Assumption-Free Uncertainty Quantification in Epigenetic Clocks

MEGAN MITCHELL[1], Lajoyce Mboning[2], Louis-S. Bouchard[2], Matteo Pellegrini[3]

[1] BIG Summer Program, Institute for Quantitative & Computational Biosciences, UCLA [2] Department of Chemistry and Biochemistry, UCLA

[3] Department of Molecular Cell and Developmental Biology, UCLA

Abstract

Epigenetic clocks – models that estimate biological age from DNA methylation (DNAm) data – typically provide point predictions without accompanying measures of uncertainty. To address this limitation, our previously published model, BayesAge, used simulated data to generate a range of plausible ages for each individual. However, this approach was computationally intensive and systematically underestimated error. In this study, we evaluated split conformal prediction and its locally weighted variant as computationally efficient and statistically rigorous methods for quantifying the uncertainty of BayesAge's predictions. We found that split conformal inference produced empirically valid confidence intervals, even in finite-sample settings, with minimal residual bias. These findings suggest that conformal inference may be valuable for supporting high-risk decision-making tasks in clinical or forensic contexts, where reliable uncertainty bounds are essential.

Background

- **Epigenetic Clocks:** DNAm is an emerging biomarker of aging that has powered an increasing number of epigenetic clocks, which estimate an individual's biological age based on age-associated CpG sites. The difference between between biological and chronological age, known as age acceleration, has been linked to diseases including cancer, cardiovascular disease, and neurological disorders. We previously introduced BayesAge, an epigenetic clock that leveraged targeted bisulfite sequencing (TBS-seq) data to produce maximum likelihood age estimates and used LOWESS regression to account for nonlinear methylation patterns.
- **Conformal Inference**: Conformal prediction is a model-agnostic, distribution-free algorithm for uncertainty quantification that offers rigorous coverage guarantees. Recently, conformalized quantile regression (CQR), which applies this method, has been shown to produce valid confidence intervals in the context of epigenetic age prediction. In this work we explored split conformal prediction, which divides data into training and calibration sets to estimate non-conformity scores and construct confidence intervals. We also tested locally weighted conformal methods to account for heteroskedastic data.

Methods

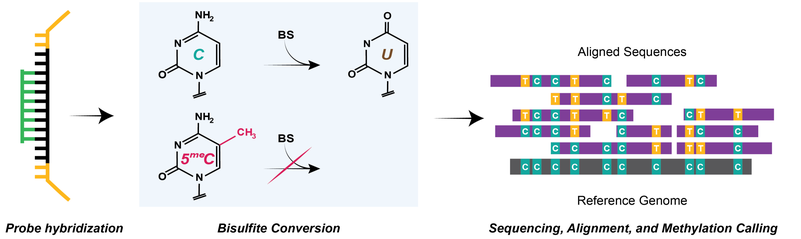

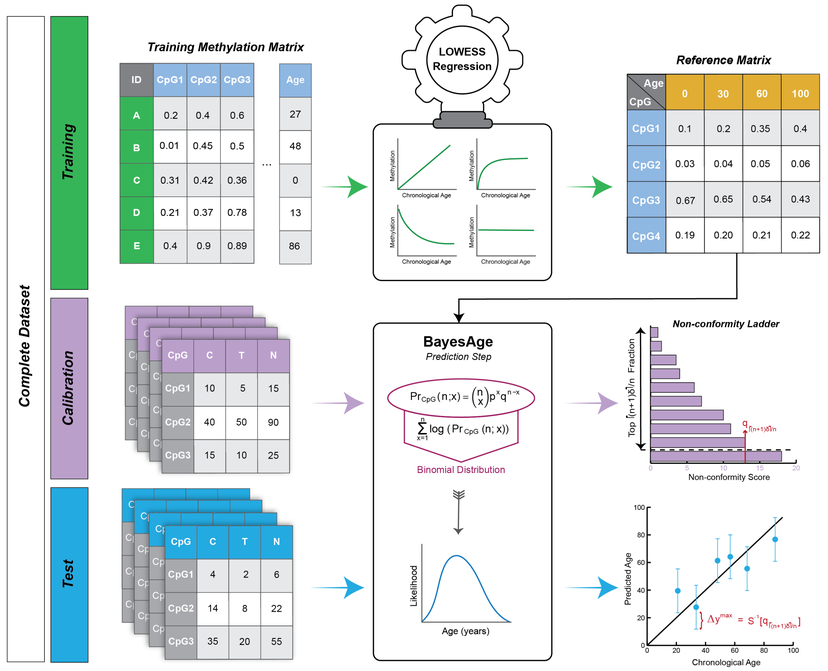

Figure 1. TBS-seq overview and BayesAge pipeline.

Results

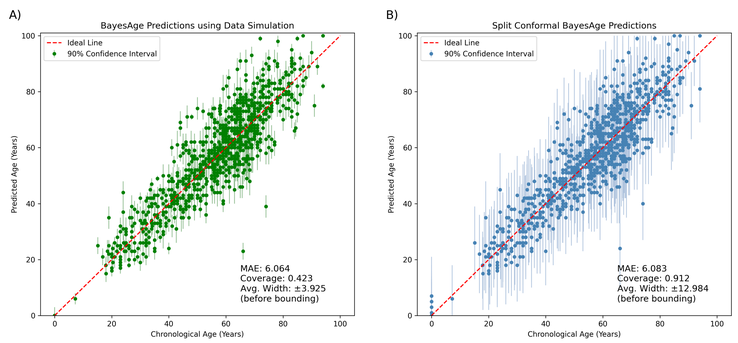

Figure 2. A) Ages predicted using the BayesAge framework with the top 16 most highly correlated CpG sites. Upper and lower bounds were computed by using the 5th and 95th percentiles of ages predicted from a list of 100 simulated samples per individual. **B)** Ages predicted using the BayesAge with split conformal inference framework using the top 16 CpG sites and 200 calibration points.

	Split Conformal		Locally Weighted Conformal	
δ	Coverage	Avg. Width (Years)	Coverage	Avg. Width (Years)
0.99	0.992	±23.74	0.992	±22.97
0.95	0.957	±16.19	0.941	±15.53
0.9	0.912	±12.98	0.896	±12.70
0.8	0.809	±9.60	0.801	±9.64
0.7	0.715	±7.67	0.694	±7.78
0.6	0.614	±6.11	0.583	±6.27
0.5	0.509	±4.82	0.485	±4.84

Table 1. Results at different confidence levels for both split and locally weighted conformal prediction.

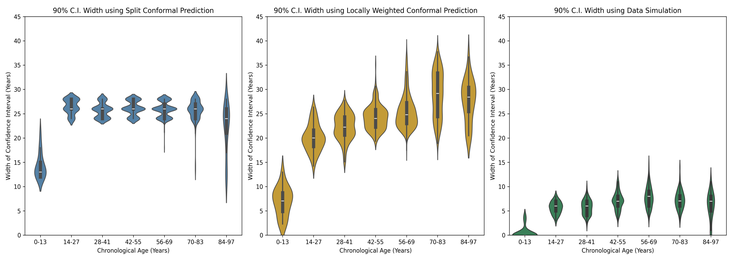

Figure 3. Interval widths by age bin using split conformal inference, locally weighted conformal inference, and data simulation.

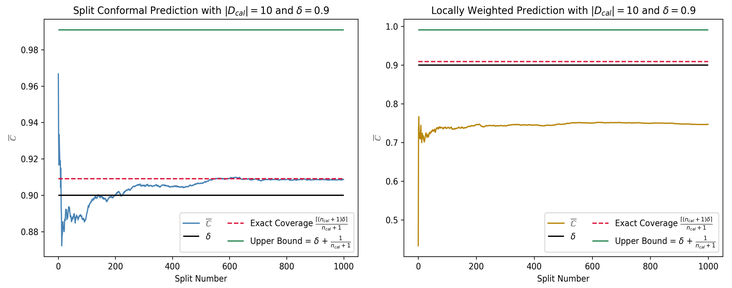

Figure 4. Running empirical coverage of split and locally weighted predictions using 10 training points, 10 calibration points, and 10 prediction points over 1000 splits.

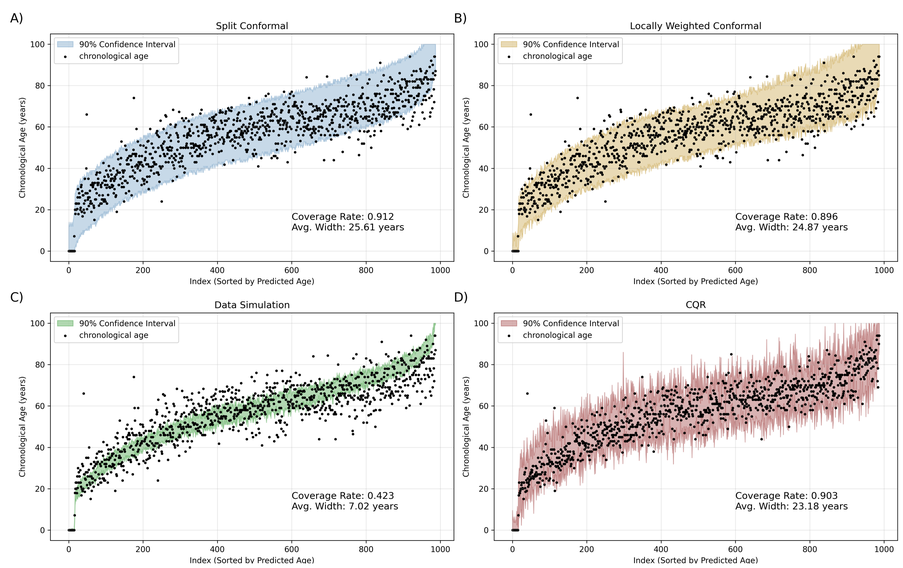

Figure 5. Coverage plots for multiple methods of uncertainty quantification. For each plot, the x-axis represents the index of each sample after sorting by predicted age. The coverage was computed by finding the fraction of samples for which the true age fell between the predicted upper and lower bounds. Ages were predicted using cross-validation with 88 samples at a time.

A) Split conformal inference: Ages were predicted using BayesAge with 200 calibration points and mean absolute error (MAE) as the non-conformity function.

B) Locally weighted conformal inference: Ages were predicted using BayesAge with 200 calibration points. The non-conformity scores were normalized by a local estimate of mean absolute deviation (MAD).

C) Data simulation: For each sample, 100 simulated samples were drawn from binomial distributions at each of the top 16 CpG sites. BayesAge was used to predict the age of each of the simulated samples, and the confidence interval was constructed using the 5th and 95th percentile of the predicted age list for each sample.

D) CQR: Implemented using the MapieQuantileRegressor package using the top 88 most highly correlated CpG sites and 200 calibration samples.

E) Lasso: Split conformal predictions using Lasso regression with an α parameter of 0.02. The number of CpG sites that were used averaged 244.5 per fold.

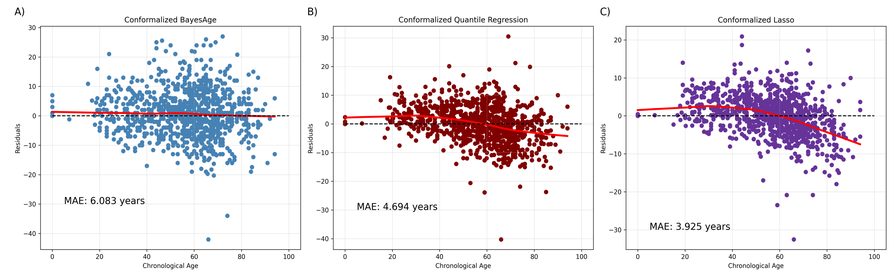

Figure 6. Residuals versus chronological age for **A)** BayesAge with split conformal predictions, **B)** CQR, and **C)** Lasso regression with split conformal prediction. LOWESS smoothing was applied between the residuals and chronological age with a default τ value of 0.67.

Conclusions

- Split conformal inference provides rigorous coverage guarantees, even in finite-sample settings. It is more computationally efficient than data simulation and can be applied to any model, regardless of the distribution of the underlying data. Locally weighted conformal inference produces confidence intervals that adapt to local variance but may not maintain coverage guarantees when sample sizes are very small.
- High-dimensional penalized linear regression methods yield lower MAE scores than BayesAge, but exhibit age-associated residuals that could affect interpretability.
- Future work might consider implementing CQR with nonparametric quantile regression techniques, such as gradient-boosted quantile regression or quantile regression forests.

Acknowledgements

I would like to thank the Bruins in Genomics summer program for this valuable research opportunity. I am so grateful to Lajoyce Mboning for her support and continued mentorship. I would also like to thank Dr. Pellegrini and Dr. Bouchard for their input and guidance on this project.

References

1. Mboning L, Rubbi L, Thompson M, Bouchard LS, Pellegrini M. BayesAge: A maximum likelihood algorithm to predict epigenetic age. Frontiers in Bioinformatics. 2024 Apr 4;4:1329144.
2. Li Y, Goodrich JM, Peterson KE, Song PX, Luo L. Uncertainty quantification in epigenetic clocks via conformalized quantile regression. Genetic Epidemiology. 2025 Jun;49(4):e70008.
3. Lei J, G'Sell M, Rinaldo A, Tibshirani RJ, Wasserman L. Distribution-free predictive inference for regression. Journal of the American Statistical Association. 2018 Jul 3;113(523):1094-111.
4. Morselli M, Farrell C, Rubbi L, Fehling HL, Henkhaus R, Pellegrini M. Targeted bisulfite sequencing for biomarker discovery. Methods. 2021 Mar 1;187:13-27.
5. MLBoost. Uncertainty Quantification (1): Enter Conformal Predictors. June 2, 2023.
6. Hannum G, Guinney J, Zhao L, Zhang LI, Hughes G, Sadda S, Klotzle B, Bibikova M, Fan JB, Gao Y, Deconde R. Genome-wide methylation profiles reveal quantitative views of human aging rates. Molecular cell. 2013 Jan 24;49(2):359-67.
7. Trapp A, Kerepesi C, Gladyshev VN. Profiling epigenetic age in single cells. Nature Aging. 2021 Dec;1(12):1189-201.