

Enhancing the performance and interpretability of epigenetic clocks

Tushar Patel, Robert Schwarz, Konstantin Riege, Miri Varshavsky, Anne Richmond, Riccardo E Marioni, Hans A Kestler, Tommy Kaplan, Steve Hoffmann, Alena van Bömmel

Enhancing the performance and interpretability of epigenetic clocks

Tushar Patel¹, Robert Schwarz¹, Konstantin Riege¹, Miri Varshavsky², Anne Richmond³, Riccardo E. Marioni³, Hans A. Kestler^{1,4}, Tommy Kaplan^{1,5}, Steve Hoffmann^{1,6,*}, Alena van Bömmel^{1,7,8,*}

¹Leibniz Institute on Aging - Fritz Lipmann Institute (FLI), Beutenbergstraße 11, Jena 07745, Germany

²School of Computer Science and Engineering, The Hebrew University of Jerusalem, Jerusalem 9190401, Israel

³Institute of Genetics and Cancer, University of Edinburgh, Edinburgh, EH4 2XU, UK

⁴Institute of Medical Systems Biology, Ulm University, Albert-Einstein-Allee 11, Ulm, 89081, Germany

⁵Barts Cancer Institute, Queen Mary University of London, London, EC1M 6BQ, UK

⁶Department of Mathematics and Computer Science, Freie Universität Berlin, Berlin, 14195, Germany

⁷Chair of Metabolic Biochemistry, Biomedical Center (BMC), LMU Medizin, Ludwig-Maximilians-Universität München, Munich, 81377, Germany

⁸Neuroimmunology and Neurodegenerative Diseases, German Center for Neurodegenerative Diseases (DZNE), Munich, 81377, Germany


*To whom correspondence should be addressed. Email: alena.vanboemmel@leibniz-flf.de

Correspondence may also be addressed to Steve Hoffmann. Email: steve.hoffmann@leibniz-flf.de

Abstract

Epigenetic clocks based on DNA methylation (DNAm) accurately predict age, but their biological underpinnings remain unclear. One primary mechanism by which DNAm might influence gene regulation is by modulating transcription factor binding activity. This study investigates the regulatory potential of predictive CpGs in established epigenetic clocks. Our analysis reveals that generally most CpGs used by epigenetic clocks do not overlap known transcription factor binding sites (TFBS), indicating that clock accuracy is not primarily driven by changes in TF binding dynamics. However, analysis of CpGs within TFBSs identifies key transcription factors potentially involved in aging, including *ZBED1*, *NFE2*, and *CEBPB*, which are enriched for age-associated CpGs, while *RELA*, *IKZF1*, and *STAT3* significantly protected against methylation changes. Leveraging TFBS-associated and age-correlated CpGs, combined with noise-stabilizing feature engineering steps, we developed an alternative TFMethyl Clock model that provides competitive predictions of chronological age. Age-predictive CpGs selected by our model enrich for target genes involved in interleukin-1 β production and fatty-acid metabolism, while being enriched at TFBSs of *NR2C2*. Furthermore, approximately three-fourths of these target genes exhibit significant age-related changes, suggesting deeper insights into possible methylation-driven aging processes. Our findings demonstrate that incorporating regulatory information into epigenetic clocks may provide mechanistic insights into the aging process while improving the interpretability and predictive power.

Graphical abstract



Received: December 3, 2025. Revised: June 4, 2026. Accepted: June 6, 2026

© The Author(s) 2026. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

Introduction

DNAm has been repeatedly shown to be a robust biomarker to predict chronological age [1–3]. While the first generation of epigenetic clocks [1, 2] already exhibited a remarkable accuracy, second-generation clocks incorporated further data to even obtain predictions on mortality or life-span [4, 5]. Fueled by these encouraging results, the palette of epigenetic clock models has become increasingly diverse in recent years, capturing different aspects of DNAm during aging [6–9]. The ever-increasing availability of training samples will likely further enhance the models' accuracy in predicting chronological age [10]. However, a striking limitation of virtually all clock models is their inability to provide deeper insights into the biological aging process. In general, finding potentially causal and verifiable downstream biological mechanisms associated with learned features remains a challenging task [3, 11, 12]. This is particularly critical if such models are applied to assess the effects of interventions aimed at improving longevity or healthspan [11].

So far, most DNAm clocks have been developed using regularized regression methods such as elastic net [13] or LASSO [14]. During the typical training phase on a specific sample cohort, a relatively small set of CpGs is identified, whose methylation levels, combined and weighted in a linear model, are enough to predict chronological age [15]. Similar to other problems addressed with machine learning (ML) strategies, training aging clocks involves a large number of features (p) but relatively few samples (n). Typically, training of clocks involves computational examination of 10^4 – 10^5 theoretical features (often correlated among themselves) with only around 10^3 samples. To address the *big p, little n* or “ $p \gg n$ ” problem, regularization terms can help achieve sufficient model performance with a minimal set of features by shrinking the weights of all other features to zero, effectively removing them from the equation. Depending on the chosen regularization method and its parameters, even small changes in methylation values in the training samples can lead to the selection of different sets of CpGs with little overlap [12].

Besides selection stability problems, features that do not exhibit a particularly high correlation with chronological age may still be selected during training—a phenomenon referred to as noise accumulation [16, 17]. Consequently, the biological interpretability of epigenetic clocks may be limited. While the pre-selection of features is a frequently suggested strategy to mitigate instabilities and the accumulation of noisy features [16], it may also be suitable to focus on particular aspects of gene regulation, e.g., on the expression of transposable elements (TEs) [18]. Further, to deplete confounders and reduce prediction bias for specific age groups, the DNAm clocks should restrict the modeling on CpG sites exhibiting aging-associated DNAm signals [7, 19].

On the other hand, the mechanism by which DNAm is thought to influence gene regulation is through modifying transcription factor (TF) binding activity [20]. In particular, DNAm may sterically hinder or selectively enhance transcription factor (TF) binding at regulatory elements, depending on the specific properties of the TF [21]. Therefore, to improve the biological interpretability of the epigenetic clocks, it may be helpful to focus on CpGs overlapping known TF binding sites (TFBS).

Here, we present a framework for building a regulatory predictor, *TFMethyl Clock*, whose features may provide additional insight into the DNAm-based regulatory processes of

aging. The proposed feature selection step focuses on aging-correlated CpG sites located within experimentally known TFBSs. Furthermore, our modeling approach is preceded by a feature clustering step to enhance the robustness of predictions. Subsequently, we investigate the extent to which TFBSs are affected by age-specific methylation changes and characterize their potential target genes.

Materials and methods

DNAm dataset

The methylation dataset used to build the final predictor combined 15 publicly available DNA methylation cohorts from human whole blood, obtained as a curated set from [22]. The 7803 samples used were mostly classified as healthy, and the cohort was overall sex balanced. The minimum and maximum ages of the cohort were 0 and 101 years, respectively, while the mean was 39.10 years and the median was 40 years. The GSE105018 contributes the maximum 1651 samples to the cohort, and the GSE55763 contributes a minimum of 12 samples (Supplementary Fig. S6). To evaluate the performance on the 30% test split—we divided the entire cohort with training samples $n = 4995$ after removing a test-cohort of $n = 2143$ with a similar age distribution, and an untouched validation GSE84727 ($n = 665$) cohort—similarly as done by [22]. Overall, no imputation method was used for calculating the missing CpG probe values, as only the measured probes in all the samples were considered for further analysis. We also discard the samples where chronological age was unknown or if the sample had missing values for the given common set of probes.

RNA-seq dataset

The RNA-seq gene-expression data used in the study originates from the GTEx Analysis V10 release cohort, downloaded on 29/11/2024. A whole-blood tissue filter was used before downloading TPM and read count data from the GTEx portal, totaling 581 samples. The age groups analyzed ranged from 20 to 80 years, with a median of 55 years. We used the midpoint of each age group as the corresponding age for each sample (e.g. 25 for 20–29 years). Reads were aligned to the GRCh38/hg38 genome using the STAR aligner (v2.7.10a). RSEM (v1.3.3) was used for the transcript quantification by the standard GTEx processing pipeline.

Transcription factor binding sites (TFBS) annotation

We curated a genome-wide TFBS annotation for all the TFs whose overlapping motifs and binding site information were present in the JASPAR and cistrome databases, respectively [23, 24]. The entire hg38 genome JASPAR TF motifs and cistrome TF Chromatin Immunoprecipitation-sequencing (ChIP-seq) datasets were downloaded from the respective online repositories. Then, for the ChIP-seq TF binding profile, we filtered for “Blood” tissue type, and further concatenated different ChIP experiments together for each TF using bedtools (v2.30.0–56) [25]. Finally, we intersected ChIP-seq peaks with the overlapping sequence motifs for each TF using bedtools, thus creating a set of TFBS for all possible TFs. From the total of 828 TF motifs available in the JASPAR database and 239 TFs with ChIP-seq binding profiles in blood samples, a final common TFBS set of 94 TFs was selected.

CpG feature pre-filtering and clustering modules

The initial methylation matrix consists of 254 028 CpGs before any filtering or clustering steps, combining all measured probes in 7803 samples. The first filter for 94 TFs binding sites was intersected in R (v4.4.0) using `GenomicRanges::findOverlaps()` function, obtaining 98 474 CpGs. The genomic positions for CpGs were determined using the `IlluminaHumanMethylationEPICv2anno.20a1.hg38` R package (v1.0.0) - originally curated from the Illumina EPIC v2 methylation array annotation. The high-variable CpGs within the same-aged samples ($n = 1451$) were identified and removed from the set using the `sd` function from `{stats}` in R. Further, Spearman's correlation coefficient for each CpG while aging was calculated using the base R `cor{stats}` function. Then, all CpGs with a cutoff of $|\rho| > 0.5$ were retained, resulting in 14 006 CpGs for the next step.

These 14 006 CpGs, based on their changing methylation pattern across samples, were grouped into 4800 clusters, thus forming meta-features. The set of median centroid methylation values for meta-features was taken as the final input features for the elastic net-regression algorithm. The clustering was performed using `kmeans{stats}` function from base R. This step was implemented to improve model robustness by reducing susceptibility to technical noise at single CpG probes.

Enrichment analysis of age-associated CpGs in TFBS

Non-overlapping TFBS counts were calculated per transcription factor using `GenomicRanges`, and observed enrichments for age-correlated CpGs were compared to null distributions generated from 100 random CpG sets of equal size. Normality was assessed with the Shapiro–Wilk test, followed by either a parametric z-test or a non-parametric empirical test, with multiple testing correction by the Benjamini–Hochberg method.

The prediction model

The algorithm used was selected from the Gaussian elastic-net regression family with 90% ridge (L2) and 10% lasso (L1) penalties, and 30-fold inner cross-validation to tune the hyperparameter (λ). The function used to train the elastic-net linear regression model `glmnet::cv.glmnet()` on the methylation data was from the `glmnet` R library (v4.1–8) [26, 27]. Importantly, the target variable (chronological-age) was log-transformed to account for exponential-like dynamics (homoscedasticity) in many CpG methylation patterns with age and obtain a more bell-shaped sample distribution, as shown before [28].

Contemporary epigenetic clocks

We used the R `methylCIPHER` package (v0.2.0) to run previously published epigenetic age predictors with no imputation setting (“imputation = F”). Specifically, before analyzing age outcome results from the Epigenetic Time of Cancer 2 (EpiTOC2) clock, we divided the observations by a hundred to correct for the offset bias. For most of the epigenetic clocks tested, almost all the probes were present in the input methylation dataset for almost all the samples. The DunedinPACE clock was absent in the `methylCIPHER` package. Also, we were unable to run it independently on the validation cohort due to a significant number of missing CpG probes.

Beta methylation noise-simulation

To simulate noise, methylation values for each CpG j were perturbed by sampling from a Beta distribution parameterized using the empirical mean (μ_j) and variance (var_j) of that CpG across samples. In total, 10 to 40% of all CpG measurements on the array were perturbed increasingly by a factor of 10% for all CpGs passing the initial filtering criteria. The Beta distribution parameters, alpha (α_j) and beta (β_j), were estimated as:

$$\alpha_j = \mu_j^2 \cdot \left(\frac{1 - \mu_j}{var_j} - \frac{1}{\mu_j} \right) \quad \beta_j = \alpha_j \cdot \left(\frac{1}{\mu_j} - 1 \right)$$

Using these parameters, perturbed methylation values were generated independently for each CpG j using the `rbeta()` function from base R (`stats`): $x_{ij} \sim \text{Beta}(\alpha_j, \beta_j)$; where i indexes samples and 665 corresponds to the number of samples in the validation dataset, and j indexes that particular CpG.

10-fold cross-validation setup and target gene-connections

To narrow down to the most important age-predictor CpGs, we performed ten repetitions of stratified leave-one-set-out cross-validation. The CV folds were divided using the `createFolds()` function from the `caret` R library (v7.0–1), having a similar age distribution.

We connected the CpGs shared across CV models to their associated gene targets using `IlluminaHumanMethylationEPICv2anno.20a1.hg38` annotations. We used a custom function to associate a single gene target to each CpG, in the case where multiple genes connect to the CpG locus. In total, filtering for genes expressed in blood (average TPM in the GTEx cohort > 0.01) resulted in 267 genes.

GO enrichment analysis for target genes

The CpG-associated gene targets were analyzed for the possible gene ontology enrichment through the `enrichGO()` function from the R library `clusterProfiler` (v4.12.0). The annotation used for performing the analysis comes from the Bioconductor annotation data package `org.Hs.eg.db` (R library v3.19.1) for the human hg38 annotation. An appropriate enrichment background was established by taking the intersection of human blood expressed genes (average TPM > 0.01) and 450k array CpG-associated genes. The enrichment q-value threshold was 0.05, and we used the Benjamini–Hochberg method for P -value adjustment.

Clustering the target genes using age expression profiles

For the gene targets, we analyzed their expression trajectories during aging by clustering them in different clusters using the `clustool` (v1.18.0). All the target genes were clustered based on the Z-scores of the age group median expression with two iterations. After the first clustering round, ~66% of CpGs associated genes remained unclustered. Hence, another round of clustering was performed, resulting in a total of five clusters comprising 56% of all genes considered.

TE distribution analysis within TFBS

To analyze the distribution of transposable elements (TEs) across affected and protected TFs, we used the `DeepTools` suite. We first applied the `computeMatrix` function (v3.5.5) in

scale-regions mode to generate a matrix of signal intensities surrounding regions of interest. Specifically, we used bigWig files representing transposable elements (TEs) signal tracks for four major TE families: LINE, SINE, LTR, and DNA transposons from the human genome hg38. These tracks were processed against two sets of genomic regions—TFs affected and protected for age methylation changes—provided as BED files. We included 2 kb upstream and 2 kb downstream flanking regions (-a 2000 -b 2000) for each TFBS, and treated any missing data as zero to avoid bias (`-missingDataAsZero`). The average methylation signal across regions was plotted as a summary track above the heatmap using the mean value.

Age differential transcription factor footprinting

ATAC-seq profiles from $n = 4$ young (<25 years) and $n = 4$ old (>70 years) human blood samples were downloaded from GEO (GSE193142) [29]. Peaks were jointly called on all samples using MACS2 (v2.2.9.1, `-nomodel`, $q < 0.01$) for downstream analyses. Tn5 insertion bias-corrected ATAC-seq signals were generated using TOBIAS ATACCorrect function (v0.17.3), and footprinting scores were computed within combined peaks for each sample using TOBIAS FootprintScores using default parameters. Age footprinting signals were aggregated across biological replicates, and differential footprinting was calculated as $\log_2(\text{old/young})$ using deepTools. Differential footprinting profiles are shown centered on regions gaining or losing DNA methylation (± 5 kb), highlighting age-associated differences in transcription factor binding at methylation-dynamic loci.

Mortality analysis

Epigenetic clocks (Hanuum, Horvath2, PhenoAge, GrimAge2, and TFMethyl) were computed for 18 858 Generation Scotland participants using EPICv1 blood DNA methylation data collected between 2006 and 2011 [30]. Associations of epigenetic age acceleration with all-cause mortality (1517 deaths; ascertained through October 2023 via national record linkage) were assessed using Cox proportional hazards models adjusted for age and sex. For each clock, epigenetic age acceleration was defined as Z-standardized residuals from linear mixed-effects models regressing out chronological age and kinship (accounting for family structure). Cox models were fitted in R using the `survival` (v3.8–3) and `coxme` (v2.2–22) packages.

Visualization

Plots were created using the following R packages: `ggplot2` v3.5.1, `ggtrastr` v1.0.2, `ggridges` v0.5.6, and `ComplexHeatmap` v2.20.0. We also used `plotHeatmap` 3.5.5 to plot the TE density within TFs.

Results

Predictive performance and biological interpretability of epigenetic clocks

To scrutinize the biological underpinnings of established epigenetic clocks, we investigated their model CpGs in terms of their regulatory potential and direct aging correlations. Using DNAm data from the human blood aging compendium [22], we calculated the proportion of CpGs overlapping with experimentally validated TFBS (see Methods) and proportions

of CpGs with a strong chronological age correlation (Rank-based Spearman's measure $|\rho| > 0.5$).

For the analyzed methylation dataset, the Horvath clocks contain 18% (Horvath1) [2] and 28% (Horvath2) [31] model CpGs strongly correlated with age and overlapping with a TFBS at the same time (Fig. 1A). For the healthspan-optimized PhenoAge model [4], we find that only 10% of CpGs fulfill both criteria. The mortality-focused GrimAge2 [5] had a similar 10.3%, and the longitudinal cohort-trained third-generation DunedinPACE [8] had the smallest 7.7% proportion of CpGs with TFBS and strong age-correlations. In turn, with 69% (49 out of 71 CpGs), the Gaussian-process clock model (GP-clock) [22] shows the most substantial proportion of correlated and TFBS-associated CpGs. A closer look at age correlations reveals that about 70% of PhenoAge and 60% of Horvath1 CpGs correlate only weakly ($|\rho| < 0.4$) with age (Supplementary Fig. S1). Owing to specific correlation filters, 36% of the GP-clock model CpGs exhibit a very strong age-correlation ($|\rho| > 0.8$). In turn, PhenoAge and Horvath1 contain no such strongly correlated model CpGs (Supplementary Fig. S1). For the Horvath1 clock, this might be a consequence of multi-tissue model training or a smaller number of measured CpGs (27K array) or both.

In general, epigenetic clocks such as Horvath1 and GP-clock show high age-prediction accuracy with median absolute errors (MdAE) or mean absolute errors (MAE) of around a few years. To test whether clock models can accurately predict chronological age solely based on weakly correlated CpGs with no known association with TFBS, we trained 100 elastic net models on 7138 blood samples from [22] using 10 000 CpGs exhibiting a weak correlation with age ($|\rho| < 0.3$) and not overlapping with any TFBS. On average, the models show an MdAE of 3.77 years on an independent healthy validation cohort ($n = 665$, age 18–81 years) (Fig. 1B). This result is comparable to that of some second-generation clocks (MdAE PhenoAge of 3.79 years) and outperforms the first-generation clocks (MdAE Horvath1 of 5.95 years).

Also, we iteratively built a series of elastic-net models by sequentially removing CpGs identified and selected by earlier models as age-predictive from the training feature set for subsequent ones. Although the supposedly “best” age-predictive features were repeatedly removed from the feature set, we observed a markedly slow increase in age prediction errors (Supplementary Fig. S2). Only after approximately 50 rounds of feature removal, i.e., removing 88 007 CpG features (averaging ~ 1700 CpGs each iteration), did prediction errors reach the error of the random model.

In summary, our analyses suggest that, provided a sufficiently large training cohort, even features with limited age correlation and potentially modest biological impact can yield very accurate age predictions. One possible explanation is that unrestricted clock models may capture global characteristics of epigenomic aging that are substantially harder to link to a specific individual molecular process.

Gene expression changes linked to TFBS-associated CpGs

To further analyze the gene-regulatory roles of TFBS-associated features in the human blood aging data [22], we divided the available CpGs into two groups based on whether they overlapped with experimentally verified TFBS or not (see Methods). Overall, after initial pre-filtering (see Meth-




Figure 1. Systematic assessment of epigenetic clocks for regulatory potential and age prediction performance. **(A)** Proportion of biologically interpretable CpGs in the clock models. The circular bar plot shows the percentage of aging-correlated ($|\rho| > 0.5$) and TFBS-overlapping CpGs in the contemporary epigenetic clocks. **(B)** Absolute errors of epigenetic clock models trained on biologically less informative features (CpGs with age $|\rho| < 0.3$ & non-overlapping TFBS) tested on the validation cohort GSE84727. Boxes represent the IQR of absolute errors, with the median of all models highlighted as a colored dashed line (MdaE = 3.77 years). The dot plot below shows the respective average age-correlation score for the CpGs used for model training.

ods), 97 670 CpGs were categorized as part of the TF class, and 132 867 CpGs as part of the non-TF class (Fig. 2A, right). Subsequently, we analyzed the general chromatin annotation of the CpGs in both classes using ChromHMM predictions [32]. As expected, most CpGs in the TF class were located in the promoter region (34%), bivalent promoters (15%), active enhancers (15%), and transcription start sites (TSSs, 11%; Fig. 2A, left). Conversely, CpGs in the non-TF class were found in Polycomb-repressed regions (15%), active enhancers (15%), and bivalent promoters (11%). As a direct consequence of this enrichment, we found that twice as many blood-expressed genes ($n = 3380$ vs. $n = 1698$) were linked to CpGs in the TF class compared to the non-TF class (Fig. 2A, right). Using human blood expression data from six age groups spanning 20–80 years in the GTEx database, we next examined how well gene expression in both classes correlates with age. Genes linked to CpGs in the TF class showed a significantly more negative correlation, with a median $\rho = -0.76$ compared to the non-TF class median $\rho = -0.49$ (Wilcoxon test P -value = $2.2e-16$, Fig. 2B). Additionally, genes associated with TF class CpGs exhibit higher variability, with a median standard deviation of 0.28 versus 0.03 in the non-TF class (Fig. 2B). To empirically verify that CpGs in the TF-class are associated with the binding dynamics of TFs, we inspected chromatin accessibility at TFBSs losing or gaining DNA methylation ($|\rho| > 0.5$) and calculated TF footprint scores (Fig. 2C). Regions gaining DNA methylation during aging showed a relative reduction in TF footprint scores when comparing old to young samples, whereas regions losing DNA methylation exhibited the opposite trend.

While these results may not be surprising, it highlights that the selection of TFBS-associated CpGs enriches critical regulatory regions, exerting more potent effects on gene expression.

More importantly, genes associated with the TF class correlate more strongly with age, exhibit a higher variability, and show a general tendency towards age-related downregulation.

Age-associated DNA methylation differences at TFBSs

Next, we examined whether age-related DNA methylation changes affect specific TFs more than others. In total, we analyzed 94 distinct human TFs active in blood that bind to CpGs from the TF class. In the first step, using blood data, we calculated the correlation between DNAm changes and age, and identified 14 006 strongly age-correlated CpGs ($|\rho| > 0.5$). Subsequently, for each TF, we determined the relative frequency of TFBSs that overlap one of these strongly age-correlated CpGs compared to all TFBSs overlapping with any CpG. For comparison, we also calculated the relative frequencies of TFBS overlapping a random set of CpGs of the same size, averaged over a hundred permutations. Among others, TFBSs for *ZBED1*, *IRF8*, and *NFE2* have a larger proportion of age-correlated CpGs ($> 20\%$) compared to random CpGs ($\sim 13\%$). In contrast, we found that the TFBSs for *USF2*, *CREB1*, and *TP53* are depleted in age-correlated CpGs compared to the random CpG set (Fig. 3, left). Based on sampling distributions derived from 100 random sets of CpGs from the TF class (Fig. 3, right), we tested for the significance of proportions of age-correlated CpGs in the binding sites of each TF (see methods). In total, we found 14 statistically significant enrichments for TFs (adj. P -value < 0.05), of which eight were enriched for age-associated methylation changes (*ZBED1*, *NFE2*, *CEBPB*, *FOXP1*, *EGR1*, *SP1*, *PAX5*, and *MAZ*). In contrast, the binding sites of *RBPJ*, *NFIC*, *RELA*, *IKZF1*, *STAT3*, and *USF2* were significantly depleted in age-




Figure 2. Effects of TF binding at the CpG loci and their target genes during aging. **(A)** CpGs were grouped into two classes based on their overlap with TFBSs, i.e. TF class and non-TF class, respectively. Left: Stacked barplots representing ChromHMM annotations for CpGs in the TF class and in the non-TF class. Right: Number of target genes associated with CpGs from TF class and non-TF class that are expressed in blood. Schematics created with Biorender.com. **(B)** Top: Histogram (left) and boxplots (right) of Spearman's correlations of gene expression and chronological age for genes associated with CpGs in the TF class (blue) and the non-TF class (brown). The median Spearman's correlation scores of the TF class target genes with age are -0.76 (Q1, 25%: -0.93, Q3, 75%: -0.09) compared to -0.49 (Q1, 25%: -0.89, Q3, 75%: 0.50) of the non-TF class. Bottom: Density plot (left) and boxplots (right) of gene expression deviations during aging for genes associated with CpGs from the TF class and non-TF class. The median standard-deviation score in the TF class is 0.28 (Q1, 25%: 0.05, Q3, 75%: 0.81) compared to 0.03 in the non-TF class (Q1, 25%: 0.001, Q3, 75%: 0.27). Black dots in the boxplots represent outliers ($> 1.5 \times \text{IQR}$). **(C)** Differential $\log_2(\text{Old/Young})$ ATAC-seq-based TF-footprinting profiles (overlaid over $n = 4$ samples) centered at TFBS with surrounding ± 5 kb regions. Only TFBS gaining or losing DNA methylation with age ($|\rho| > 0.5$) are shown.

associated CpGs (Fig. 3, right). Interestingly, the expression of 10 of these 14 TFs ($P = 0.013$, Monte-Carlo simulation) itself was very strongly correlated with age in human blood ($|\rho| > 0.9$)—increasing for *EGR1* and *RELA*, and decreasing for *ZBED1*, *NFE2*, *SP1*, *PAX5*, *MAZ*, *IKZF1*, *STAT3*, and *USF2*.

The change in TF gene expression, combined with the methylation changes at its binding sites, indicates that the TFs could govern the expression of downstream genes in an age-dependent manner. Notably, five of these 14 TFs (*CEBPB*, *EGR1*, *SP1*, *RELA*, and *STAT3*) are listed by GenAge [33] as human aging-associated genes from experimental studies. In particular, *CEBPB*, *EGR1*, and *SP1*, identified as DNAm-affected in our analysis, play a role in regulating fat metabolism, acute stress response, and general transcriptional control [33]. In contrast, *STAT3* and *RELA*, whose binding sites are depleted in age-associated methylation changes, were functionally linked to chronic inflammation and immune dysregulation, respectively [34, 35]. Upon closer inspection, we found a small fraction of *RELA* binding sites (12%) undergoing substantial DNAm changes during aging ($|\rho| > 0.5$). Although not statistically significant, genes in the vicinity of these TFBS often showed larger gene expression changes with age as compared to genes with *RELA* binding sites that exhibited only weak $|\rho| < 0.5$ DNAm changes (Wilcoxon,

$P = 0.089$, Supplementary Fig. S3). Among the potential *RELA* targets with substantial DNAm changes, *IRAK4*, *OSM*, and *DUSP2* show strong correlation with age ($|\rho| \sim 0.9$, P -value < 0.05) and are involved in immune-related processes [36–38]. Similarly, we found target genes of *STAT3* with strong age correlation ($|\rho| \sim 0.9$, P -value < 0.05) and TFBS methylation changes. Among them are *ING4*, *MAPKAP1*, and *MAP3K5*, which are involved in inflammation [39, 40]. Interestingly, *STAT3* was previously reported to contribute to maintaining a youthful epigenetic state and promoting progenitor-like properties as shown in articular chondrocytes through its regulation of DNA methylation patterns [41].

Taken together, we could identify TFs that are vulnerable to DNAm changes during aging at their binding sites and TFs that may be protected from the DNAm changes at their binding sites. Moreover, genes linked to DNAm-vulnerable TFBSs also showed larger correlation with age.

Accuracy and robustness of TFBS-based epigenetic clock

Given the interesting properties of age-correlated CpGs overlapping with TFBSs, we next developed an epigenetic clock using these features (Fig. 4A). We initially selected CpGs




Figure 3. The enrichment of age-associated CpGs in TFBSs. Left: Comparison of TFBS overlap with age-associated CpGs vs. random CpGs. Bar plots show the proportion of transcription factor binding sites (TFBSs) affected by age-associated CpGs (dark brown) and by an equal number of randomly sampled CpGs (light brown) across selected TFs. Only the top 20 TFs with the highest observed enrichment and the bottom 20 with the lowest enrichment are shown. The x-axis indicates the proportion of affected TFBSs, and the y-axis lists TFs ranked by enrichment. Middle: Ridge density plots show the distribution of the proportion of TFBSs affected by randomly sampled non-age-associated CpGs ($n = 100$ iterations) for each transcription factor (TF). Vertical lines represent the observed proportion of TFBSs affected by age-associated CpGs. TFs with significant enrichment ($\text{padj} < 0.05$) are highlighted. The x-axis indicates the proportion of affected TFBSs for individual TFs. Right: Motifs based on TFBS overlapping age-associated CpGs.

from the methylation array that overlap with an experimentally validated TFBS (TF class, $n = 98\,474$). Then we filtered those for high age-correlation ($|\rho| > 0.5$, $n = 14\,006$, [Supplementary Fig. S4](#)). To further improve predictions and to reduce noise caused by lifestyle choices, environment, or by other factors (such as probe hybridization, melting temperature, or cellular composition) [7], we removed highly variable CpGs, defined by a standard deviation greater than 10% within the same age group ($n = 1451$, 0.57% of total CpGs) ([Supplementary Fig. S4](#), with an example CpG). To improve the robustness of any model built from this data, e.g. against missing values and technical noise, we additionally clustered the final set of 14 006 CpGs based on their age-associated methylation profiles using k-means (optimized $k = 4800$, see Materials and methods) ([Supplementary Fig. S5](#)).

After completing these preprocessing steps, we trained an elastic net regression model using the median DNAm value of each cluster to predict the chronological age (see Materi-

als and methods). The 4800 clusters used as input for training had an average cluster size of 3 CpGs. During the training and testing, we employed the DNAm blood Illumina 450k dataset from [22]. The dataset (see Materials and methods, [Supplementary Fig. S6](#)), comprising 7138 samples from 14 studies, was split into a training set ($n = 4997$) and a test set ($n = 2\,141$). The model achieved a correlation of r (Spearman's ρ) = 0.98 along with comparably low errors (MdAE = 1.86 years, RMSE = 4.72 years) in the test set (Fig. 4B). Subsequently, we merged the training and test sets to train our final age-predictive epigenetic clock model, “TFMethyl Clock.”

This final model, comprising 268 selected clusters covering 548 CpGs ([Supplementary Fig. S7](#)), was evaluated on a held-out validation cohort (GSE84727, $n = 665$) and compared to established epigenetic clocks. A comparison of our model with the established clocks reveals a pairwise overlap of at most 14 CpGs (3%), with the Horvath2 model ([Supplementary Fig. S8](#)). In the held-out validation cohort, our




Figure 4. Development and performance of TFMethyl Clock – TFBS-based epigenetic clock model. **(A)** In the TFMethyl Clock architecture, CpGs with large age-correlation and overlapping a TFBS are selected before modeling. Selected features are clustered based on their aging trajectories, and these clusters are used for modeling. Schematics created with Biorender.com. **(B)** Predicted age by TFMethyl Clock model (vertical axis) vs. the actual chronological age (horizontal axis) on the test set ($n = 2143$) in years. The dashed line represents the perfect positive correlation. **(C)** Distribution of the absolute errors (in years) on the validation cohort ($n = 665$) for established epigenetic clocks and for TFMethyl Clock, with the median absolute errors (MdAE) above each boxplot. The scatter plot shows the predicted age by TFMethyl Clock vs actual chronological age on the validation cohort. **(D)** Impact of CpG noise on epigenetic age predictions by TFMethyl Clock and GP-clock. Scatter plots show predicted vs. actual chronological ages using the original DNA methylation values, and using sequentially replacement of 10%, 20%, 30%, and 40% probed CpGs with random values drawn from fitted beta distributions. The dashed line represents perfect positive linear correlation. Boxplots depict the mean absolute error (MAE, in months) for each clock under different noise conditions. **(E)** Prediction errors (lines) with 95% confidence intervals (shaded area) by actual chronological age for selected epigenetic clocks (grey) and TFMethyl Clock.

model achieved a correlation of 0.97 (MdaE = 2.07 years, RMSE = 3.47 years) (Fig. 4C). In terms of both the MdaE and RMSE, our model outperforms all other tested clocks. The GP-clock, trained on the same training cohort as ours, was the runner-up with an MdaE of 2.31 years and RMSE of 3.72 years. Horvath2, a clock specifically designed to predict chronological age from blood and skin, achieved an MdaE of 2.82 years. The first-generation clocks, Horvath1 (MdaE = 5.95 years) and Hannum (MdaE = 6.29 years), followed suit. The epiTOC2 clock [42], originally designed to investigate the mitotic age of samples, had a weaker performance on the validation cohort with an MdaE of 16.84 yrs; while the GrimAge2, with 39.95 years of median error, performed worse (Fig. 4C). Importantly, this clock was developed to predict all-cause mortality hazard ratios (HR), not chronological age. In turn, when tested in the Generation Scotland cohort [30] for association of epigenetic age acceleration with all-cause mortality (see Methods), the TFMethyl clock returned a mild but significant HR (1.07, $P = 6 \times 10^{-4}$). In contrast, the Horvath2 clock predicted an HR of 1.05, while the best performance was achieved by PhenoAge (HR = 1.32) and GrimAge2 (HR = 1.59) (Supplementary Fig. S14).

To test the robustness of our cluster-based TFMethyl model for chronological age predictions, we randomly selected 10 to 40% of methylation array probes (25 403 to 101 612 CpGs) and replaced their original DNAm values in the validation samples with random values from a beta distribution fitted to the empirical shape parameters of a particular CpG site (see Methods). We then compared the changes in the performance of the GP-clock and our model on the validation set. The errors of the GP-clock [22] rapidly increased with the addition of noise. In particular, the errors increased to 34, 51, 67, and 80 months when replacing 10%, 20%, 30%, and 40% array CpGs, respectively. In turn, our model demonstrated relatively stable performance, with the errors increasing by no more than 21 months upon replacing 40% of the array CpGs (Fig. 4D). However, this is expected since the goal of the GP-clock was to develop a lean model that could be measured by targeted PCR.

It has been reported that most DNAm-based epigenetic clocks perform sub-optimally in predicting samples from extremely young and old individuals [43]. Our validation cohort benchmark aligns with this observation. While the largest errors in young samples were observed for the Hannum clock, a constant increase in errors for samples of age 60 + was observed for all studied clocks (Hannum, PhenoAge, Horvath2, and Gaussian, Fig. 4E). Compared to the two middle-age quartiles, Horvath2 showed a 2.7-fold increase in error for the oldest quartile, i.e. the top 25% of the oldest individuals. In the same analysis, PhenoAge and GP-clock had a 1.6-fold and 1.9-fold increase in error, respectively. Meanwhile, our approach showed the lowest increase of 1.16-fold.

Based on these benchmarks, our model performs better than established epigenetic clocks in terms of absolute errors, biases in both young and old ages, and robustness to noise.

Biological aging signatures reflected by TFMethyl clock

To assess our framework's capacity to achieve biological interpretability, we selected CpGs that are stable predictors of epigenetic age by training different models using a 10-fold cross-validation (CV) approach. Then, we extracted features from

these models with non-zero coefficients in all the CV folds, from here on called age-predictive CpGs ($n = 661$).

Subsequently, we analyzed the enrichment of age-predictive CpGs at TFBSs of 83 overlapping TFs. Compared to the background of input CpGs, binding sites of *ATF2*, *EBF1*, *NR2C2*, and *REST* showed the most substantial enrichment (Fisher's exact test, P -value < 0.05), suggesting a regulatory involvement in age-predictive methylation changes (Supplementary Fig. S9). The corresponding Gene-TF network revealed *NR2C2* harboring 221 age-predictive CpGs in its binding sites across the genome. For 41 CpGs, we also found that the expression values of their associated genes strongly correlate with age ($|\rho| > 0.8$, Fig. 5A).

Building on these findings, we analyzed all 267 blood-expressed genes that are potentially affected by the 661 age-predictive CpGs (see Methods, detailed in Supplementary Table S1). In summary, we found a significant enrichment in pathways (adjusted P -value = 0.02) associated with the positive regulation of interleukin-1 β production (*NOD1*, *NOD2*, *STAT3*, *TNF*, *FZD5* and *NLRP12*) and long-chain fatty acid metabolic process (*ACSL5*, *ACSL6*, *ADTRP*, *CYP11B1*, *PLA2G4C*, *CYP11A1*, *HACL1*, and *ELOVL2*, Fig. 5B). Strikingly, it has been shown that the *ELOVL2*, a gene involved in the elongation of polyunsaturated fatty acids, correlates with DNAm changes during aging across various tissues, making it one of the most potent biomarkers for aging [44, 45]. The *ELOVL2* gene is annotated to the CpG site cg16867657, which exhibits the largest average coefficient (0.72, nearly double the second biggest coefficient) in our models, indicating a strong age-associated hypermethylation.

To validate the influence of target genes associated with these age-predictive CpGs on the dynamics of the aging process, we examined their expression patterns in the whole blood data from the GTEx collection, comprising 581 samples from individuals aged 20 to 80 years. Notably, the data were obtained from individuals different than the DNAm cohort. Comprehensively, 200 out of all 267 (~75%) potential target genes had an age-Spearman correlation larger than > 0.5 (Supplementary Fig. S10). To further explore the specificity (e.g. up- or down-regulation) and timing of gene expression changes during aging, we used an unsupervised clustering approach for the gene expression trajectories. First, we grouped the samples into 10-year age intervals, calculated the median expression value for each gene within each group, and then applied the *clust* algorithm [46] to identify gene clusters with similar expression patterns over time. This clustering allows us to group and visualize shared patterns in age-related gene expression changes and reveals their dynamics across the lifespan. After two rounds of clustering, we got five distinct gene clusters (C1-C5), comprising 56% ($n = 149$) of the blood-expressed genes in our target set (Fig. 5C). The corresponding methylation trends for these monotonic gene-expression clusters (Supplementary Fig. S11) reveal a somewhat mixed increasing or decreasing trend with age, overall suggesting a locus-specific gene regulation. The largest cluster (C1, $n = 45$) exhibited a linear downregulation during aging and comprised genes such as *IRS2*, *CYP11B1*, and *NLRP12*. The clusters C2 ($n = 45$) and C3 ($n = 22$) showed a non-linear increase in expression during aging and comprised genes such as *PLA2G4C*, *ACSL6*, *NOD2* (C2), and *ACSL5* (C3) (Supplementary Fig. S12).

In particular, C3 contains genes that undergo an expression surge after the age of 60, whereas only a minor increase




Figure 5. Target genes of age-predictive CpGs selected by TFMethyl Clock. **(A)** Target genes of transcription factors whose binding sites significantly enrich for age-predictive CpGs. Node size reflects the target-gene set of each TF in the dataset, indicating its regulatory influence. Edges denote regulatory relationships from TFs to target genes via corresponding CpGs. Gene node color represents the gene-expression correlation with age. **(B)** Gene Ontology (GO) enrichment analysis for target genes of age-predictive CpGs. The size of the GO term node represents the number of genes associated with the respective GO term. The color gradient indicates the direction and magnitude of the averaged model beta coefficients. Only genes with TPM > 0.01 were included in the enrichment analysis. Schematics created with Biorender.com. **(C)** Gene expression trajectories of target genes of age-predictive CpGs across age groups. Each plot displays the expression dynamics (z-scored) in five distinct co-expression clusters (C1–C5) based on whole-blood RNA-seq data. Pie charts represent the size of each cluster relative to all the genes.

in expression can be observed before the age of 50. In contrast, cluster C4 ($n = 21$) exhibits a sudden downregulation between age 40 and 60 following a comparably stable expression in early life (example gene *PPP1R12C* shown in [Supplementary Fig. S12](#)). Lastly, the C5 ($n = 16$) is characterized by stable expression until midlife, followed by rapid downregulation after the age of 60, as is the case for *CYP1A1* ([Supplementary Fig. S12](#)). The non-linear trajectories have been recently described in several studies in multi-omics, including DNAm and transcriptome during aging in both humans and mice [47, 48]. Notably, we found several age-predictive CpGs whose target genes also display non-linear expression changes, supporting recent findings on the non-linearity of aging [47].

Discussion

Here, we provided evidence that models with CpGs depleted in critical regulatory regions and even lacking strongly corre-

lated aging signals can produce excellent age predictions. This finding draws parallels with previous reports showing that just an increased age-associated noise is sufficient to build accurate aging clocks [49]. Therefore, it may be less surprising that investigations into the biological explainability of DNAm clocks often remained inconclusive [3, 50, 51]. To remedy this shortcoming, we have investigated the potential to build models with CpGs that particularly overlap TFBs and have a higher individual correlation with age. CpG methylation changes are known to mediate regulatory changes via TF inhibition [20]. Our ATAC-seq-based footprinting analysis supports the view that CpG methylation may also alter TF-binding patterns during aging or, alternatively, indicate these differential binding patterns. Filtering of CpGs further limits the feature space and consequently reduces the possible inconsistency of feature selection in the elastic net. Interestingly, the applied constraints not only led to improved performance on the test data but also in the independent validation cohort. Among the tested clocks, our TFMethyl Clock model displays a superior perfor-

mance in terms of chronological age predictions. Instead of training on CpGs directly, we generated CpG meta-features using a simple k-means clustering. This attribute of our approach further alleviates the $n \gg p$ problem, deals more effectively with the multicollinearity of CpG features, and, in our view, is more intuitive than other strategies, such as principal component (PC) based strategies, e.g. used in PC clocks [52]. This may contribute to TFMethyl Clock's robustness to noise and performance for samples at the extremes of the age spectrum, unlike most other epigenetic clocks.

More importantly, our strategy led to the identification of TFBS for specific TFs that significantly enrich age-correlated CpGs, while others appear to be protected. Among the more frequently affected TFBSs are the binding sites for the Zinc Finger BED-Type Containing 1 protein (ZBED1). The protein encoded in the pseudoautosomal region 1 of the X and Y chromosomes acts as a transcription factor for genes involved in cell proliferation. Additionally, it acts as an E3-type small ubiquitin-like modifier (SUMO) ligase. Belonging to a larger family of related proteins that may use the same binding sites, further research is necessary to understand the potential involvement of ZBED1 or other family members in the aging process.

On the other hand, age-associated CpGs were significantly depleted in binding sites of the transcription factor *USF2*, which is known for its involvement in cellular senescence [53]. Likewise, the TFBSs of *CREB1*, a methylation-sensitive TF that can bind to unmethylated sites to upregulate LTR repeat transcription [20], were depleted in age-associated CpGs. Further, we observed that all TFs depleted for age-associated CpGs show an enrichment for LTRs among other transposable elements (Supplementary Fig. S13). While the significant depletion of age-associated CpGs in these TFBSs is challenging to interpret, one might speculate that these TFBSs are subject to tighter control of DNA methylation, which actively protects these regions against age-associated dynamics and activation, at least in the respective cell type. We believe that these results encourage follow-up studies to focus on methylation-based properties of individual TFs while aging.

Notably, three-quarters of genes associated with age-predictive CpGs found by TFMethyl Clock also exhibit age-related changes in expression. This likely results from our selection process, which effectively enriches regulatory regions. The analysis of genes linked to age-predictive CpGs revealed enrichments in pathways related to interleukin-1, inflammation, and the modulation of immune response—critical pathways impacted by the aging process [54]. Genes such as *NODs*, *STAT3*, and *TNF* are well-established components of innate immunity and inflammation [55, 41, 56]. The second enriched category, “long-chain fatty acid metabolic process,” aligns with emerging new literature connecting lipid metabolism to lifespan regulation and cellular senescence [57]. In particular, the recovery of the *ELOVL2* locus, a gene whose methylation is known to be a robust age biomarker, serves as a positive control in support of our modeling approach. Likewise, we found an enrichment of age-predictive CpGs in binding sites of *NR2C2*—a factor that has recently been associated with premature aging, and significantly shortened lifespan and accelerated cellular senescence phenotype [58, 59].

Tracing the expression trajectories of genes associated with age-predictive CpGs, we identified several gene expression clusters that did not enrich any particular biological function. Interestingly, the majority of the gene clusters exhibited non-

linear behavior during aging, typically shifting during midlife or in later life. Further studies are necessary to determine whether these changes can be validated and support previously described non-linearities in aging [47].

In this study, we used methylation data from the Illumina 450k array, which only covers a small portion of the approximately 28 million CpGs in the human genome. Using the latest methylation arrays or applying the presented strategy to sequencing-based methylation measurements could offer a more comprehensive view.

In contrast to other approaches, such as the GP-clock, TFMethyl Clock heavily relies on additional biological information, i.e. comprehensive and accurate TFBS data. This dependency naturally limits its application to well-annotated genomes and is prone to bias. As a direct consequence of its design, it does not consider other regulatory CpGs that are not directly involved in TF binding. Since TF binding is often tissue-specific, we believe studying other tissues is necessary to better understand global, recurrent aging signatures, as well as those that are more tissue-specific. This approach would be especially helpful in understanding the role of the identified transcription factors' dynamics in aging.

Ethics

All components of GS received ethical approval from the NHS Tayside Committee on Medical Research Ethics (REC Reference Number: 05/S1401/89). GS has also been granted Research Tissue Bank status by the East of Scotland Research Ethics Service (REC Reference Number: 20-ES-0021). All participants provided broad and enduring written informed consent for biomedical research. This study was performed in accordance with the Helsinki Declaration.

Acknowledgements

We thank all members of Hoffman lab for insightful discussions, and the anonymous reviewers for their valuable critics and ideas. Graphical abstract was created with Biorender.com and SVGRepo.com.

Author contributions: T.P. (Data curation [equal], Formal Analysis [equal], Methodology [equal], Software [equal], Visualization [equal], Writing – original draft [equal], Writing – review & editing [equal]), R.S. (Data curation [supporting], Writing – review & editing [supporting]), K.R. (Data curation [equal], Software [equal]), M.V. (Data curation [supporting]), A.R. (Data curation [supporting]), R.E.M. (Data curation [supporting], Formal Analysis [supporting]), H.K. (Supervision [supporting], Writing – review & editing [supporting]), T.K. (Data curation [supporting], Writing – review & editing [supporting]), S.H. Conceptualization [equal], Funding acquisition [lead], Methodology [equal], Supervision [equal], Writing – original draft [equal], Writing – review & editing [equal] A.v.B. (Conceptualization [lead], Funding acquisition [equal], Methodology [equal], Supervision [lead], Writing – original draft [equal], Writing – review & editing [equal]).

Supplementary data

Supplementary data is available at NAR online.

Conflict of interest

None declared.

Funding

The work of T.P., A.v.B. and S.H. was supported by the DFG ProMoAge RTG 2155. The FLI is a member of the Leibniz Association and is financially supported by the Federal Government of Germany and the State of Thuringia. S.H. was supported by the DFG projects 418080850, 418087534, 498394458, and 460757154, German Network for Bioinformatics Infrastructure (W-de.NBI-017), the Carl Zeiss Foundation in the context of the IMPULS consortium (P2019-01-006), and the Leibniz Research Alliance „Resilient Aging“ (LFV-2021-2-LIR).

G.S. was funded by a grant from the Chief Scientist Office of the Scottish Government Health Directorates (CZD/16/6) and the Scottish Funding Council (HR03006). Genotyping and DNA methylation profiling of the GS samples was carried out by the Genetics Core Laboratory at the Edinburgh Clinical Research Facility, University of Edinburgh, Scotland, and was funded by the Medical Research Council UK and the Wellcome Trust (Wellcome Trust Strategic Award ‘Stratifying Resilience and Depression Longitudinally’ Reference 104036/Z/14/Z). DNA methylation profiling was also funded in part by the Wellcome Trust (Reference 220857/Z/20/Z). Funding to pay the Open Access publication charges for this article was provided by Leibniz Institute on Ageing - Fritz Lipmann Institute.

Data availability

This research did not generate any new data. All datasets are freely available from public databases. The methylation dataset used was borrowed from [22]. The RNA-seq gene expression data are from the GTEx portal (https://gtexportal.org/home/downloads/adultgtex/bulk_tissue_expression).

The code for utilizing TFMethyl clock and analysis is deposited at https://github.com/Hoffmann-Lab/Enhancing_the_performance_and_interpretability_of_epigenetic_clocks (<https://doi.org/10.5281/zenodo.20611646>). Processed data generated in the work is available via Zenodo at <https://doi.org/10.5281/zenodo.18672348>.

According to the terms of consent for Generation Scotland participants, access to data must be reviewed by the Generation Scotland Access Committee. Applications should be made to genscot@ed.ac.uk and normally take up to 6 weeks for approval. Further details can be found at <https://genscot.ed.ac.uk/for-researchers/access/>.

References

- Hannum G, Guinney J, Zhao L *et al.* Genome-wide methylation profiles reveal quantitative views of human aging rates. *Mol Cell* 2013;49:359–67. <https://doi.org/10.1016/j.molcel.2012.10.016>
- Horvath S. DNA methylation age of human tissues and cell types. *Genome Biol* 2013;14:R115. <https://doi.org/10.1186/gb-2013-14-10-r115>
- Teschendorff AE, Horvath S. Epigenetic ageing clocks: statistical methods and emerging computational challenges. *Nat Rev Genet* 2025;26:350–68. <https://doi.org/10.1038/s41576-024-00807-w>
- Levine ME, Lu AT, Quach A *et al.* An epigenetic biomarker of aging for lifespan and healthspan. *Aging (Albany, NY)* 2018;10:573–91.
- Lu AT, Binder AM, Zhang J *et al.* DNA methylation GrimAge version 2. *Aging (Albany, NY)* 2022;14:9484–549.
- Galkin F, Mamoshina P, Kochetov K *et al.* DeepMAGE: a methylation aging clock developed with deep learning. *Aging Dis* 2021;12:1252–62. <https://doi.org/10.14336/AD.2020.1202>
- Ying K, Liu H, Tarkhov AE *et al.* Causality-enriched epigenetic age uncouples damage and adaptation. *Nat Aging* 2024;4:231–46. <https://doi.org/10.1038/s43587-023-00557-0>
- Belsky DW, Caspi A, Corcoran DL *et al.* DunedinPACE, a DNA methylation biomarker of the pace of aging. *eLife* 2022;11:e73420. <https://doi.org/10.7554/eLife.73420>
- Weidner CI, Lin Q, Koch CM *et al.* Aging of blood can be tracked by DNA methylation changes at just three CpG sites. *Genome Biol* 2014;15:R24. <https://doi.org/10.1186/gb-2014-15-2-r24>
- Zhang Q, Vallergera CL, Walker RM *et al.* Improved precision of epigenetic clock estimates across tissues and its implication for biological ageing. *Genome Med* 2019;11:54. <https://doi.org/10.1186/s13073-019-0667-1>
- Moqri M, Herzog C, Poganik JR *et al.* Biomarkers of aging for the identification and evaluation of longevity interventions. *Cell* 2023;186:3758–75. <https://doi.org/10.1016/j.cell.2023.08.003>
- Mei X, Blanchard J, Luellen C *et al.* Fail-tests of DNA methylation clocks, and development of a noise barometer for measuring epigenetic pressure of aging and disease. *Aging (Albany, NY)* 2023;15:8552–75.
- Zou H, Hastie T. Regularization and variable selection via the elastic net. *J R Stat Soc, Ser B* 2005;67:301–20. <https://doi.org/10.1111/j.1467-9868.2005.00503.x>
- Tibshirani R. The lasso method for variable selection in the Cox model. *Statist Med* 1997;16:385–95. [https://doi.org/10.1002/\(SICI\)1097-0258\(19970228\)16:4<385::AID-SIM380>3.0.CO;2-3](https://doi.org/10.1002/(SICI)1097-0258(19970228)16:4<385::AID-SIM380>3.0.CO;2-3)
- Horvath S, Raj K. DNA methylation-based biomarkers and the epigenetic clock theory of ageing. *Nat Rev Genet* 2018;19:371–84. <https://doi.org/10.1038/s41576-018-0004-3>
- Fan J, Samworth R, Wu Y. Ultrahigh dimensional feature selection: beyond the linear model. *J Mach Learn Res* 2009;10:2013–38.
- Fan J, Fan Y. High dimensional classification using features annealed independence rules. *Ann Statist* 2008;36:2605–37. <https://doi.org/10.1214/07-AOS504>
- Morandini F, Lu JY, Rechsteiner C *et al.* Transposable element 5mC methylation state of blood cells predicts age and disease. *Nat Aging* 2025;5:193–204. <https://doi.org/10.1038/s43587-024-00757-2>
- Dabrowski JK, Yang EJ, Crofts SJC *et al.* Probabilistic inference of epigenetic age acceleration from cellular dynamics. *Nat Aging* 2024;4:1493–507. <https://doi.org/10.1038/s43587-024-00700-5>
- Kaluscha S, Domcke S, Wirbelauer C *et al.* Evidence that direct inhibition of transcription factor binding is the prevailing mode of gene and repeat repression by DNA methylation. *Nat Genet* 2022;54:1895–906. <https://doi.org/10.1038/s41588-022-01241-6>
- Yin Y, Morgunova E, Jolma A *et al.* Impact of cytosine methylation on DNA binding specificities of human transcription factors. *Science* 2017;356:eaaj2239. <https://doi.org/10.1126/science.aaj2239>
- Varshavsky M, Harari G, Glaser B *et al.* Accurate age prediction from blood using a small set of DNA methylation sites and a cohort-based machine learning algorithm. *Cell Rep Methods* 2023;3:100567. <https://doi.org/10.1016/j.crmeth.2023.100567>
- Rauluseviciute I, Riudavets-Puig R, Blanc-Mathieu R *et al.* JASPAR 2024: 20th anniversary of the open-access database of transcription factor binding profiles. *Nucleic Acids Res* 2024;52:D174–82. <https://doi.org/10.1093/nar/gkad1059>
- Taing L, Dandawate A, L’Yi S *et al.* Cistrome Data Browser: integrated search, analysis and visualization of chromatin data

- Nucleic Acids Res* 2024;52:D61–6.
<https://doi.org/10.1093/nar/gkad1069> <https://doi.org/10.1093/bioinformatics/btq033>
25. Quinlan AR, Hall IM. BEDTools: a flexible suite of utilities for comparing genomic features. *Bioinformatics* 2010;26:841–2. <https://doi.org/10.1093/bioinformatics/btq033>
 26. Friedman J, Hastie T, Tibshirani R. Regularization paths for generalized linear models via coordinate descent. *J Stat Soft* 2010;33:1–22. <https://doi.org/10.18637/jss.v033.i01>
 27. Tay JK, Narasimhan B, Hastie T. Elastic net regularization paths for all generalized linear models. *J Stat Soft* 2023;106:1–31. <https://doi.org/10.18637/jss.v106.i01>
 28. Lu AT, Fei Z, Haghani A *et al*. Universal DNA methylation age across mammalian tissues. *Nat Aging* 2023;3:1144–66. <https://doi.org/10.1038/s43587-023-00462-6>
 29. Morandini F, Rechsteiner C, Perez K *et al*. ATAC-clock: an aging clock based on chromatin accessibility. *Geroscience* 2024;46:1789–806. <https://doi.org/10.1007/s11357-023-00986-0>
 30. Walker RM, McCartney DL, Carr K *et al*. Data resource profile: whole-blood DNA methylation resource in Generation Scotland (MeGS). *Int J Epidemiol* 2025;54:dyaf091. <https://doi.org/10.1093/ije/dyaf091>
 31. Horvath S, Oshima J, Martin GM *et al*. Epigenetic clock for skin and blood cells applied to Hutchinson Gilford Progeria Syndrome and ex vivo studies. *Aging (Albany, NY)* 2018;10:1758–75.
 32. Vu H, Ernst J. Universal annotation of the human genome through integration of over a thousand epigenomic datasets. *Genome Biol* 2022;23:9. <https://doi.org/10.1186/s13059-021-02572-z>
 33. de Magalhães JP, Abidi Z, Dos Santos GA *et al*. Human Ageing Genomic Resources: updates on key databases in ageing research. *Nucleic Acids Res* 2024;52:D900–8.
 34. García-García VA, Alameda JP, Page A *et al*. Role of NF- κ B in ageing and age-related diseases: lessons from genetically modified mouse models. *Cells* 2021;10:1906.
 35. Lee H, Herrmann A, Deng J-H *et al*. Persistently activated Stat3 maintains constitutive NF-kappaB activity in tumors. *Cancer Cell* 2009;15:283–93. <https://doi.org/10.1016/j.ccr.2009.02.015>
 36. Suzuki N, Suzuki S, Millar DG *et al*. A critical role for the innate immune signaling molecule IRAK-4 in T cell activation. *Science* 2006;311:1927–32. <https://doi.org/10.1126/science.1124256>
 37. Wolf CL, Pruett C, Lighter D *et al*. The clinical relevance of OSM in inflammatory diseases: a comprehensive review. *Front Immunol* 2023;14:1239732. <https://doi.org/10.3389/fimmu.2023.1239732>
 38. Lancaster GI, Kraakman MJ, Kammoun HL *et al*. The dual-specificity phosphatase 2 (DUSP2) does not regulate obesity-associated inflammation or insulin resistance in mice. *PLoS One* 2014;9:e111524. <https://doi.org/10.1371/journal.pone.0111524>
 39. Yang Y, Liu Y, He X *et al*. ING4 alleviated lipopolysaccharide-induced inflammation by regulating the NF- κ B pathway via a direct interaction with SIRT1. *Immunol Cell Biol* 2020;98:127–37. <https://doi.org/10.1111/imcb.12308>
 40. Morris BJ, Chen R, Donlon TA *et al*. Lifespan extension conferred by mitogen-activated protein kinase kinase kinase 5 (MAP3K5) longevity-associated gene variation is confined to at-risk men with a cardiometabolic disease. *Aging (Albany, NY)* 2021;13:7953–74.
 41. Sarkar A, Liu NQ, Magallanes J *et al*. STAT3 promotes a youthful epigenetic state in articular chondrocytes. *Aging Cell* 2023;22:e13773. <https://doi.org/10.1111/accel.13773>
 42. Teschendorff AE. A comparison of epigenetic mitotic-like clocks for cancer risk prediction. *Genome Med* 2020;12:56. <https://doi.org/10.1186/s13073-020-00752-3>
 43. El Khoury LY, Gorrie-Stone T, Smart M *et al*. Systematic underestimation of the epigenetic clock and age acceleration in older subjects. *Genome Biol* 2019;20:283.
 44. Slieker RC, Relton CL, Gaunt TR *et al*. Age-related DNA methylation changes are tissue-specific with ELOVL2 promoter methylation as exception. *Epigenetics Chromatin* 2018;11:25. <https://doi.org/10.1186/s13072-018-0191-3>
 45. Chen D, Chao DL, Rocha L *et al*. The lipid elongation enzyme ELOVL2 is a molecular regulator of aging in the retina. *Aging Cell* 2020;19:e13100. <https://doi.org/10.1111/accel.13100>
 46. Abu-Jamous B, Kelly S. Clust: automatic extraction of optimal co-expressed gene clusters from gene expression data. *Genome Biol* 2018;19:172. <https://doi.org/10.1186/s13059-018-1536-8>
 47. Shen X, Wang C, Zhou X *et al*. Nonlinear dynamics of multi-omics profiles during human aging. *Nat Aging* 2024;4:1619–34. <https://doi.org/10.1038/s43587-024-00692-2>
 48. Olecka M, van Bömmel A, Best L *et al*. Nonlinear DNA methylation trajectories in aging male mice. *Nat Commun* 2024;15:3074. <https://doi.org/10.1038/s41467-024-47316-2>
 49. Meyer DH, Schumacher B. Aging clocks based on accumulating stochastic variation. *Nat Aging* 2024;4:871–85. <https://doi.org/10.1038/s43587-024-00619-x>
 50. Bell CG, Lowe R, Adams PD *et al*. DNA methylation aging clocks: challenges and recommendations. *Genome Biol* 2019;20:249. <https://doi.org/10.1186/s13059-019-1824-y>
 51. Levine ME, Higgins-Chen A, Thrush K *et al*. Clock work: deconstructing the epigenetic clock signals in aging, disease, and reprogramming. *Biorxiv*, <https://doi.org/10.1101/2022.02.13.480245>, 15 February 2022, preprint: not peer reviewed.
 52. Higgins-Chen AT, Thrush KL, Wang Y *et al*. A computational solution for bolstering reliability of epigenetic clocks: implications for clinical trials and longitudinal tracking. *Nat Aging* 2022;2:644–61. <https://doi.org/10.1038/s43587-022-00248-2>
 53. Kang T, Moore EC, Kopania EEK *et al*. A natural variation-based screen in mouse cells reveals USF2 as a regulator of the DNA damage response and cellular senescence. *G3 (Bethesda)* 2023;13:jkad091. <https://doi.org/10.1093/g3journal/jkad091>
 54. Li X, Li C, Zhang W *et al*. Inflammation and aging: signaling pathways and intervention therapies. *Sig Transduct Target Ther* 2023;8:239. <https://doi.org/10.1038/s41392-023-01502-8>
 55. Almeida-da-Silva CLC, Savio LEB, Coutinho-Silva R *et al*. The role of NOD-like receptors in innate immunity. *Front Immunol* 2023;14:1122586. <https://doi.org/10.3389/fimmu.2023.1122586>
 56. Idriss HT, Naismith JH. TNF alpha and the TNF receptor superfamily: structure-function relationship(s). *Microsc Res Tech* 2000;50:184–95. [https://doi.org/10.1002/1097-0029\(20000801\)50:3\(184::AID-JEMT2\)3.0.CO;2-H](https://doi.org/10.1002/1097-0029(20000801)50:3(184::AID-JEMT2)3.0.CO;2-H)
 57. Johnson AA, Stolz A. The role of lipid metabolism in aging, lifespan regulation, and age-related disease. *Aging Cell* 2019;18:e13048. <https://doi.org/10.1111/accel.13048>
 58. Lee Y-F, Liu S, Liu N-C *et al*. Premature aging with impaired oxidative stress defense in mice lacking TR4. *Am J Physiol: Endocrinol Metab* 2011;301:E91–8.
 59. Gong D, Liu Z, Kang R *et al*. Deficiency of NR2C2 accelerates senescence of testicular Leydig cells and infertility in male mice. *Cell Signalling* 2026;138:112284. <https://doi.org/10.1016/j.cellsig.2025.112284>