predictset

Conformal prediction and uncertainty quantification in R

Charles Coverdale

London, United Kingdom

April 2026

The question

How do we attach finite-sample, distribution-free coverage guarantees to any fitted model in R?

Eleven methods, one model-agnostic interface, any learner from `lm` to `xgboost`.

Why it matters

- **Regulators** (EU AI Act, FDA SaMD) now require calibrated uncertainty alongside point predictions
- **Clinical prediction** needs per-patient intervals that are valid without model assumptions
- **Credit and insurance** need group-conditional coverage that a marginal guarantee cannot deliver
- **Operations research** needs sequential intervals that adapt when the data-generating process drifts

What is already out there

- **`probably`** (tidymodels): regression only, no classification, no Jackknife+ or CV+[1]
- **`conformalInference`** (GitHub only): research code, not on CRAN, unmaintained since 2019
- **Classification, Mondrian, weighted, ACI**: no prior CRAN implementation at all
- **Bespoke scripts**: most practitioners roll their own, with no shared test harness

The gap: **no single CRAN package covers the standard conformal workflow end to end**.

[1] Kuhn & Vaughan (2023), *probably: Tools for Post-Processing Class Probability Estimates*, R package.

What predictset offers

1. **Coverage**: 11 methods spanning split, CV+, Jackknife+, CQR, APS, RAPS, LAC, Mondrian, weighted, ACI, conformal p-values
2. **Interface**: every method takes a formula, a fitted model, or a user-defined train/predict pair; same three-line call across regression and classification
3. **Provenance**: pure R, no compiled code, four runtime imports, 318 tests, CRAN since March 2026

Methods follow @angelopoulos2023gentle and the primary sources cited per function[2].

[2] Angelopoulos & Bates (2023), *Conformal Prediction: A Gentle Introduction*, Foundations and Trends in Machine Learning 16(4).

Regression: split conformal prediction intervals

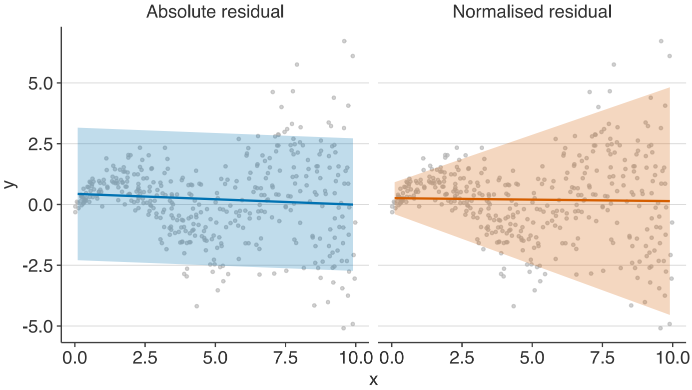

Split conformal gives constant-width bands; normalised scores scale with local variance. Exports: `conformal_split()`, `conformal_cv()`, `conformal_jackknife()`, `conformal_cqr()`.

Classification: adaptive prediction sets (APS)

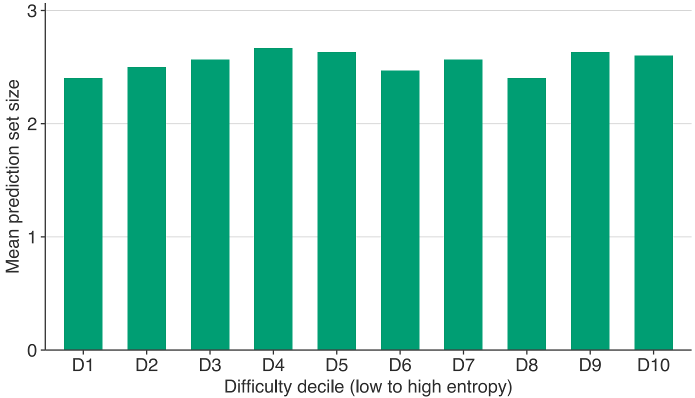

APS set size grows with per-instance uncertainty while holding marginal coverage at the target. Exports: `conformal_aps()`, `conformal_raps()`, `conformal_lac()`.

Mondrian: group-conditional coverage

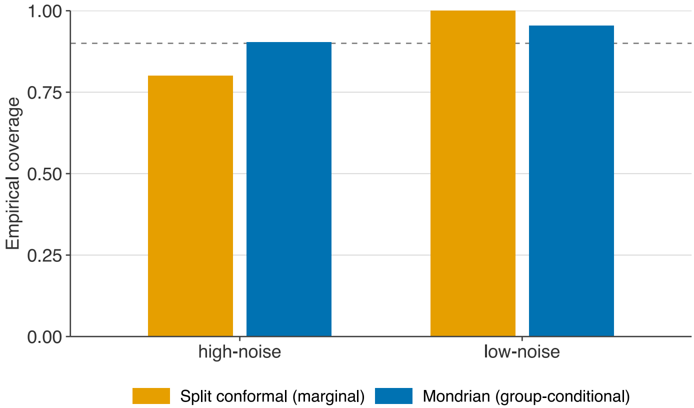

Mondrian calibrates a separate quantile per group, recovering the target within each subgroup. Exports: `conformal_mondrian()`, `conformal_mondrian_class()`.

Weighted conformal and adaptive (ACI)

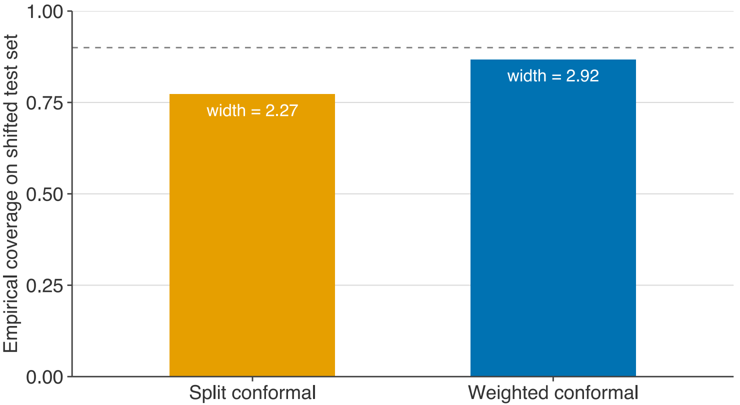

Weighted reweights scores by a density ratio for covariate shift; ACI updates target level online under drift. Exports: `conformal_weighted()`, `conformal_aci()`.

Central formulas

Nonconformity score (absolute residual form):

$$s_i = |y_i - \hat{f}(x_i)|, \qquad i \in \mathcal{I}_{\text{cal}} \tag{1}$$

Conformal quantile (calibration order statistic):

$$\hat{q}_\alpha = s_{(\lceil (n_{\text{cal}}+1)(1-\alpha) \rceil)} \tag{2}$$

Coverage guarantee (Vovk et al. 2005; Lei et al. 2018):

$$\mathbb{P}\left\{Y_{n+1} \in \hat{C}(X_{n+1})\right\} \;\geq\; 1 - \alpha \tag{3}$$

Package at a glance

Function families:

- **Regression**: `conformal_split`, `conformal_cv`, `conformal_jackknife`, `conformal_cqr`
- **Classification**: `conformal_aps`, `conformal_raps`, `conformal_lac`
- **Conditional / adaptive**: `conformal_mondrian`, `conformal_mondrian_class`, `conformal_weighted`, `conformal_aci`
- **Diagnostics**: `coverage`, `coverage_by_group`, `coverage_by_bin`, `interval_width`, `set_size`, `conformal_compare`

Deps: `cli`, `grDevices`, `graphics`, `stats`. R $\geq$ 4.1.0.

Three-tier model interface

Every method accepts:
1. A formula `y ~ .` 2. A fitted `lm`, `glm`, `ranger` 3. A `make_model()` wrapper for `xgboost`, custom learners

Minimal working example

```r
library(predictset)

# Fit, calibrate, predict: any model, any method, three lines
fit  <- conformal_split(y ~ ., data = train,
                        x_new = x_test, alpha = 0.10)
cov  <- coverage(fit, y_test)
new  <- predict(fit, newdata = future_data)

# Swap method: same call, change one word
fit2 <- conformal_jackknife(y ~ ., data = train,
                            x_new = x_test, alpha = 0.10)
```

The conformal object caches the fitted model and calibration quantile, so new data is scored without refitting.

Case study: fairness via group-conditional coverage

Data. ProPublica COMPAS recidivism dataset, 6167 defendants in Broward County, Florida, screened 2013-14[3].

Question. *Can conformal prediction deliver equalised per-group coverage where a marginal guarantee cannot?*

Why this case. Canonical benchmark in algorithmic-fairness literature. A 90 per cent marginal guarantee is consistent with 99 per cent coverage on one subgroup and 60 per cent on another. @romano2020malice treat equalised coverage as a minimum fairness standard.

[3] Angwin et al. (2016), *Machine Bias*, ProPublica.

COMPAS: Mondrian APS equalises per-race coverage

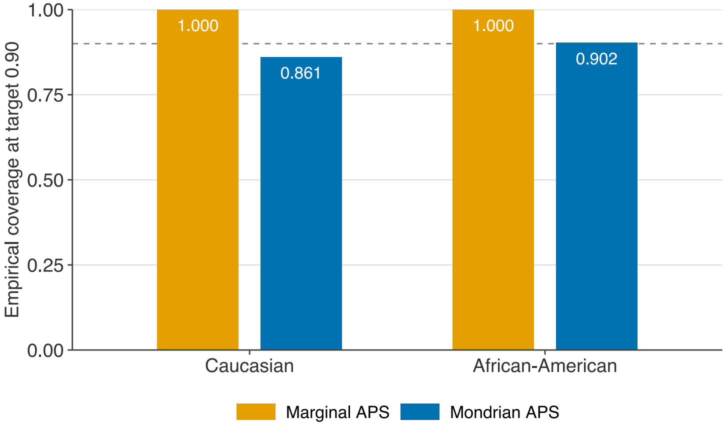

Per-race coverage on COMPAS two-year recidivism at target $1 - \alpha = 0.90$. *Marginal APS over-covers both subgroups; Mondrian APS tracks the target per group.*

CV+ and Jackknife+ recover width without losing coverage

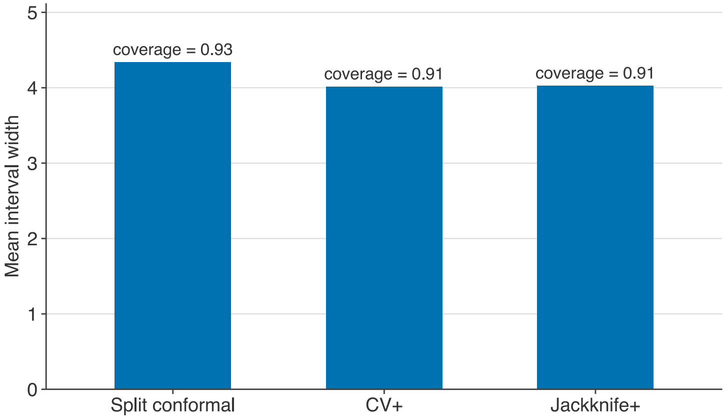

Mean interval width at target coverage 0.90. CV+ and Jackknife+ reclaim roughly 9 per cent of the width that split conformal sacrifices by halving the training data.

What predictset does not yet do

- **No parametric or Bayesian intervals**: use `rstanarm`, `brms`, or base `predict()` for those
- **Time-series dependence**: only ACI is implemented; full sequential-data conformal methods are pending
- **Weighted conformal requires user-supplied weights**: no bundled density-ratio estimator
- **Classification set size depends on base calibration**: miscalibrated probabilities yield larger sets

v0.4.0 roadmap: conformal procedures for dependent data beyond ACI, likelihood-ratio estimators for weighted conformal, tighter `tidymodels` integration.

Contact, code, paper

Charles Coverdale

`charles.f.coverdale@gmail.com`
`github.com/charlescoverdale/predictset`
`charlescoverdale.github.io/publications`

Acknowledgements. COMPAS data: ProPublica (public release). Thanks to the conformal-prediction community, in particular the work of Vovk, Lei, Romano, Barber, Candes, Angelopoulos, and Tibshirani, CRAN maintainers, and reviewers of *predictset* v0.3.0.

Scan for paper PDF