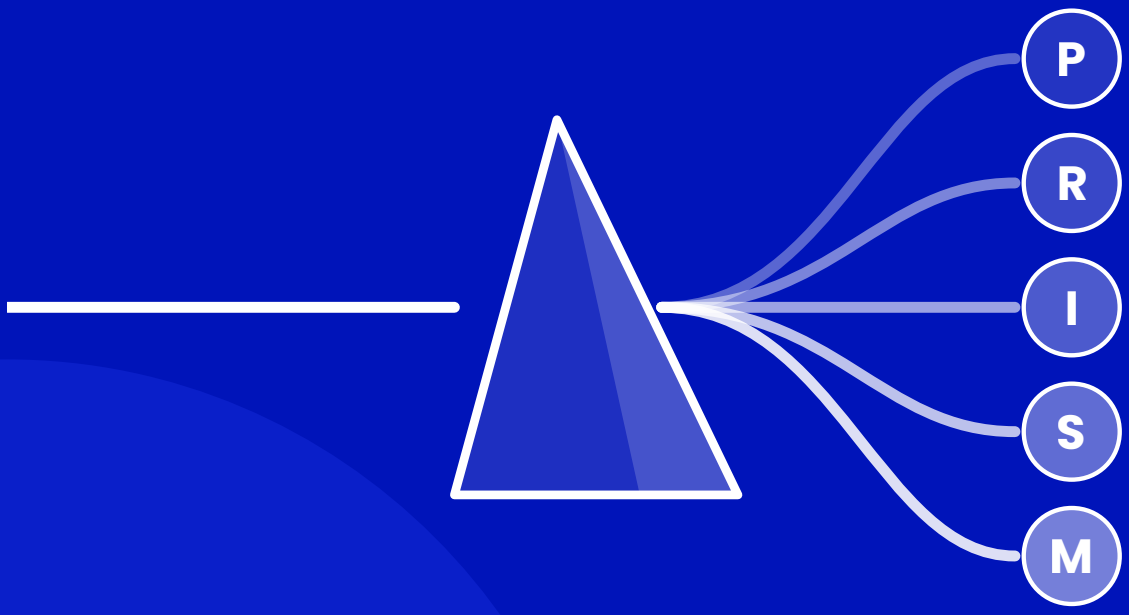
EARLY-WARNING SURVEILLANC

PRISM-EW™ Whitepaper

Probabilistic Rt-Integrated Surveillance Model for Early Warning — version 1.1



AUTHOR:

Oluwatobi Lucas Akinbode, Statistical Data Scientist · AI Researcher · Predictive Healthcare Analytics Specialist

Executive summary

PRISM-EW™ turns routine daily case counts into one explainable early-warning score for each region: the **Outbreak Pressure Index (OPI™)**, from 0 to 100, with four alert levels. It answers the question surveillance teams ask every week: *is an outbreak building here, how sure are we, and why?*

Case counts show what has already happened. They do not show whether transmission is accelerating, whether current activity is unusual for the season, or what is likely next. Those answers usually come from separate tools (an Rt estimate, a forecast, an aberration detector) that must be read side by side. PRISM-EW™ joins them in one five-stage pipeline:

- **P** – **Preprocess:** clean the counts and remove weekday reporting effects.
- **R** – **Reproduction:** estimate the effective reproduction number (Rt) with a Bayesian credible interval.
- **I** – **Incidence:** forecast the next 21 days with a three-model probabilistic ensemble.
- **S** – **Signal:** run standard outbreak detectors (EARS, CUSUM, seasonal baseline).
- **M** – **Multi-signal fusion:** combine five signals into the OPI™, with a point-by-point explanation of every score.

A companion **risk-factor module** examines how individual-level behaviours relate to each other, using Spearman and distance correlation, false-discovery control and interactive networks. It builds on HIV risk-behaviour research at Duke University.

STATUS (OCTOBER 2026):

- A working web application runs every stage in the user's browser, so uploaded data are never sent to a server. It is deployable on Firebase Hosting.
- All methods are implemented from first principles and checked by 29 automated tests. They match Python reference libraries (scipy, dcor) to 10 decimal places.
- On synthetic surveillance data, the ensemble had the lowest average forecast error of all models, and planted outbreaks reached the High alert level within days.
- Validation on real historical outbreak data is the next milestone (Section 8).

> PRISM-EW™ does not claim to have invented Rt estimation, EARS, CUSUM or the individual forecasting models. Its contribution is the integrated, explainable architecture that combines them into one operational early-warning signal.

1. Introduction

1.1 Context

Public-health agencies monitor infectious diseases across counties, districts, states and regions, usually as daily or weekly case counts. Those counts are noisy:

- Reports dip at weekends and catch up on Mondays.
- Days go missing.
- Seasonal waves raise baseline levels every year.
- The most recent days carry the most uncertainty.

PRISM-EW™ treats early warning as a problem of **predictive surveillance intelligence**. It combines current activity, transmission dynamics, statistical abnormality, historical burden and near-term risk into one signal that a surveillance team can act on.

1.2 Stakes

- **Delayed detection:** an outbreak may only become obvious after case counts have already risen sharply.
- **Reporting noise:** missing data and reporting cycles distort trends.
- **Fragmented indicators:** Rt, forecasts, anomaly signals and historical burden are often read separately.
- **Reactive decisions:** action depends on past cases rather than likely future pressure.
- **Weak prioritisation:** with limited staff, teams need to know which regions need attention first.

2. The problem PRISM-EW™ solves

2.1 Common failure modes

- **Reactive monitoring:** outbreaks are identified after substantial increases have already occurred.
- **Uncorrected data:** missing observations and weekday reporting patterns produce misleading trends.
- **Isolated transmission measures:** Rt is not read together with burden, forecasts and detection signals.
- **Limited foresight:** historical counts say little about likely activity over the next few weeks.
- **No unified warning:** many signals exist, but there is no transparent method for ranking regions by outbreak pressure.

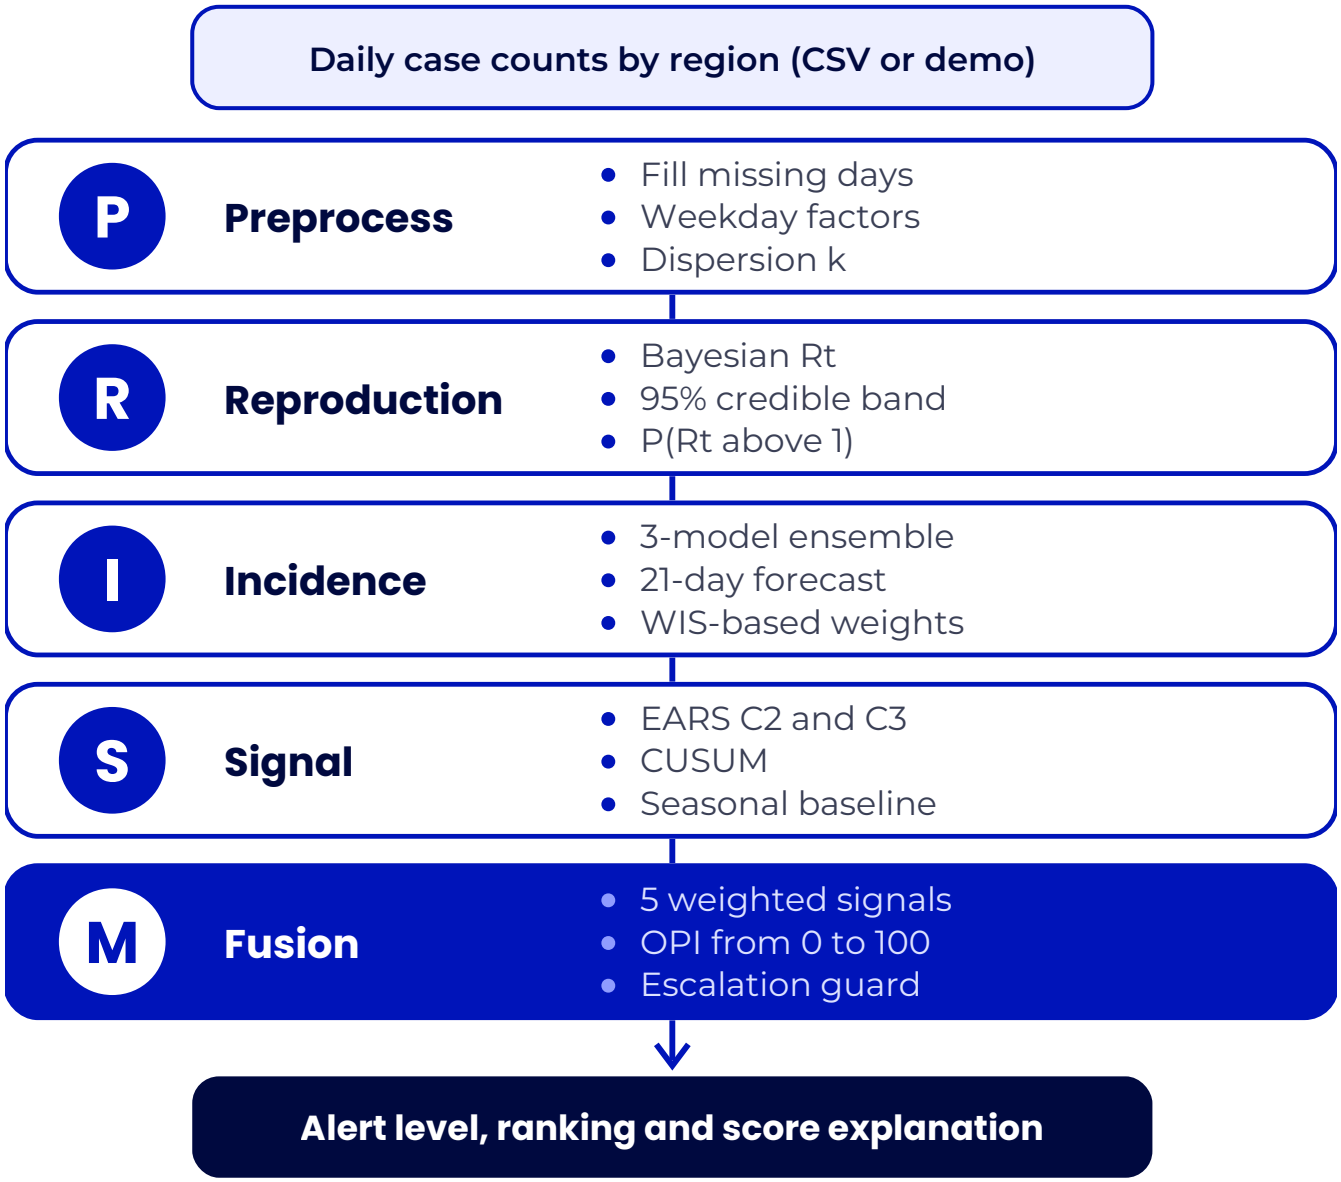
2.2 Root causes

- Over-reliance on raw case counts.
- Limited integration of established statistical surveillance methods.
- Weak treatment of uncertainty.
- Forecasting and outbreak detection run as separate systems.
- Little comparison of current activity with the same season in past years.
- Difficulty turning statistical outputs into operational alerts.

3. Architecture

Every region, and the total across all regions, passes through the same five stages in order. Each stage hands its results to the next and contributes at least one signal to the final score.

Case counts pass through five stages to become one alert



PRISM-EW™ architecture · five stages from case counts to alert

The stage initials spell PRISM: **P**reprocess, **R**eproduction, **I**ncidence, **S**ignal and **M**ulti-signal fusion. Section 4 specifies the method, parameters and outputs of each stage.

4. Stage specifications

Each stage below lists its objective, its method with default parameters, its outputs and its decision value. All parameters shown are defaults and can be changed in the application.

4.1 P — Preprocess: surveillance data reliability

OBJECTIVE: prepare case counts for reliable analysis.

METHOD:

- **Missing days:** filled by linear interpolation; a missing tail carries the last 7-day mean forward.
- **Weekday reporting factors:** each day's count divided by its centred 7-day average, averaged over the last 12 weeks and normalised to a mean of 1. Dividing by these factors gives the adjusted series used by every later stage.
- **Overdispersion:** the negative-binomial dispersion k, estimated by method of moments. It measures how much more variable the counts are than pure chance.

OUTPUTS: cleaned series, 7-day averages, rates per 100,000 people, weekday factors and dispersion.

DECISION VALUE: a weekend dip or a missing report is not mistaken for a decline.

4.2 R — Reproduction: transmission intelligence

OBJECTIVE: determine whether transmission is expanding or contracting.

METHOD: the Bayesian renewal-equation estimator of Cori et al. (2013).

SERIAL INTERVAL: discretised gamma distribution (default mean 4.7 days, SD 2.9 days).

PRIOR AND WINDOW: a Gamma(1, 5) prior on Rt and a 7-day sliding window, giving a conjugate posterior.

DISEASE PRESETS: influenza-like (3.6 / 1.6 days), COVID-19 Omicron-era (4.7 / 2.9), RSV (7.5 / 2.1), measles (11.7 / 2.0) and cholera (5.0 / 3.0).

$$R_t \mid \text{data} \sim \text{Gamma}\left(1 + \Sigma\, I_k,\ 1 / \left(1/5 + \Sigma\, \Lambda_k\right)\right)$$

, summing over the last $\tau = 7$ days, where $\Lambda_k = \Sigma_s w_s \cdot I_{k-s}$ is the infection pressure from earlier cases weighted by the serial interval w_s.

OUTPUTS: daily Rt posterior mean, 95% credible interval, and P(Rt > 1).

DECISION VALUE: probabilistic evidence of whether spread is increasing, with its uncertainty stated.

4.3 I — Incidence: probabilistic forecasting

OBJECTIVE: estimate likely incidence over the next 21 days.

METHOD: three structurally different models, each simulating 800 Monte Carlo paths.

- **Renewal model:** Rt drawn from the posterior evolves as a mean-reverting random walk on the log scale; cases are generated through the renewal equation with negative-binomial noise.
- **Damped Holt model:** exponential smoothing with a damped trend on log(1 + cases); parameters are chosen by grid search.
- **Quasi-Poisson trend:** a log-linear Poisson regression on the last 28 days, with parameter uncertainty and a trend damped by 3% per day.

Paths are returned to the reported scale by reapplying the weekday factors. The models are **blended, not selected**. Each model's weight is inversely proportional to its weighted interval score (WIS, a standard error measure for probabilistic forecasts) in a rolling-origin backtest: six weekly forecast origins, scored at 1–14 days ahead. Every weight has a floor of 10%.

OUTPUTS: median forecast, 50% and 95% prediction intervals, accuracy scores for each model, and 14-day surge probability.

DECISION VALUE: teams can anticipate pressure, and each region shows how much to trust its forecast.

4.4 S — Signal: statistical outbreak detection

OBJECTIVE: identify statistically unusual activity.

METHOD: four detectors, each with an alarm threshold.

Detector	What it compares	Alarm when
EARS C2	Today against days t−9 to t−3	C2 > 3
EARS C3	Sum of the last three C2 excesses above 1	C3 > 2
CUSUM	Running sum of C2 above a reference of 1 SD; catches slow, sustained rises	> 5
Seasonal baseline (Farrington-style)	This week's total against the same ±3-week window in up to 5 prior years	z > 2.33

Each detector is scaled by its threshold. The results are fused into one score from 0 to 1: 0.6 × the strongest detector + 0.25 × the second + 0.15 × the third.

OUTPUTS: daily alarm flags and a fused detection score.

DECISION VALUE: separates unusual increases from routine fluctuation and expected seasonal activity.

4.5 M — Multi-signal fusion: outbreak pressure intelligence

OBJECTIVE: turn many epidemiological signals into one transparent, rankable warning.

METHOD: a weighted, fully decomposable index, specified in Section 5.

OUTPUTS: the OPI™ (0–100), alert level, regional ranking, and each signal's point contribution.

DECISION VALUE: one measure for prioritising regions, where every point can be traced to its cause.

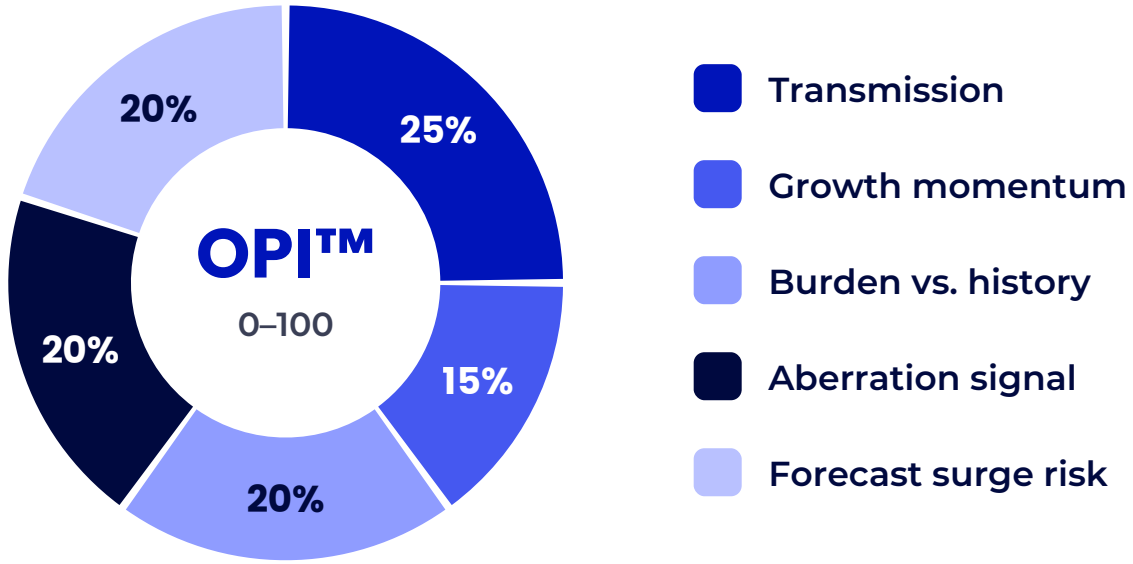
5. The Outbreak Pressure Index (OPI™)

The OPI™ is a weighted sum of five signals, each first mapped to a score between 0 and 1. Because it is linear, every point of a region's score can be traced to the signal that produced it.

$$\text{OPI} = 100 \times \Sigma\ w_i \cdot s_i$$

, summed over the five signals i = 1–5, with the weights w_i summing to 1.

5.1 Signal definitions and default weights



Signal	Question it answers	Score s_i (0–1)	Weight w_i
Transmission	Is each case causing more than one new case?	½·P(Rt > 1) + ½·clip((R̄t − 0.9) / 0.5, 0, 1); full magnitude credit at Rt = 1.4	25%
Growth momentum	How fast are cases rising this week?	logistic(r / 0.03), where r = ln($MA7_t$ / $MA7_{t-7}$) / 7; growth of 3% a day scores about 0.73	15%

Signal	Question it answers	Score s_i (0–1)	Weight w_i
Burden vs. history	Is the current level high for this region?	Percentile rank of today's 7-day average among all earlier days	20%
Aberration signal	Has a detector raised an alarm?	Highest fused detection score over the last 3 days	20%
Forecast surge risk	Will the next two weeks be clearly worse?	P(next 14 days > 1.2 × last 14 days), taken from the ensemble's simulated futures	20%

The weights are policy choices. Users can change them in the application, and they are renormalised to sum to 1 at run time.

5.2 Alert levels

Low	Guarded	Elevated	High
0–24	25–49	50–69	70–100

OPI™	Level	Suggested response
0–24	● Low	Routine monitoring
25–49	● Guarded	Closer watch; review data quality
50–69	◆ Elevated	Investigate and prepare
70–100	■ High	Act: escalate to response leads

Each level is shown with an icon and a text label as well as a colour, so it is readable without colour vision and in print.

5.3 Escalation guard

A region with low historical burden can have a fast-starting outbreak that the burden signal underweights. The guard is one explicit rule:

> ! If any detector is in alarm **and** P(Rt > 1) > 90%, the OPI™ is raised to at least **50 (Elevated)**.

When the guard fires, the application says so beside the region's score.

5.4 Nowcast OPI™ trajectory

For auditing, the application recomputes the four signals that need no forecast for each of the last 120 days. This gives a retrospective OPI™ trend, so analysts can see when pressure began to build.

5.5 Worked example: Riverside (synthetic demonstration region)

Signal	Evidence	Score	Weight	Points
Transmission	Rt 1.52 (95% CrI 1.50–1.54); P(Rt > 1) = 100%	1.00	0.25	25.0
Growth momentum	+9.4% a day; doubling in about 7 days	0.96	0.15	14.4
Burden vs. history	100th percentile of the region's history	1.00	0.20	20.0
Aberration	All detectors in alarm	1.00	0.20	20.0
Forecast surge risk	P(next 14 days > 1.2 × last 14) = 100%	1.00	0.20	20.0
Total				**≈ 99 → High**

> These figures come from the application's synthetic demonstration data, not from a real jurisdiction.

6.3 Outputs

- **Association network:** variables coloured by category; line width shows strength and line colour shows direction; strongly linked variables are placed together.
- **Variable focus view:** every association of a chosen variable, ranked, with both measures side by side.
- **Spearman-vs-distance-correlation scatter:** points well above the diagonal reveal non-monotonic links.
- **Sortable table of all pairs:** exportable as CSV.

6.4 Subgroup analysis

Any suitable column can define subgroups:

- **Text columns** with 2–20 values, such as population, sex or district.
- **Categorical variables**, such as marital status.
- **Numeric variables split into thirds**, such as age bands.

Analysts can then compare associations across groups and check for Simpson's paradox, where pooling groups creates or hides an association. Category labels come from the data (Category|Name column headers), so the module is not tied to any one survey.

7. Data requirements and inputs

Item	Surveillance pipeline (P–R–I–S–M)	Risk-factor module
Unit	Case counts per region per day or week	One row per respondent
Format	CSV, long (date, region, cases, population) or wide (one column per region)	CSV, one numeric column per question; optional text columns for subgroups
Minimum history	60 days per region (120+ recommended); 1+ year turns on the seasonal detector	30 respondents (100+ recommended)
Weekly data	Accepted; spread evenly across days (weekday effects cannot then be estimated)	n/a
Population	Optional; enables rates per 100,000	n/a
Disease parameters	Serial-interval mean and SD (presets or custom)	n/a
Missing values	Interpolated	Pairwise deletion

> The surveillance pipeline needs only aggregate counts, not individual records. Individual survey data for the risk-factor module should be de-identified and used under appropriate ethical approval (Section 11).

8. Validation and evidence

The methods have been verified for correctness and evaluated on realistic synthetic data. Evaluation on real historical outbreaks has not yet been done; Section 8.3 sets out how it will be.

8.1 Correctness verification (completed)

Check	Result
Automated test suite	**29 of 29 tests pass**
Distance correlation and Spearman's ρ vs. Python dcor 0.7 and scipy 1.18	Identical to 10 decimal places
Rt estimator on a simulated renewal process with known Rt = 1.25	Recovered within 1.15–1.35; P(Rt > 1) > 99%
Weekday factors on a known reporting pattern	Recovered to within 0.05
Quasi-Poisson trend on known 5% daily growth	Growth rate recovered to 2 decimal places
Distance-correlation independence test under the null	False-positive rate below 8% at nominal 5% (conservative)
Data ingestion	Long, wide and weekly CSVs; gaps; text and yes/no columns

8.2 Synthetic surveillance evaluation (completed)

Setup. PRISM-EW™ includes a simulator for testing. It generates 8 regions over 2 years, with:

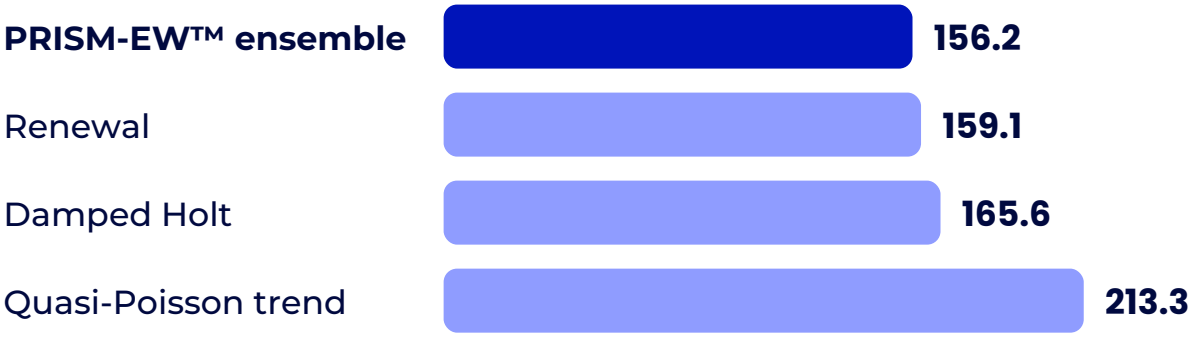
- a stochastic renewal process with immunity that builds and wanes, and winter seasonality;
- about 22% of infections reported, with weekday reporting effects;
- 0.4% of days missing at random;

- two planted outbreaks, in which transmission was multiplied by 1.5 in Riverside and 1.38 in Coastal Bay.

Forecasts were backtested from 6 weekly origins, scored at 1–14 days ahead, and averaged over the 8 regions.

Model	Mean WIS (lower is better)	50% interval coverage	95% interval coverage
PRISM-EW™ ensemble	**156.2**	45%	86%
Renewal	159.1	36%	80%
Damped Holt	165.6	47%	87%
Quasi-Poisson trend	213.3	32%	73%

Mean WIS (lower is better)



Findings:

- The ensemble had the lowest average error of all four forecasts.
- In 5 of 8 regions it was within 5% of the best single model for that region, and it was never the worst model in any region.
- The 50% intervals were close to their target. The 95% intervals were too narrow: 86% coverage on average, falling to 68–77% in the two regions with fast-accelerating outbreaks. Widening the intervals during rapid growth is a priority improvement.

Detection of the planted outbreaks:

- **Riverside:** its OPI™ had already begun rising with a seasonal upswing (34 → 56 in the week before). A detector alarm followed 1 day after the transmission shift began, and the region reached High 2 days after.
- **Coastal Bay:** the region was already Elevated or High from seasonal activity (OPI™ 62–84). The first detector alarm came 5 days after the shift.

Both regions finished at OPI™ 99 (High).

> Because both shifts happened during seasonal rises, these timings show responsiveness, not a clean measure of lead time. That is why real-data evaluation is the next step.

8.3 Real-data evaluation protocol (planned)

- **Data:** public historical surveillance series with documented outbreaks. Examples include CDC FluView/ILINet (US influenza-like illness by region), WHO FluNet, and archived COVID-19 and measles case series.

Design:

- Replay each season day by day, using only data available at each date.
- Run all parameters frozen in advance; no tuning on the test seasons.
- Compare against simple baselines: a naive forecast ("the next 14 days equal the last 14") and EARS C2 used on its own.

Metrics:

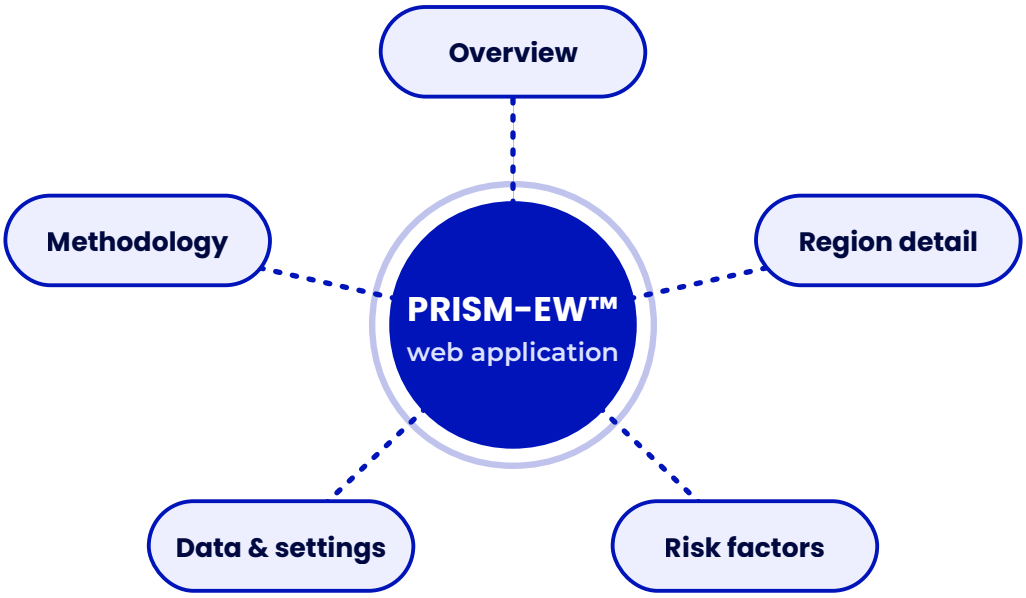
- **Lead time:** days between the OPI™ first reaching Elevated (and High) and the epidemic peak or official outbreak declaration.
- **Alert quality:** sensitivity, false-alarm rate per region-year, and positive predictive value at each alert level.
- **Forecast skill:** WIS relative to the naive baseline, and 50%/95% interval coverage.

Results will be reported in this format:

Dataset and seasons	Median lead time to Elevated (days)	Sensitivity	False alarms per region-year	Relative WIS vs. naive	95% coverage
To be completed					

9. The PRISM-EW™ web application

9.1 Modules



Module	What it shows
Overview	Headline OPI™ and alert level; Rt, cases per day, growth rate and surge-probability tiles; a sortable regional ranking with 60-day trend lines; the combined forecast

Module	What it shows
Region detail	Forecast with 50%/95% intervals (log scale and per-model views); Rt with its credible band; detector alarms; the OPI™ point breakdown (“Why this score”); nowcast OPI™ trajectory; forecast-accuracy table
Risk factors	Association network, variable focus view, Spearman-vs-distance-correlation scatter, table of all pairs, and subgroup filters (Section 6)
Data & settings	CSV upload with validation messages and template; synthetic data generator; disease presets; forecast horizon, Rt window, simulation count and OPI™ weights
Methodology	Every formula, threshold, limitation and reference, updated live with the current settings

Exports: forecast CSV for each region, full analysis as JSON, and all risk-factor pairs as CSV.

9.2 Architecture, privacy and reliability

- **Local processing:** a static React and TypeScript application hosted on Firebase. All statistics run in the user’s browser in a Web Worker, so the page stays responsive and uploaded data never reach a server.
- **No backend:** no database, no login and no storage of case or survey data.
- **Reproducible:** every simulation uses a fixed random seed, so the same data and settings always give the same results.
- **Fast:** the full surveillance pipeline for 8 regions plus the total runs in about 1 second; the 666-pair risk-factor analysis of 1,200 respondents runs in about 3 seconds.
- **Portable engine:** the statistical engine has no user-interface dependencies, so it can also run in Node.js, Cloud Functions or a scheduled batch job.
- **Accessible:** light and dark themes, phone-width layouts, and alert levels shown with icon and text as well as colour.
- **Quality assurance:** an automated test suite, plus a CI/CD pipeline that runs the tests before every deployment.

10. Positioning and differentiators

Established tools each cover part of the early-warning task well. PRISM-EW™ is designed to sit alongside them and integrate their ideas, not to replace them.

Tool or approach	Main strength	What PRISM-EW™ adds
EpiEstim / EpiNow2 (R)	Rt estimation (EpiNow2 also corrects for reporting delays)	Rt read alongside burden, forecasts and detection in one score
CDC FluSight and forecasting hubs	Multi-team ensemble forecasts, evaluated by WIS	A local, self-scored ensemble for any jurisdiction's own data
R surveillance package, ESSENCE	Aberration detection (Farrington, EARS, CUSUM)	Detection fused with transmission and forecast risk, with explanation
Dashboards of raw counts	Easy access and familiarity	Uncertainty, prioritisation and early warning

Differentiators

- **Integrated early warning:** preparation, transmission, forecasting, detection and scoring in one pipeline.
- **Probabilistic throughout:** credible intervals, prediction intervals and probabilities rather than single point estimates.
- **Self-weighting ensemble:** model weights come from each region's own backtest performance.
- **Explainable alerts:** every OPI™ point is attributed to a named signal.
- **Privacy by architecture:** analysis runs in the browser; no data leave the device.
- **Disease-adaptable:** serial-interval presets and configurable weights.
- **Two lenses:** population-level outbreak pressure, plus individual-level risk-factor structure.

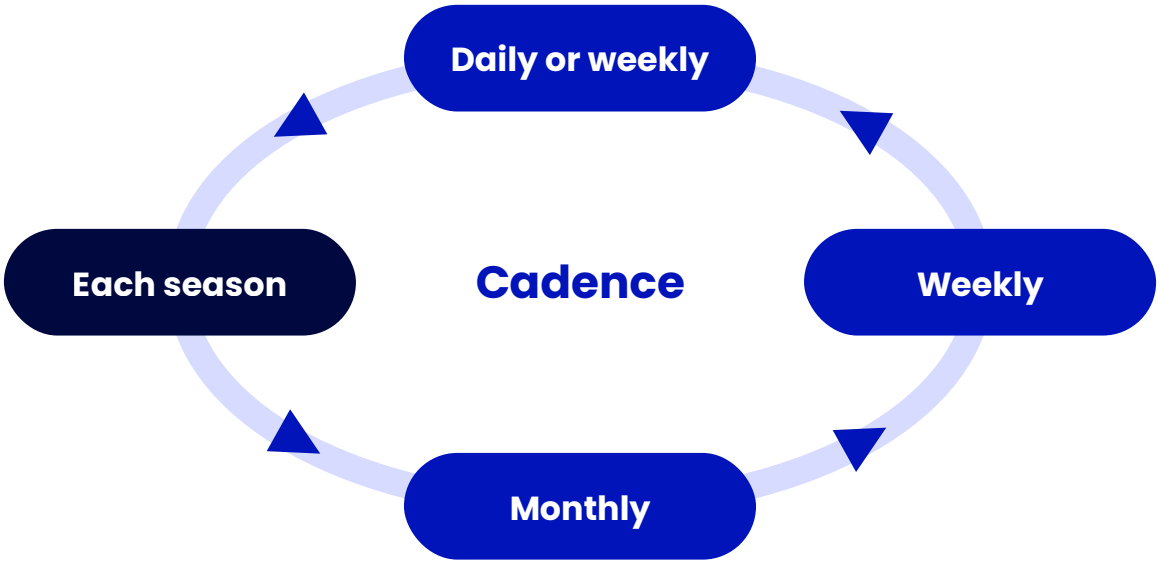
> PRISM-EW™ does not claim to have invented Rt estimation, EARS, CUSUM, distance correlation or the individual forecasting models. Its contribution is the integrated, explainable architecture.

11. Operating model and governance

11.1 Roles

Role	Responsibility
Public-health lead	Surveillance priorities, escalation decisions, response coordination
PRISM-EW™ lead	Data preparation, configuration, scoring and reporting
Epidemiologists	Interpreting transmission and detection signals; deciding what to investigate
Data and statistics team	Data quality, forecast validation, model monitoring and recalibration
Health authorities	Investigation, preparedness and response

11.2 Cadence



- **Daily or weekly:** update case data and review regional alert levels.
- **Weekly:** review Rt, forecasts, OPI™ rankings and detector alarms.
- **Monthly:** review forecast accuracy and interval coverage in the accuracy table.

- **Each season:** recalibrate weights, cut-points and disease parameters against observed outcomes.

11.3 Governance principles

- **Human in the loop:** an alert level triggers human review and investigation, never automatic action.
- **Documented configuration:** every change to weights, cut-points or disease parameters is recorded with a date and reason, so past alerts can be reproduced.
- **Data protection:** the surveillance pipeline uses aggregate counts. Individual survey data for the risk-factor module should be de-identified and analysed under institutional review board (IRB) approval or equivalent.
- **Communicating uncertainty:** reports show intervals and probabilities, not only point values.

12. Application blueprints and illustrative scenarios

12.1 Blueprints

Use	How PRISM-EW™ helps	Preset serial interval
Influenza surveillance	Seasonal transmission, unusual increases, near-term surge risk	3.6 / 1.6 days
COVID-19 surveillance	Transmission, burden, forecasting and detection combined	4.7 / 2.9 days
RSV monitoring	Regional shifts in paediatric respiratory pressure	7.5 / 2.1 days
Measles surveillance	Fast detection of unusual incidence in low-burden areas (escalation guard)	11.7 / 2.0 days
Cholera surveillance	Geographic patterns and areas needing investigation	5.0 / 3.0 days
HIV prevention research	Risk-factor networks to target testing and PrEP screening (Section 6)	n/a

12.2 Illustrative scenarios

- **Regional respiratory outbreak:** Rt rises, cases accelerate and surge risk is high, so the region moves to High.
- **Early transmission shift:** burden is still historically low, but a detector alarms and P(Rt > 1) exceeds 90%, so the escalation guard lifts the region to Elevated for investigation.
- **Seasonal activity:** high counts are compared with the same weeks in earlier years, separating expected seasonal activity from abnormal increases.

- **Multi-region prioritisation:** the OPI™ ranking directs limited investigation staff to the regions under most pressure.

> These scenarios illustrate how the methods behave. They are not completed independent deployments.

13. Limitations

PRISM-EW™ supports decisions; it does not replace epidemiological judgement or field investigation.

- **Reporting delays are not yet modelled.** In live feeds the most recent days are often incomplete, which can make a rise look like it is slowing. Nowcasting is on the roadmap.
- **Rt assumptions:** Rt assumes a known serial interval and stable testing and ascertainment. A large change in testing biases Rt.
- **Interval width:** in synthetic tests, 95% forecast intervals covered 86% of outcomes rather than 95%, and less during fast acceleration.
- **Backtest look-ahead:** backtests reuse weekday factors estimated from the full series, so reported accuracy is slightly optimistic.
- **Local calibration needed:** the weights and cut-points are policy choices and should be calibrated against past outbreaks in each jurisdiction.
- **Synthetic evidence only so far:** real-data performance is not yet established (Section 8.3).
- **Risk-factor associations are cross-sectional and unadjusted.** They show co-occurrence, not causation, and pooled groups can create or hide associations.

14. Roadmap and current status

1

Phase 1 – Methodology and architecture (complete).

Five-stage P-R-I-S-M design, OPI™ specification, and correctness verification.

2

Phase 2 – Web application (complete as a working prototype, October 2026).

Dashboards, regional analysis, forecasting, detection, risk-factor module, and synthetic evaluation.

3

Phase 3 – Real-data validation (next).

Retrospective evaluation on public historical outbreaks using the Section 8.3 protocol. Interval calibration during rapid growth. A pilot with a partner health department or research group.

4

Phase 4 – Operational extensions.

- Reporting-delay correction (nowcasting).
- Automated data feeds.
- Configurable alert notifications, for example when a region changes level.
- Optional team accounts with stored configurations.
- Further disease models.

15. Creator



Oluwatobi Lucas Akinbode is a statistical data scientist and AI researcher, and a Ph.D. student in Applied Science & Technology (Data Science & Analytics) at North Carolina A&T State University. His work spans predictive healthcare analytics, statistical methodology and scientific computing. It includes:

- HIV risk-behaviour network analysis at Duke University;
- environmental-health modelling with the NIH *All of Us* Research Program;
- interval-valued data analysis and uncertainty quantification;
- cloud-based statistical computing platforms for more than 500 students.

He designed and built PRISM-EW™.

Contact: olakinbode@aggies.ncat.edu · codelifeonline.org · github.com/tobzz30

Collaborators sought: epidemiologists, public-health researchers, statisticians, health departments and surveillance professionals interested in piloting or validating PRISM-EW™.

CALL TO ACTION

Preprocess. Reproduce. Forecast. Signal. Merge.

Partner with PRISM-EW™ on a retrospective early-warning assessment of your surveillance data. It will measure how early the OPI™ would have flagged past outbreaks, and set out a measurable path toward predictive, responsive public-health surveillance.

PRISM-EW™: turning disease surveillance data into predictive early-warning intelligence.

References

1. Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society, Series B*. 1995;57(1):289–300.
2. Bracher J, Ray EL, Gneiting T, Reich NG. Evaluating epidemic forecasts in an interval format. *PLoS Computational Biology*. 2021;17(2):e1008618.
3. Cori A, Ferguson NM, Fraser C, Cauchemez S. A new framework and software to estimate time-varying reproduction numbers during epidemics. *American Journal of Epidemiology*. 2013;178(9):1505–1512.
4. Gardner ES Jr, McKenzie E. Forecasting trends in time series. *Management Science*. 1985;31(10):1237–1246.
5. Hutwagner L, Thompson W, Seeman GM, Treadwell T. The bioterrorism preparedness and response Early Aberration Reporting System (EARS). *Journal of Urban Health*. 2003;80(Suppl 1):i89–i96.
6. Noufaily A, Enki DG, Farrington P, Garthwaite P, Andrews N, Charlett A. An improved algorithm for outbreak detection in multiple surveillance systems. *Statistics in Medicine*. 2013;32(7):1206–1222.
7. Székely GJ, Rizzo ML, Bakirov NK. Measuring and testing dependence by correlation of distances. *Annals of Statistics*. 2007;35(6):2769–2794.
8. Akinbode OL, Howard L, Ostermann J, Thielman NM. Can non-HIV risk behaviors serve as a proxy measure of HIV risk? Poster, 20th Annual Duke CFAR Fall Scientific Retreat, Durham, NC; September 2024.



Oluwatobi Akinbode

PRISM-EW™

olakinbode@aggies.ncat.edu · codelifeonline.org · github.com/tobzz30

6. Risk-factor association module

A separate module analyses **individual-level survey data** (one row per respondent) rather than aggregate case counts. It asks which behaviours and characteristics occur together, for example whether non-HIV risk behaviours such as alcohol use or general risk-taking can stand in for HIV risk. This builds on the author's network analysis of HIV risk in Kilimanjaro, Tanzania, presented at the 2024 Duke CFAR Fall Scientific Retreat.

6.1 Two complementary measures for every pair of variables

ρ

Spearman's ρ (−1 to 1): rank-based, so it measures *monotonic* association: whether two answers rise or fall together in one direction. It suits Likert items, counts and yes/no answers. p-values come from the t-approximation.

dCor

Distance correlation (0 to 1; Székely, Rizzo & Bakirov, 2007): zero only when two variables are independent. It therefore also detects non-monotonic links, such as partner numbers peaking at a middle age. p-values come from a conservative asymptotic test.

6.2 Statistical safeguards

- **Multiple-testing control:** with hundreds of pairs some will look significant by chance, so q-values are adjusted with the Benjamini–Hochberg false discovery rate, separately for each measure.
- **Non-monotonic flag:** a pair is flagged when distance correlation is significant ($q < 0.05$), at least 0.15, and at least 0.10 above $|\rho|$. These are links that Spearman's ρ understates or misses.
- **Missing answers:** each pair uses every respondent who answered both questions.