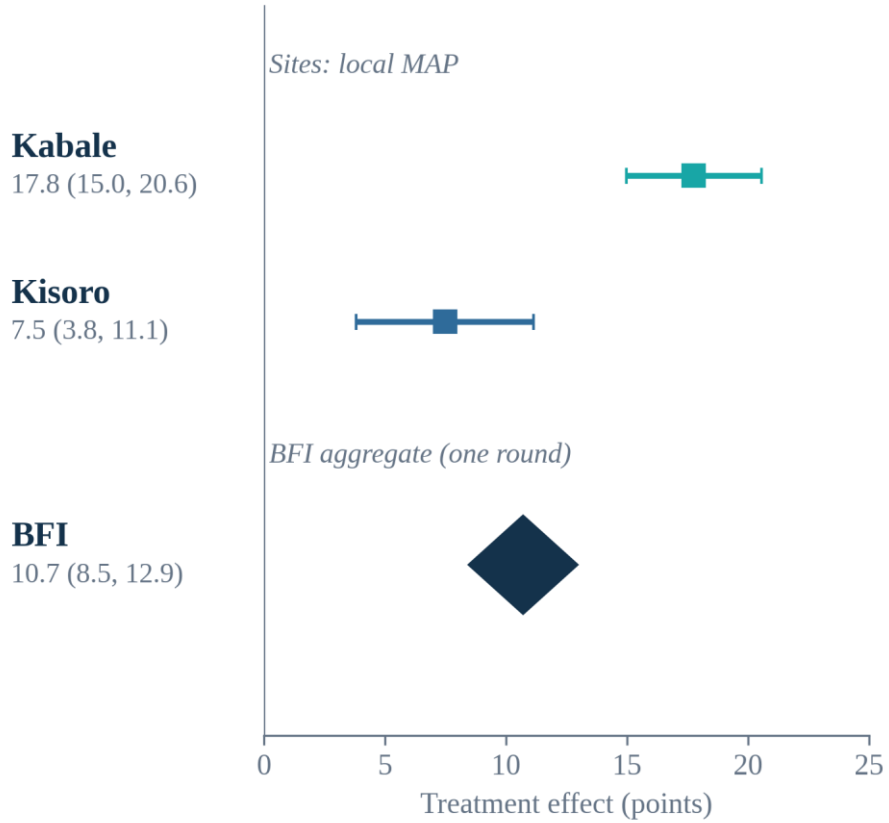
72% vs 97%

FCI below average · difference 25 pp (95% CI 18–32)

5 / 5

TOVA variables favour the intervention (visual and auditory)

Federated estimates are compared against this benchmark



Our data: treatment effect on the continuous outcome, 2 sites (Kabale, Kisoro)

Conclusions

- The maternal education intervention improved children's cognitive performance (KABC-II) at 8 years
- Large, consistent benefit: +14 FCI points (≈ 0.9 SD); all 5 TOVA variables also favour the intervention
- Effect seen at both sites (17.8 and 7.5); BFI aggregate 10.7 (95% CI 8.5–12.9) from one round of site summaries
- Caveats: heterogeneity between sites; the package prevents attribution to single components

Benchmark: published cluster-adjusted pooled analysis (Kakwangire et al., Pediatrics 2024;153:e2023063352). Method: Jonker et al., Stat Med 2024;43:2421–2438.

Thank you for the attention!