Cell-type specific epigenetic clocks to quantify biological age at cell-type resolution

Huige Tong [1,$], Xiaolong Guo [1,$], Macsue Jacques [2], Qi Luo [1] , Nir Eynon [2] and Andrew E. Teschendorff [1,*]

1. CAS Key Laboratory of Computational Biology, CAS-MPG Partner Institute for Computational Biology, Shanghai Institute of Nutrition and Health, Shanghai Institute for Biological Sciences, University of Chinese Academy of Sciences, Chinese Academy of Sciences, 320 Yue Yang Road, Shanghai 200031, China.
2. Australian Regenerative Medicine Institute (ARMI), Faculty of Medicine, Nursing and Health Sciences, Monash University, Level 1, 15 Innovation Walk, Clayton, VIC 3800, Australia.

[$] Equal contribution

* Corresponding author: Andrew E. Teschendorff   andrew@sinh.ac.cn

Abstract

The ability to accurately quantify biological age could help monitor and control healthy aging. Epigenetic clocks have emerged as promising tools for estimating biological age, yet so far, most of these clocks have been developed from heterogeneous bulk tissues, and are thus composites of two aging processes, one reflecting the change of cell-type composition with age and another reflecting the aging of individual cell-types. There is thus a need to dissect and quantify these two components of epigenetic clocks, and to develop epigenetic clocks that can yield biological age estimates at cell-type resolution. Here we demonstrate that in blood and brain, approximately 35% of an epigenetic clock's accuracy is driven by underlying shifts in lymphocyte and neuronal subsets, respectively. Using brain and liver tissue as prototypes, we build and validate neuron and hepatocyte specific DNA methylation clocks, and demonstrate that these cell-type specific clocks yield improved estimates of chronological age in the corresponding cell and tissue-types. We find that neuron and glia specific clocks display biological age acceleration in Alzheimer's Disease with the effect being strongest for glia in the temporal lobe. The hepatocyte clock is found accelerated in liver under various pathological conditions. In contrast, non-cell-type specific clocks do not display biological age-acceleration, or only do so more marginally. In summary, this work highlights the importance of dissecting epigenetic clocks and quantifying biological age at cell-type resolution.

Keywords: DNA methylation; epigenetic clocks; cell-type deconvolution; biological aging; Alzheimer's Disease; obesity

Introduction

Individuals of the same chronological age may display marked differences in health and future disease risk [1]. These differences are attributed to variations in biological age, and whilst a number of biomarkers for biological age have been developed, these are mostly measured in, or restricted to, easily accessible tissues like blood [1-4]. More recently, epigenetic clocks based on DNA methylation (DNAm) marks [5-7] have emerged as powerful tools to predict chronological age, and to some extent also biological age [3,8-16]. However, many challenges remain that limit the practical utility and interpretation of these epigenetic clocks [9,10,17].

Prominent among these challenges is cell-type heterogeneity (CTH) [18,19]: most if not all epigenetic clocks have been developed or trained from heterogeneous bulk-tissues that are composed of many different cell-types. This CTH has two undesirable effects.

First, it can distort biological age estimates by underlying variations in cell-type fractions. As a concrete example, whole blood, as well as solid tissues with a significant immune-cell infiltration, are all characterized by an age-related shift from naïve to mature T-cell fractions [20,21] [22,23], which as shown by Jonkman et al can confound biological age estimates in immune cell subtypes [20]. Another example is the case of brain-tissue, with

early studies reporting correlations of an accelerated Horvath-clock[24] with neurodegenerative features (e.g. neurotic plaques or amyloid load) [25,26] in the prefrontal cortex, or with genetic variants associated with neurodegenerative diseases [27-29], but with more recent studies attributing such age-accelerations to changes in underlying cell-type proportions [30,31] [32]. The second undesirable effect of CTH is that, even if we adjust for it *a-posteriori*, as it is often done in blood tissue by adjusting epigenetic clock values for variations in estimated immune cell fractions [24], the ensuing epigenetic clock estimates are still averages over all the distinct cell-types in the tissue. Thus, standard *a-posteriori* adjustment for cell-type fractions does not inform us about the biological aging of individual cell-types. This is a serious limitation for understanding which cell-types age faster and contribute most to age-related diseases, or for identifying the cell-types that respond best to cellular rejuvenation and other healthy-aging promoting interventions [33-38].

Thus, current epigenetic clocks are composites, reflecting a mixture of two underlying age-related processes: an extrinsic one reflecting age-related variations of a tissue's cell-type composition, and an intrinsic one, reflecting age-related DNAm changes in the individual cell-types of the tissue, the latter being itself a composite process, as each cell-type may age at a different rate. Importantly, the relative contribution of extrinsic and intrinsic aging processes to an epigenetic clock's accuracy has not yet been quantified, and the only attempts so far at quantifying biological age at cell-type resolution have come from studies that have applied epigenetic clocks derived from bulk-tissue (e.g. liver) to sparse and unreliable single-cell DNAm data (e.g. hepatocytes) [39].

Here we advance the epigenetic clock field in two ways. First, we dissect epigenetic clocks into their extrinsic and intrinsic components, rigorously quantifying these aging components in two of the most frequently profiled tissues (whole blood and brain), demonstrating that extrinsic aging accounts for a substantial proportion of an epigenetic clock's predictive accuracy. Second, we present a computational strategy for building cell-type specific epigenetic clocks that can yield biological age estimates at cell-type resolution. This strategy explicitly adjusts for CTH whilst training the clock, leveraging recent state-of-the-art cell-type deconvolution methods [21,40] to identify cell-type specific age-associated DNAm changes [41]. As such this is one of the first studies to characterize age-associated DNAm changes at cell-type resolution. We

illustrate our strategy in brain and liver tissue, demonstrating how cell-type specific epigenetic clocks from these tissues can improve tissue-specific estimation of chronological and biological age.

Results

Dissecting and quantifying the two components of epigenetic clocks

Epigenetic clocks are composites of two underlying age-related processes, an extrinsic one that reflects age-associated variations in cell-type composition and an intrinsic one that captures the age-associated DNAm changes in individual cell-types (**Fig.1a**). Thus, the predictive accuracy (R^2 value) of a clock can be composed as the sum of an extrinsic and intrinsic accuracy, so that the extrinsic R^2 value can be estimated from the predictive accuracies of two clocks: one adjusted for CTH, and another which is not (**Fig.1a**). To quantify the relative rates of extrinsic and intrinsic aging, we assembled a large collection of whole blood Illumina DNAm datasets, using two of the largest cohorts to train Elastic Net age-predictors unadjusted and adjusted for CTH, respectively (**Methods**). Adjustment for CTH was performed at the resolution of 12 immune cell subtypes using a validated DNAm reference matrix encompassing naïve and memory B-cells, naïve and memory CD4T-cells, naïve and memory CD8T-cells, T-regulatory cells, natural-killer cells, monocytes, neutrophils, eosinophils and basophils [21] (**Methods**). We also built an elastic net clock adjusting for 9 immune-cell subtypes, by merging together the estimated fractions of the naïve and memory lymphocyte subsets, allowing us to assess the importance of the known age-related shift from naïve to mature T-cell subsets in driving an epigenetic clock's accuracy (**Methods**) [21]. We then estimated R^2 values of the 3 clocks in 11 independent whole blood DNAm datasets (**Methods**). This revealed that the unadjusted clock achieved substantially higher R^2 values than the two adjusted clocks, specially when compared to the clock adjusted for all 12 immune cell types (**Fig.1b, SI figs.S1-2**). We estimated that the intrinsic aging component accounts for approximately 61% of the predictive accuracy of the unadjusted clock (**Fig.1b**), which means that extrinsic aging accounts for a substantial fraction of a clock's accuracy (39%). Of note, not adjusting for relative shifts in naïve to memory subsets, we would have over-and-under estimated the intrinsic and extrinsic aging components to be approximately 84% and 16%, respectively (**Fig.1b**). By assembling a large collection of brain DNAm datasets (**Methods, SI table S1**) and adjusting for 7 brain cell-type fractions (GABAergic excitatory neurons, glutamatergic inhibitory neurons, astrocytes, microglia, oligodendrocytes, endothelial cells and stromal cells), as estimated with the HiBED algorithm [42], we verified that the intrinsic component of aging in brain is around 68%, similar to the estimate obtained in blood (**Fig.1c, SI**

fig.S3). To see which brain cell-types are driving the extrinsic aging component, we performed a fixed and random effects meta-analysis over 13 Illumina DNAm brain datasets, encompassing over 2,000 samples (**SI table S1**), which revealed an age-associated increase of excitatory neurons, whilst inhibitory neurons and astrocytes decreased (**Fig.2**). Thus, this age-related shift in neuronal subtypes explains why the extrinsic component of aging in brain is also quite substantial (around 32%). Of note, these findings are broadly consistent with recent single-nucleus DNAm studies reporting reduced inhibitory neuron fractions in the brain of aged donors [43,44]. Overall, these results obtained in blood and brain tissue highlight the critical importance of adjusting for CTH when building epigenetic clocks.

A computational strategy to build cell-type specific epigenetic clocks

In order to dissect the intrinsic aging component, we next devised a computational strategy to build cell-type specific epigenetic clocks (**Fig.3a**). After estimation of the cell-type fractions in a given bulk-tissue DNAm dataset, we apply the CellDMC algorithm [41] to a training subset of the data to identify age-associated differentially methylated cytosines in each cell-type of interest (age-DMCTs) (**Fig.3a**). Assuming there is a substantial number of age-DMCTs for a given cell-type, we then apply an Elastic Net Regression model [45] to train predictive models for chronological age parameterized by a penalty parameter (**Fig.3b**). This training is restricted to age-DMCTs of the given cell-type and can be performed either on the residuals of the DNAm data matrix, obtained by regressing out the effect of cell-type composition, or on the unadjusted DNAm data matrix (**Methods**). In the former case, we define the clocks as being "intrinsic", otherwise they are "semi-intrinsic". The best performing models in a blinded model-selection set define our cell-type specific intrinsic and semi- intrinsic epigenetic clocks, one for each cell-type (**Fig.3b**). These clocks are then validated in independent DNAm datasets.

The reasoning behind defining separate intrinsic and semi-intrinsic clocks can be understood with simulation models. We generated simulated DNAm mixtures of 3 underlying cell-types (granulocytes, monocytes and lymphocytes), and in which age-DMCTs are introduced into only one cell-type (lymphocytes) under four different scenarios (**Methods**): (i) age-DMCTs do not map to cell-type specific CpGs and no cell-type fractions change with age, (ii) as (i) but with the lymphocyte fraction changing with age, (iii) age-DMCTs discriminate lymphocytes from the other two cell-types and no cell-type fraction changes with age, (iv) as (iii) but with the lymphocyte fraction changing with age. For each scenario, we generated training and validation sets, deriving lymphocyte specific intrinsic and semi-intrinsic clocks, subsequently validating them in the independent test sets. Interestingly, whilst the intrinsic clock displayed marginally better performance in the scenarios where

cell-type fractions did not change with age, the semi-intrinsic clock was substantially better in cases where the lymphocyte fraction did change with age (**Fig.1c**). Thus, the use of residuals in scenarios where a given cell-type fraction changes with age can dilute out the biological signal in that cell-type, and under these circumstances, the cell-type specific semi-intrinsic clock would be preferable.

Construction and validation of neuron-specific epigenetic clocks

Next, we applied this strategy to the case of brain tissue. For training the clocks, we used one of the largest collections of prefrontal cortex (PFC) samples profiled with Illumina 450k technology [46], encompassing 416 samples with a wide age range (18-97 years) (**SI table S1, Fig.4a**). Although the HiBED algorithm [42] inferred fractions for 7 brain cell-types (**SI fig.S4a**), for subsequent CellDMC analysis, power considerations [41,47] required us to collapse the inferred fractions to 3 broad cell-types: neurons, glia and endothelial/stromal (**Fig.4a**). Applying CellDMC with the estimated fractions, disease status and sex as covariates, resulted in 16,814 neuron-specific and 11,157 glia-specific age-DMCTs (**Fig.4a, Methods**). Visualization of representative age-DMCTs confirmed the cell-type specific nature of their age-associated DNAm changes (**Fig.4b**). Focusing on the neuron-DMCTs and adjusting the DNAm data for variations in cell-type fractions, we next applied the Elastic Net regression framework to learn different models specified by a penalty parameter. For model (i.e. optimal parameter) selection, we used a completely independent Illumina DNAm dataset from Philstrom et al [48] which included 67 PFC control samples with an age range 49-100 years (**Fig.4c, SI table S1**). We declared the model achieving best generalization performance in the Philstrom dataset, as defining our neuron-specific intrinsic ("Neu-In") clock (**Fig.4c, SI fig.S4b, SI table S2**). This neuron-specific clock was further validated in 5 independent datasets of sorted neuron samples, encompassing a total of 143 samples (**Fig.4d, SI fig.S4c, SI table S1**). An analogous neuron "semi-intrinsic" clock ("Neu-Sin") was obtained by applying the same training procedure to the same age-DMCTs in the Jaffe et al DNAm data but now without adjusting the DNAm data for variations in cell-type fractions (**Methods, Fig.4d, SI fig.S4d, SI table S3**). To assess specificity of the clocks to neurons and brain tissue, we assembled a large collection of Illumina DNAm datasets encompassing other brain cell-types (glia), other sorted non-brain cell-types and bulk tissue-types (**SI table S1, Methods & Supplementary Methods**). As a further control, we also benchmarked our Neu-In/Sin clocks to two other clocks: a clock obtained by running the Elastic Net on all CpGs but adjusting the DNAm data for cell-type fractions (**Methods, Fig.4d, SI fig.S4e, SI table S4**), thus defining a brain, but not neuron-specific, clock ("BrainClock"), and the multi-tissue Horvath clock. We restricted the assessment of all clocks to control samples only, to avoid potential confounding by disease. First, we

observed that the neuron intrinsic and semi-intrinsic clocks also displayed good correlations with chronological age in sorted glia datasets, although the largest of these glia datasets displayed a significantly weaker association, suggesting that our neuron-clocks are indeed neuron-specific (**Fig.4d**). In line with this, a multivariate linear regression of correlation strength against cell-type (neuron vs glia), weighted by cohort size, revealed a significantly higher correlation in neurons compared to glia for the neuron-intrinsic clock but not for the other 3 clocks (**SI fig.S5a**). When assessed between brain and non-brain tissue/cell-types, we observed that the two neuron-specific clocks as well as the brain-specific clock achieved better predictions of chronological age in the brain-tissue datasets compared to non-brain ones, whilst, as expected for a multi-tissue clock, Horvath's clock performed similarly between brain and non-brain tissues (**Fig.4e-f**). Importantly, and despite the relatively small number of sorted neuron datasets, the neuron-specific clocks also displayed stronger correlations with age in the sorted neuron datasets compared to the brain-specific clock (**SI fig.S5b**). Thus, overall, our data support the view that the neuron-specific clocks are indeed optimized to predict chronological age in brain cell subtypes and that they are less confounded by changes in cell-type composition compared to e.g. the brain-specific clock.

Neuron and glia-specific clocks are accelerated in Alzheimer's Disease
Next, we asked if the clocks predict age-acceleration in neurodegeneration. To this end, we collated eleven DNAm-studies of Alzheimer's Disease encompassing two brain regions, frontal lobe (FL) and temporal lobe (TL) (**SI table S5**), amounting to a total of 213 AD cases and 185 controls in FL, and 456 cases and 341 controls in TL. As before, we first performed a meta-analysis of cell-type fractions to see if any of these fractions change in AD (**Methods**). This revealed a significant decrease of excitatory neurons and corresponding increases of oligodendrocyte/OPC and microglia fractions in AD (**SI fig.S6**). Thus, only the excitatory neuron fraction is observed to change with both age and AD (**Fig.2 and SI fig.S6**), although interestingly the increased fraction seen with age is reversed in AD cases. For each of the 4 clocks in each brain region we then performed a fixed and random effects meta-analysis (**Methods**), revealing a weak but statistically significant age-acceleration in AD cases for the neuron-intrinsic clock in the TL, but not for brain-specific and Horvath's clocks (**Fig.5a-b**). The association of the neuron-intrinsic clock in the PFC/FL was marginally significant. Given that the I^2 test for heterogeneity showed great consistency in the effect sizes between cohorts and regions (**Fig.5a-b**), we extended the meta-analysis to include all datasets from both regions, now revealing significant associations for both intrinsic and semi-intrinsic neuron-clocks (**Fig.5c**). Because all of the AD datasets considered in this analysis are

from bulk-tissue, we wondered if a similar finding would apply to the other major cell-type in brain, i.e. glia. To this end, we used the previous Jaffe & Philstrom DNAm datasets to build and validate analogous glia intrinsic (Glia-In) and semi-intrinsic (Glia-Sin) clocks (**Methods**, **SI fig.S7, SI tables S6-S7**), subsequently computing their DNAm-Ages in the AD datasets. Interestingly, this also revealed age-acceleration effects in the glia compartment of both FL and TL regions, with the effect being strongest for the Glia-Sin clock (**Fig.5a-c**). Comparing the predicted age-acceleration of the Neu-In with the Glia-In clock, and that of the Neu-Sin with the Glia-Sin clocks, revealed moderate but significant correlations between the Neu and Glia clock estimates (**SI fig.S8**). Thus, age-acceleration patterns in neurons and glia are similar, despite the two clocks being made up of entirely distinct sets of CpGs. In summary, these data indicate that both neuron and glia specific clocks display age-acceleration in AD, and that this acceleration is more prominent in the TL glia fraction.

Construction and validation of a hepatocyte-specific epigenetic clock

As a second example, we considered the case of liver tissue, aiming to build a hepatocyte-specific epigenetic clock. We used our EpiSCORE algorithm[49] and liver DNAm reference matrix defined over 5 cell-types (hepatocytes, cholangiocytes, Kupffer cells, endothelial cells and lymphocytes) [40] to estimate corresponding cell-type fractions in a series of 210 normal liver specimens (age range=18-75, mean age=47, sd=12) with available Illumina EPIC DNAm profiles [50] (**Methods**). Liver cell-type fractions changed with chronological age, although this was mainly restricted to cholangiocytes and endothelial cells, a pattern that was replicated in independent cohorts (**SI fig.S9**). Next, we assessed if we can reliably identify age-DMCTs. To this end we applied CellDMC[41] to all 210 samples, revealing that most age-DMCTs were in either hepatocytes or cholangiocytes, with hepatocyte and cholangiocyte age-DMCTs displaying strong overlap but with opposite directionality of change (**SI fig.S10a-d**). To validate these age-DMCTs, we computed the average DNAm over hepatocyte age-hypermethylated DMCTs in an independent set of 55 primary hepatocyte cultures from donors of different ages (**Methods**), which revealed the expected positive correlation with age (**SI fig.S10e**). Correspondingly, hepatocyte age-hypomethylated DMCTs displayed the expected anti-correlation, with the pattern exactly inverted for cholangiocyte age-DMTCs (**SI fig.S10e**). To further validate the hepatocyte and cholangiocyte age-DMCTs, we collected three additional liver DNAm datasets (**Methods**) and repeated the CellDMC analysis in each one. This confirmed the strong overlap between hepatocyte and cholangiocyte age-DMCTs in each validation cohort (**SI fig.S10f, SI fig.S11**). This strong validation supports the view that we can confidently identify age-associated DMCTs in hepatocytes and that these

display opposite directionality of change in cholangiocytes.

Having demonstrated that we can successfully identify hepatocyte age-DMCTs, we next applied a 10-fold cross-validation procedure to build a hepatocyte-specific semi-intrinsic clock (**Methods**). Of note, this procedure involved re-applying the CellDMC algorithm on the full set of CpGs but now using only on 9 folds, with the leave-one-out bag being used for model selection (**Methods**). This procedure was iterated 10 times, each time using a different bag as the leave-one-out test set, a procedure that avoids overfitting [51], which resulted in a final age-predictor ("HepClock") (**SI table S8, Fig.6a**). We also trained an analogous clock but now using age-DMCs derived from a linear model that only adjusts for cell-type fractions, the resulting clock thus being liver-specific but not hepatocyte-specific ("LiverClock", **Methods, SI table S9**). To validate and benchmark these clocks, we applied them to the DNAm dataset of primary hepatocyte cultures (**Methods**). Although R-values (Pearson Correlation Coefficients) were similar for all clocks, HepClock attained a significantly lower median absolute error of 4.9 years compared to values of 20 and 10 for the liver and Horvath clocks, respectively (**Fig.6b**). This suggests that improved predictions of chronological in sorted hepatocytes is indeed possible with a corresponding hepatocyte specific clock. Next, we applied HepClock and LiverClock to a collection of DNAm datasets from bulk-liver tissue, sorted immune-cell types, as well as other tissue-types including blood, skin, prostate, colon and kidney (**Methods**). This confirmed the clear specificity of HepClock and LiverClock to liver tissue (**Fig.6c**). We note that Horvath's clock was excluded from this analysis, because some of the liver DNAm datasets were used for the original training of the Horvath clock itself [24].

Hepatocyte specific clock predicts biological age-acceleration in liver pathologies

We next applied HepClock, LiverClock and Horvath's clock to 3 separate DNAm studies representing a variety of liver pathologies, in order to assess the ability of these clocks to predict biological age-acceleration. One study encompassed 115 liver samples from patients with non-alcoholic fatty liver disease (NAFLD) [50], another set comprised 67 liver tissue samples from obese individuals (BMI>35) [52], with the 3rd cohort comprising 94 liver tissue samples from patients with Alpha-1 Antitrypsin Deficiency (AATD) [53], a genetically inherited disorder that increases the risk of cirrhosis (**Methods**). Whilst on the NAFLD samples all three clocks were accelerated, the observed acceleration was strongest for the Liver and HepClocks (**Fig.6d**). On the clinically obese samples only HepClock displayed a very significant age-acceleration effect, with Horvath's clock being marginally associated (**Fig.6d**). In the AATD-cohort, only the Liver and HepClocks displayed age-acceleration (**Fig.6d**). In

another independent liver-tissue DNAm dataset encompassing healthy non-obese and obese individuals [54,55], HepClock and Horvath clock age-acceleration measures displayed a correlation with BMI, but not so for the LiverClock (**Fig.6e)**, mirroring the results in the previous cohort of 67 obese individuals (**Fig.6d**). The age-acceleration effect in the hepatocytes of NAFLD and obese livers is noteworthy, because the hepatocyte fraction itself decreases substantially in NAFLD and obese individuals (**SI fig.S12**). Overall, as assessed over the 4 studies, the hepatocyte-specific epigenetic clock improves the sensitivity to detect biological age acceleration in the hepatocytes of liver tissue in variety of different liver pathologies including an inherited one where Horvath's clock did not predict age-acceleration.

Minimal overlap with blood-based clocks

Having demonstrated how cell-type specific clocks predict age-acceleration in relevant conditions and cell-types, it is important to study the nature of the specific CpGs making up these clocks. We first assessed overlaps of our clock CpGs with previous clocks including Horvath [24], Zhang [11] and PhenoAge [3]. In general, overlaps were minimal if not zero (**SI fig.S13**). Likewise, the overlap with the recently proposed causal DamAge and AdaptAge clocks [56] was minimal (**SI fig.S13**). These small overlaps are unsurprising since all of these other clocks have been derived from blood-tissue (or predominantly so). This may indicate that there is an inherent limitation in using blood to build epigenetic clocks that are informative of biological age in non-blood tissue-types, echoing previous views [32].

The Neu-In clock CpGs are enriched for AD EWAS hits and mQTLs

Given the associations of our clocks with AD and BMI, we next asked if our clock CpGs overlap with AD or BMI associated differentially methylated cytosines (DMCs), as determined by corresponding EWAS [57,58] (**Methods**). Whilst for the brain-clocks we did observe a significant overlap with AD-DMCs, this was not evident for the Hep/LiverClocks and BMI-DMCs (**SI fig.S14**). We also asked if any of our clock CpGs are associated with mQTLs using the ARIES database of mQTLs in blood [59] (**Methods**). This revealed a significant enrichment for mQTLs among brain clock CpGs, whilst it was only marginal in the case of the Hep/LiverClocks (**SI fig.S14**), probably because of the smaller number of Hep/LiverClock CpGs reflecting the smaller training set in the case of liver tissue. In the case of the Neu-In clock, we observed 2 clock-CpGs defining cis-mQTLs where the linked SNP has been associated with AD [60] (**SI fig.S14, SI table S10**). Interestingly, these cis-mQTLs map to the gene *CASS4,* a CAS scaffolding protein family member involved in the

formation of neuritic plaques and neurofibrillary tangles, as well as in the disruption of synaptic connections in AD [61,62]. Thus, a minor component of the Neu-In clock involves CpGs whose DNAm-levels are also influenced by genetic variants associated with AD.

Discussion

Here we have quantified, within the context of epigenetic clocks, the relative contributions of extrinsic and intrinsic aging, demonstrating in two different tissue-types (blood and brain), that the extrinsic component accounts for around 35% of an epigenetic clock's accuracy. Thus, whilst most of an epigenetic clock's accuracy is intrinsic, driven by age-associated DNAm changes in individual cell-types, the age-related shifts in tissue composition contributes a substantial amount. For instance, based on our analysis, epigenetic clocks like Horvath or Zhang clocks [11,24], which typically attain an R^2 value of 0.8 or higher, would likely only display R^2 values of approximately 0.5 if one were to rederive the clocks adjusting for 12 immune-cell type fractions. Whilst in the case of blood, the extrinsic component is driven by an age-related increase of the mature to naïve T-cell fractions, in brain, this is driven by an age-related shift in the excitatory to inhibitory fractions.

Our second key contribution is the dissection of the intrinsic component of aging through construction of cell-type specific epigenetic clocks. As shown here in the case of brain and liver tissue, the resulting cell-type specific clocks demonstrate clear cell-type and tissue-specificity, achieving improved predictions of chronological age in the corresponding sorted cells and tissue-types. Most importantly, our cell-type specific clocks yield improved estimates of biological age in the relevant cell-types and tissues (here brain and liver), as compared to Horvath's multi-tissue clock and the non-cell-type specific (but tissue-specific) brain and liver clocks, the latter result highlighting the importance of distinguishing cell-type specific from tissue-specific clocks.

The case of Alzheimer's Disease is important to discuss further, since there has been substantial controversy and debate as to whether epigenetic clocks display age-acceleration in AD or not. Here we have significantly advanced the field by constructing neuron and glia-specific clocks, and demonstrating an age-acceleration effect for both neurons and glia, with the effect being strongest for the Glia-Sin clock in the temporal lobe, suggesting that epigenetic age-acceleration associated with AD may be particularly prominent in the glia TL compartment. As far as the age-acceleration of the Neu-In and Neu-Sin clocks is concerned, it is worth noting

that the excitatory neuron fraction was observed to increase with age, whilst it was decreased in AD. This pattern would suggest that the observed age-acceleration of the neuron-compartment in AD is not due to shifts in underlying neuronal subtype proportions, but instead caused by DNAm changes in these specific neuronal subtypes. It is also noteworthy that one of the CpGs in the Neu-In clock defines a cis-mQTL linked with a SNP that has been associated with AD and which maps to *CASS4*, a gene which has been previously implicated in neuritic plaque formation and loss of synaptic connectivity in AD [61,62]. This is important, because it is conceivable that clock-CpGs whose DNAm levels track age-acceleration in AD due to environmental risk factors (e.g. social isolation, sleep deprivation) may also be under the influence of genetic variants that determine AD-risk, which would thus point towards a shared causal pathway to AD [17]. The distinctive lack of overlap between our clock-CpGs and those of the recently proposed causal DamAge and AdaptAge clocks [56] likely owes to the fact that these causal-clocks have been derived from blood, a tissue which may not be relevant for studying causal epigenetic changes in brain-tissue and AD.

The ability to quantify biological age at cell-type resolution can also help address whether age-acceleration in various cell-types of a tissue correlate with each other or not. Here we explored this important question in the context of the neuron and glia clocks, revealing that the age-acceleration of these two clocks are indeed significantly correlated. Despite this, the age-acceleration measures were also sufficiently distinct, with the glia clocks displaying stronger acceleration in AD compared to the neuron-clocks. In the case of liver, it should be noted that most of the identified age-associated DNAm changes occurred in hepatocytes and cholangiocytes, with the underlying DNAm changes in these two cell-types displaying a very strong anti-correlation. Thus, a putative cholangiocyte clock would be composed of overwhelmingly the exact same CpGs making up the hepatocyte clock, and thus be effectively analogous to it, which is why we did not consider it here.

Whilst the computational framework presented here could be extended to construct cell-type specific epigenetic clocks for other tissue and cell-types, it is worth highlighting some of the key limitations. First, it must be stressed that our strategy hinges critically on the ability to detect age-associated DNAm changes in individual cell-types and that this requires appropriately powered datasets [23,47]. Indeed, the need for large sample sizes (typically at least many hundreds) has recently been emphasized by us [47]. The required sample size to ensure adequate power is a complex function of many parameters, including the number of cell-types (i.e. cellular resolution) at which the inference is to be made and the distribution (notably means and variances) of their proportions across samples. At higher cellular resolution, there is a corresponding need that the cohort be big enough so that all possible combinations of cell-type fractions are realized, which in practice may be difficult

because of intrinsic correlations between cell-types. Thus, the need for fairly large DNAm datasets currently limits our strategy to a few tissue-types (e.g. brain/liver) and a low cellular resolution (e.g. neuron, glia and stroma). In order to be able to quantify biological age in additional subsets, say glia subtypes in brain, or for the different immune cell subtypes in liver, one would need on the order of thousands of samples, which in future could be realized by merging DNAm datasets together. For instance, for an age-related disease like NAFLD that is associated with molecular alterations in hepatocytes and which leads to a marked reduction in the hepatocyte fraction with age, other cell-types (e.g. Kupffer macrophages) may also undergo biological aging, marking the persistent inflammation associated with NAFLD [63], but here we were underpowered to detect age-DMCTs in macrophages. Another key unresolved issue is the relative quantification of cell-type specific versus shared/common age-associated DNAm changes. For instance, as shown by us, most age-associated DNAm changes in blood seem to be shared between distinct immune cell-types [64], implying that the number of truly cell-type specific age-associated DNAm changes may be small in a tissue like blood. Ideally, cell-type specific clocks should be built from those CpGs that only change in one cell-type of the tissue and not others, yet confidently identifying these is challenging and would require the aforementioned very large DNAm datasets. Certainly, an in-depth study in a tissue like blood, where large numbers of DNAm datasets are available would merit further investigation, although given that the relative fractions of naïve and mature T-cell subtypes change with age [21], this may require on the order of tens of thousands of samples.

In summary, we have quantified the relative contributions of extrinsic and intrinsic aging to an epigenetic clock's predictive accuracy, whilst also presenting a computational strategy to dissect the intrinsic component, building cell-type specific epigenetic clocks that clearly demonstrate cell-type and tissue specificity in relation to prediction of chronological and biological age. The construction of such cell-type specific clocks will be critical to advance our understanding of biological aging at cell-type resolution and for evaluating in which cell-types cellular rejuvenation interventions work best.

Methods

Brief description of main datasets analyzed
Human liver DNA methylation datasets:

Johnson et al [50]: This is an Illumina EPIC DNAm dataset of liver-tissue, comprising 341 samples, with 16 duplicates among them. The dataset includes samples from four stages representing the progression of non-alcoholic fatty liver disease (NAFLD): stage 0, stage 3, stage 3-4, and stage 4. Stage 0 is normal liver. Data is available at GEO (http://www.ncbi.nlm.nih.gov/geo/; accession number GSE180474). We downloade a preprocessed data matrix following QC, including P-values of detection for each probe. To select high-quality probes for clock construction, we utilized the dropLociWithSnps() function of minfi v.1.44.0 to remove SNP probes [65]. Sex chromosome CpGs were also excluded to minimize gender disparities. Any beta value with a P-value greater than 0.01 was considered as a missing value, and any probe with a coverage-fraction less than 0.99 was discarded. Missing values were then imputed using impute.knn (k=5) of impute v.1.72.3 [66]. Type-2 bias was corrected using BMIQ algorithm [67], and finally duplicates were averaged, resulting in a final data matrix of 325 samples encompassing 812,168 CpGs on autosomes. Of these 325 samples, 210 were from normal healthy liver, whereas the remaining 115 represent NAFLD of grade-3 or higher. In this work, the 210 normal liver samples were used to build and validate the Liver and HepClocks in a 10-fold CV-strategy, whereas the 115 NAFLD samples were used to test for biological age-acceleration.

Park et al [68]: Illumina HM850 (EPIC) dataset encompassing 56 primary hepatocyte cultures. Raw IDAT files were downloaded from GSE123995. Data was processed as before. Any probe with missing values was excluded, resulting in 809,289 probes. A sample was removed due to very low predicted hepatocyte fraction (as estimated using EpiSCORE).

Bonder et al [52]: Illumina HM450k dataset profiling 67 obese liver samples. All individuals have a BMI greater than 35. Raw IDAT files were downloaded from GSE61446. minfi::preprocessRaw() was employed to process the raw IDAT data. The remaining steps in the preprocessing pipeline align with the aforementioned description, with a probe coverage threshold set at 0.90. A total of 411,233 probes remained.

Bacalini et al [69]: Illumina HM450 dataset profiling 40 normal liver samples. Raw IDAT files were downloaded from GSE107038 and processed as the other datasets above. Any probe with missing values was excluded, resulting in 446,651 probes. A sample was removed due to very low predicted hepatocyte fraction (as estimated using EpiSCORE).

Horvath and Ahrens et al [54,55]: Illumina HM450 datasets comprising 77 samples from healthy non-obese (n=44) and healthy obese (n=33) individuals. Processed data was downloaded from GEO under accession numbers GSE61258 and GSE48325. For both datasets, we only used probes with no missing values (i.e. significant detection P-value < 0.01) across all samples. Datasets were then normalized with BMIQ. The

merged dataset encompassed 77 samples and 417,213 probes. A subset of samples were obtained following bariatric surgery, but since surgery had no effect on clock estimates [54,55] we kept these samples in. Of note, the subset of 44 healthy non-obese samples were used to assess the predictive accuracy of the clocks with chronological age.

Wang et al AATD dataset [53]: Illumina HM850 (EPIC) dataset encompassing 108 liver samples from 94 Alpha-1 Antitrypsin Deficiency (AATD) patients and 14 AATD patients with cirrhosis. Only the normal liver samples from the 94 AATD patients without cirrhosis were used. Processed normalized data was downloaded from GEO under accession number GSE119100. This normalized data is already adjusted for type-2 probe bias (adjusted with Subset-quantile Within Array Normalization (SWAN)), and probes that failed to achieve significant detection (p<0.01) in all samples alongside cross-reactive probes were filtered out, leaving a total of 866,836 probes.

Human brain DNA methylation datasets used for training and model selection
Jaffe: The 450k dataset from Jaffe et al [46], which profiled DNAm in human dorsolateral prefrontal cortex from controls and patients with schizophrenia (SZ), was obtained from the NCBI GEO website under the accession number GSE74193 (https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE74193). The file "GSE169156_RAW.tar" which contains the IDAT files was downloaded and processed with minfi package [65]. Probes with >25% NAs (defined by $P > 0.01$) were discarded. Samples younger than 18 years old were also discarded. The filtered data was subsequently normalized with BMIQ [67], resulting in a normalized data matrix for 473,536 probes and 416 samples, of which 226 are controls and 190 are SZ cases, age range 18-97 years.

Pihlstrøm: The EPIC dataset from Pihlstrøm et al [48], which profiled DNAm in human frontal cortex, was obtained from the NCBI GEO website under the accession number GSE203332 (https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE203332). The file "GSE203332_RAW.tar" which contains the IDAT files was downloaded and processed with minfi package. Probes with NAs (defined by $P > 0.01$) were discarded. The filtered data was subsequently normalized with BMIQ, resulting in a normalized data matrix for 832,899 probes and 442 samples, of which 67 are controls with an age-range of 49 to 99 years.

Human brain DNA methylation datasets used for validation
Due to the large number of brain DNAm datasets analyzed, these are all summarized in **Supplementary Methods**.

Human DNA methylation datasets from other tissues

All the whole blood datasets used here have been preprocessed and analyzed previously by us [21]. In addition to whole blood, we also analyzed the following DNAm datasets from other tissue-types:

Colon [70]: Illumina HM450 dataset comprising 96 normal colon mucosa. Raw data is available from GSE131013. IDAT files were downloaded and processed with minfi package. Probes with NAs (defined by $P > 0.01$) were discarded. The filtered data was subsequently normalized with BMIQ, resulting in a normalized data matrix for 448,286 probes and 240 samples. 144 tumor samples were discarded.

Prostate [71]: Illumina HM450 dataset comprising 51 normal prostate samples. Raw data is available from GSE76938. IDAT files were downloaded and processed with minfi package. Probes with NAs (defined by $P > 0.01$) were discarded. The filtered data was subsequently normalized with BMIQ, resulting in a normalized data matrix for 471,112 probes and 123 samples and 72 prostate cancer samples were discarded.

Kidney [72]: Illumina HM450 dataset profiling 31 normal kidney tissue samples. Raw data is available from GSE79100. There are three technical replicates per tissue. IDAT files were downloaded and processed with minfi package. Probes with NAs (defined by $P > 0.01$) were discarded. The filtered data was subsequently normalized with BMIQ, resulting in a normalized data matrix for 454,671 probes and 93 samples.

Skin [73]: Illumina HM450 dataset profiling 322 normal skin samples from a peri-umbilical punch biopsy. From GSE90124, we downloaded the processed type-2 probe bias adjusted DNAm data with 450,531 probes and 322 samples.

Estimating cell-type fractions in whole blood, liver and brain.

For whole blood, we estimated fractions for 12 immune cell subtypes (naïve and memory B-cells, naïve and memory CD4+ T-cells, naïve and memory CD8+ T-cells, T-regulatory cells, natural killer cells, monocytes, neutrophils, eosinophils, basophils) using EpiDISH and our validated DNAm reference matrix for these 12 cell-types [21]. In the case of liver tissue we used our validated EpiSCORE algorithm and the associated liver DNAm reference matrix defined for hepatocytes, cholangiocytes, endothelial cells, Kupffer macrophages and lymphocytes [49]. In the case of brain-tissue, we applied the HiBED algorithm [42], since the EpiSCORE brain DNAm reference matrix led to the microglia fraction always being zero. We applied HiBED to estimate proportions for 7 brain cell-types including GABAergic excitatory neurons, glutamatergic inhibitory neurons, astrocytes, microglia, oligodendrocytes, endothelial cells and stromal cells. When applying CellDMC we collapsed these estimates to 3 broad cell-types: neurons, glia and endothelial/stromal.

Relative quantification of extrinsic and intrinsic aging in whole blood and brain.

Whole Blood: As training data we used two of the largest whole blood DNAm datasets, the Airway [74] and Lehne [75] cohorts, which we have previously normalized and analyzed [21]. Briefly, the Airway EPIC DNAm dataset (GSE147740) included 1032 samples with age information. The Lehne dataset is a 450k DNAm dataset (GSE55763) encompassing 2707 samples in total. The two datasets were merged, resulting in 314,436 CpGs and 3739 samples. In the Airway dataset, age is cofounded by gender (Pearson correlation test, P-value =0.003896). To train the unadjusted and adjusted clocks, we applied the Elastic Net regression to the training dataset with alpha=0.5 and a 10-fold cross-validation to infer the optimal penalty parameter lambda. In the unadjusted case, we applied the Elastic Net to the standardized DNAm residuals after adjusting for sex and cohort only. In the adjusted case, we applied the Elastic Net to the standardized DNAm residuals after adjusting for sex, cohort and the 12 estimated immune-cell type fractions. We also trained another clock adjusted for 9 immune cell-type fractions by merging the naïve and memory lymphocyte subsets together, in other words, by adding the fractions of naïve and memory B-cells to yield a fraction for B-cells, and similarly for CD4T and CD8T-cells. The 3 optimal clocks were then applied to 11 independent whole blood datasets (using only control/healthy samples, see **SI table.S1**), all previously normalized as in Luo et al [21]. In the case of the adjusted clocks, the clocks were applied to the standardized DNAm residuals after adjustment for the 12 or 9 immune-cell fractions. Finally, the predicted ages were linearly regressed to the true chronological ages in each cohort, resulting in R^2 values for the unadjusted and the two adjusted clocks. The ratio of the adjusted clock's R^2 value to that of the unadjusted clock, R^2_{ADJ}/R^2_{UNADJ} quantifies the fraction of the unadjusted clock's accuracy that can be attributed to intrinsic aging. For the extrinsic component, its fractional contribution is given by $1 - R^2_{ADJ}/R^2_{UNADJ}$.

Brain: For training, we used the Jaffe et al dataset [46] encompassing 491 samples, including 300 controls and 191 samples from schizophrenia patients, with an age range 0.003 ~ 96.98 years. We used HiBED to estimate the fractions for 7-cell-types (ExN, InN, Oligo/OPC, Micro, Astro, Endo, Strom), as described earlier. To train the Elastic Net clocks with alpha=0.5, we used a 10-fold cross-validation to infer the optimal penalty parameter lambda. In the unadjusted case, we applied the Elastic Net to the standardized DNAm residuals obtained after adjusting for sex and disease-status, whilst in the adjusted case, it was applied on the standardized DNAm residuals obtained after adjusting for sex, disease-status and 7 brain cell-type fractions. The two clocks were then applied to 11 independent brain DNAm datasets (using control/healthy samples only, see **SI table.S1**), obtaining in each cohort

corresponding R^2_{ADJ} & R^2_{UNADJ} values. The fraction of the unadjusted clock's accuracy attributed to intrinsic and extrinsic aging was then estimated as described above for whole blood.

Construction of the neuron/glia/brain-specific epigenetic clocks

We used HiBED [42] to estimate fractions for 7 brain cell-types in the Jaffe and Philstrom datasets, in each case summing up fractions to yield estimates for 3 broad cell-types: neurons (Neu), glia (Glia) and endothelial/stromal. We then applied CellDMC [41], as implemented in the EpiDISH package [76], to the Jaffe dataset in order to infer age-associated DMCTs in each of these 3 broad cell-types. Briefly, the CellDMC model postulates the following relationship for the DNAm level β_{cs} of a CpG c in sample s

$$\beta_{cs} = \sum_{t=1}^{3} \hat{f}_{ts} \beta_{cts} = \sum_{t=1}^{3} \hat{f}_{ts} (\mu_{ct} + Age_s \gamma_{ct}) + \varepsilon_{cs}$$

where $\hat{f}_{ts}$ are the estimated cell-type fractions for cell-type t in sample s. This can be re-expressed as a multivariate regression model with interaction terms between age and cell-type fractions

$$\beta_{cs} = \sum_{t=1}^{3} \mu_{ct} \hat{f}_{ts} + \sum_{t=1}^{3} \gamma_{ct} \hat{f}_{ts} Age_s + \varepsilon_{cs}$$

This estimates the regression coefficients μ_{ct} and γ_{ct} so as to minimize the error in a least squares sense. When applying CellDMC to the Jaffe dataset, we also adjusted for schizophrenia (SZ) status and gender, i.e we ran the model:

$$\beta_{cs} = \sum_{t=1}^{3} \mu_{ct} \hat{f}_{ts} + \sum_{t=1}^{3} \gamma_{ct} \hat{f}_{ts} Age_s + \theta_c Sex_s + \varphi_c SZ_s + \varepsilon_{cs}$$

We included both SZ cases and controls, adjusting for case/control status, so as to increase power. Having identified the significant neuron-DMCTs (those CpGs c with a significant γ_{ct} with $t=neuron$, t-test FDR<0.05), we next trained Elastic Net regression (glmnet R package) [45] models (alpha=0.5) for age, each parameterized by a different penalty parameter (lambda) value. This training was done on the residuals obtained by regressing out the cell-type fractions from the DNAm beta matrix. To select the optimal model, we then applied the age-predictors for each choice of lambda to the Philstrom DNAm dataset (using the residuals after adjusting this DNAm data matrix for cell-type fractions), selecting the lambda that minimized the root-mean-square error (RMSE). This defines a neuron-specific epigenetic clock that is applicable to residuals after adjusting for cell-type fraction, a clock we call "Neu-In"

(neuron-intrinsic).

We also constructed a similar clock called "Neu-Sin" (neuron semi-intrinsic), following the same procedure as before, but now not adjusting the DNAm data matrix for cell-type fractions. In other words, whilst the Neu-In and Neu-Sin clocks are both built from CpGs that change with age in neurons, the Neu-In clock uses the residuals for construction, whilst the Neu-Sin clock does not. A third non-cell-type specific clock (Brain) was also built using the same procedure as before, but not using CellDMC to identify CpGs changing with age, instead identifying age-DMCs by running a linear model of DNAm versus age adjusting for cell-type fractions, i.e. in this case the regression model is

$$\beta_{cs} = \sum_{t=1}^{3} \mu_{ct} \hat{f}_{ts} + \gamma_c\, Age_s + \theta_c Sex_s + \varphi_c SZ_s + \varepsilon_{cs}$$

and one finds age-DMCs, as those CpGs c with a significant γ_c (t-test, FDR<0.05). The age-predictor itself was then also trained on residuals obtained after regressing out the effect of cell-type fractions. Of note, the model selection step to define the Neu-Sin and Brain clocks was also done using the independent Philstrom DNAm dataset.

Finally, an analogous procedure describing the construction of the Neu-In and Neu-Sin clocks was followed to build corresponding glia-specific intrinsic and semi-intrinsic clocks (Glia-In & Glia-Sin).

Validation of the neuron/glia/brain specific clocks

To validate and demonstrate specificity of the Neu-In/Sin, Glia-In/Sin and Brain clocks, we applied them to sorted neuron, sorted glia, brain and non-brain DNAm cohorts. When applying the Neu-In, Glia-In and Brain clocks, we adjusted the DNAm datasets for underlying cell-type fractions, which were estimated using HiBED in the case of brain tissue, EpiDISH in whole blood cohorts and EpiSCORE in other non-brain tissues.

Construction of HepClock and LiverClock

We first describe the construction of the hepatocyte-specific clock (HepClock). We selected 210 normal liver samples from Johnson et al.'s dataset for training and model selection. EpiSCORE was used to estimate fractions for 5 liver cell-types including hepatocytes, as described above. The 210 samples were then randomly divided into 10 bags (each bag with 21 samples), with stratified sampling within each age group to ensure a similar age distribution in each bag. To identify CpGs changing with age in hepatocytes, we then applied the CellDMC algorithm [41] to a subset composed of 9

bags, resulting in a set of hepatocyte age-DMCTs. The specific CellDMC model run was

$$\beta_{cs} = \sum_{t=1}^{5} \mu_{ct}\hat{f}_{ts} + \sum_{t=1}^{5} \gamma_{ct}\,\hat{f}_{ts}Age_s + \varepsilon_{cs}$$

Having identified the significant hepatocyte-DMCTs (those CpGs c with a significant γ_{ct} with t=*hepatocyte,* t-test FDR<0.05), and using the same 9 bags, we then trained a range of lasso regression models parameterized by a penalty parameter lambda. Each of these models (parameterized by lambda) were subsequently applied to the left-out-bag to yield age-predictions for the samples in this left-out-bag. This procedure was repeated 10-times, each time using a different subset of 9 bags for training and one left-out-bag for model selection. The age-predictions over all 10 folds were then combined and compared to the true ages computing the RMSE for each choice of lambda. The model minimizing the RMSE was then selected as the optimal model. Of note, the above procedure applies CellDMC a total of 10 times, once for each fold. The number of hepatocyte age-DMCTs is variable between folds, so we identified the minimum number (k) obtained across all folds, and subsequently only used the top-k hepatocyte age-DMCTs from each fold before running the lasso regression model. Lasso was implemented using the glmnet function from the glmnet R-package (version 4.7.1) [45], with alpha set to 1. Having identified the optimal model (lambda parameter), we finally reran the lasso model with this optimal parameter on the full set of 210 samples using the top-k hepatocyte age-DMCTs as inferred using CellDMC on all samples.

The LiverClock was built in an analogous fashion with only one key difference: instead of applying CellDMC in each fold, we ran a linear regression of DNAm against age with all the liver cell-type fractions as covariates. That is, the model in this case is

$$\beta_{cs} = \sum_{t=1}^{5} \mu_{ct}\hat{f}_{ts} + \gamma_c\,Age_s + \varepsilon_{cs}$$

and one finds age-DMCs, as those CpGs c with a significant γ_c (t-test, FDR<0.05). After optimization of lambda, the final age lasso predictor was built from such age-DMCs derived using all 210 samples.

Simulation analysis on mixtures of sorted cell-types

We used a simulation framework to assess and compare the intrinsic and semi-intrinsic paradigms for cell-type specific clock construction. We first generated basis DNAm data matrices for 3 hypothetical cell-types "A", "B" and "C" each one

defined over 10,000 CpGs and 200 samples, drawing DNAm beta-values from Beta(10,40), i.e. from a Beta-distribution with a mean value of 0.25. We then declared the first 1000 CpGs to be cell-type specific for a cell-type "A", by drawing the DNAm-values for two other cell-types ("B" and "C") from Beta(40,10), i.e. from a Beta-distribution with mean 0.75. Thus, the first 1000 CpGs are unmethylated in cell-type "A" compared to "B" and "C". Likewise, another 1000 non-overlapping CpGs were defined to be unmethylated in cell-type "B" and another 1000 non-overlapping CpGs for cell-type "C". The remaining 7000 CpGs are unmethylated in all cell-types. To generate realistic mixtures of these 3 cell-types, we randomly sampled cell-type fractions from true estimated fractions in whole blood by applying our EpiDISH algorithm [76] to estimate 7 immune cell-type fractions in a large dataset of 656 whole blood samples [77]. We declared 3 cell-types to be granulocytes, monocytes and lymphocytes, by adding together the fractions of B-cells, CD4T-cells, CD8 T-cells and NK-cells to represent "lymphocytes" and adding together the fractions of neutrophils and eosinophils to represent "granulocytes". To define the ages of the 200 samples we bootstrapped 200 values from the ages of 139 donors from the BLUEPRINT consortium [78] in order to mimic an age distribution of a real human cohort. Before mixing together the artificially generated DNAm data matrix for the 3 hypothetical cell-types, we introduced age-DMCTs in the cell-type defining lymphocytes. This was done by drawing DNAm-values for these age-DMCTs to increase with the age of the sample. In more detail, because the minimum age of the BLUERPINT samples was 30, we first assigned M-values according to the formula:

$$Mval = Mval0 + (age - 30) * effS + \sigma$$

where *effS=0.1* and *Mval0=log2(0.1/(1-0.1))* and σ representing a Gaussian deviate with a standard deviation value equal to *effS*. These values were chosen to mimic age-related DNAm changes as seen in real data, although with a smaller noise component in order to better understand and interpret downstream results.

In total we considered 4 different scenarios depending on the choice of age-DMCTs in lymphocytes and whether the fraction of "lymphocytes" increases with age or not. In scenarios (i) and (ii), 1000 age-DMCTs were chosen to not overlap with any of the 3000 cell-type specific marker CpGs. In scenarios (iii) and (iv), 900 age-DMCTs in lymphocytes were chosen to be lymphocyte-specific markers. Reason we did not pick all 1000 lymphocyte markers, is to leave 100 markers that remain cell-type specific to allow for reference construction and estimation of cell-type fractions using EpiDISH. In scenarios (i) and (iii), cell-type fractions were chosen so that no fraction changed significantly with age, whilst in scenarios (ii) and (iv) the lymphocyte-fraction was chosen to increase with age.

For each of the 4 scenarios, we generated training and validation mixture sets. DNAm reference matrices were built from different realizations of the basis DNAm data

matrices but using the same parameter values. These reference matrices were then applied to the training dataset to infer cell-type fractions. With the estimated cell-type fractions we then applied CellDMC [41] to the training set to infer age-DMCTs. Since age-DMCTs are only introduced in lymphocytes, we focus on the predicted age-DMCTs in this cell-type, applying an Elastic Net (alpha=0.5) penalized regression framework in a 10-fold cross-validation framework to learn a predictor of chronological age using either an intrinsic or semi-intrinsic paradigm: in the semi-intrinsic paradigm the Elastic Net predictor was trained on the original DNAm-values, whilst in the intrinsic case it was applied to residuals obtained by regressing out the effect of the cell-type fractions. The optimal penalty parameter was selected via the 10-fold CV procedure and the optimal predictor then applied to the validation dataset. Of note, in the semi-intrinsic setting the validation set is defined on the original DNAm-values, whereas the intrinsic model is applied to the residuals of the validation set. For each scenario we ran 50 different simulations, recording the accuracy (root mean square error-RMSE and Pearson Corrrelation Coefficient-PCC) of the semi-intrinsic and intrinsic lymphocyte-specific clock predictors in the corresponding validation set.

Meta-analyses

In the case of brain tissue datasets, we performed two separate meta-analyses. One of these explored if brain cell-type fractions change with age, whilst the other was done to explore if our cell type specific clocks (extrinsic age-acceleration-EAA) display correlations with Alzheimer's Disease (AD) status. The former analysis was done for 7 brain cell types over 13 independent DNAm datasets encompassing frontal and temporal lobe regions, in each dataset adjusting for disease status and sex. The latter meta-analysis was done for frontal lobe (n=5) and temporal lobe (n=6) separately, and then again treating both regions together (n=11). In each cohort, associations between EAA and AD status was assessed using a multivariate linear regression with sex and age as covariates, AD status encoded as ordinal with 0=control, 1=AD case. We extracted the corresponding effect size, standard error, Student's t test, and P-value for each regression, Meta-analysis was then performed using both fixed effect and random effect inverse variance method using the metagen function implemented in the meta R-package [79,80].

Enrichment analyses

We assessed enrichment of clock-CpGs against EWAS-DMCs for AD and BMI, as determined by the EWAS catalog [57] and EWAS-atlas [58]. Enrichment analysis was done using a one-tailed Fisher exact test. We also assessed enrichment of our clock-CpGs for blood mQTLs as determined by the ARIES mQTL database [59]

focusing on the mQTLs identified in the adult population. Significance was evaluated using a one-tailed Fisher exact test. Overlap with GWAS AD SNPs from Bellenguez et al [60] and GWAS BMI SNPs from Pulit et al [81] was done by restricting to SNPs that were part of mQTLs that included clock-CpGs.

Code Availability

All the cell-type specific clocks presented here are freely available as part of the R-package CTSclocks, downloadable from *https://github.com/HGT-UwU/CTSclocks*

Data Availability

The Illumina DNA methylation datasets analyzed here are all freely available from GEO (www.ncbi.nlm.nih.gov/geo). Details: Liver-NAFLD (n=325, GEO: GSE180474); Hepatocyte cultures (n=56, GEO: GSE123995); Liver-obese (n=67, GEO: GSE61446); Liver-Horvath/Ahrens (n=77, GEO: GSE61258 & GSE48325); Liver tissue (40 normal, GEO: GSE107038); Liver-AATD (n=94, GEO: GSE119100); MESA (214 purified CD4+ T-cells and 1202 Monocyte samples, GEO: GSE56046 & GSE56581); Colon (96 normal and 144 adenocarcinoma, GEO: GSE131013); Prostate (123 normal and 63 benign, GEO: GSE76938); Kidney (93 normal, GEO: GSE79100); Skin (322 normal, GEO: GSE90124); Brain sorted-Pai (33 normal and 67 BD, SCZ, GEO: GSE112179); Brain sorted-Gasparoni (32 normal and 30 AD, GEO: GSE66351); Brain sorted-Kozlenkov (29 normal and 59 Heroin or suicide, GEO: GSE98203); Brain sorted-Guintivano (58 normal, GEO: GSE41826); Brain sorted-Hannon (103 normal, GEO: GSE234520); Brain sorted-Witte (20 normal and 36 BD, SCZ, MDD, GEO: GSE191200); Brain-Jaffe (226 normal and 190 SCZ, GEO: GSE74193); Brain-Philstrom (67 normal and 375 Various mental disease and neurodegenerative disease, GEO: GSE203332); Brain-Semick (187 normal and 82 AD, GEO: GSE125895); Brain-Pidsley (71 normal and 179 AD, GEO: GSE43414); Brain-Smith (138 normal and 148 AD, GEO: GSE80970); Brain-Horvath (95 normal and 205 AD, HD, GEO: GSE72778); Brain-Gasparoni (52 normal and 76 AD, GEO: GSE66351); Brain-Watson (34 normal and 34 AD, GEO: GSE76105); Brain-Brokaw (358 normal and 450 AD, GEO: GSE134379); Brain-Murphy (38 normal and 37 MDD, GEO: GSE88890); Brain-Rydbirk (37 normal and 41 MSA, GEO: GSE143157); Brain-Torabi (50 normal and 75 SCZ, GEO: GSE128601); Brain-Xu (25 normal and 23 AUD, GEO: GSE49393); Brain-Huynh (19 normal and 27 Multiple sclerosis, GEO: GSE40360); Brain-Viana2 (13 normal and 14 SCZ, GEO: GSE89703); Brain-Viana3 (17 normal and 16 SCZ, GEO: GSE89705); Brain-Viana4 (28 normal and 21 SCZ, GEO: GSE89706); Brain-Markunas (221 normal, GEO:

GSE147040); Brain-Viana1 (17 normal and 16 SCZ, GEO: GSE89702); Brain-Stefania1 (37 normal and 21 Suicide, GEO: GSE137222); Brain-Stefania2 (15 normal and 18 Suicide, GEO: GSE137223); Whole blood-Airway (1032 normal, GEO: GSE147740); Whole blood-Barturen (101 normal and 473 SARS, GEO: GSE179325); Whole blood-Flanagan (184 normal, GEO: GSE61151); Whole blood-Hannon1 (304 normal and 332 SCZ, GEO: GSE80417); Whole blood-Hannon2 (405 normal and 260 SCZ, GEO: GSE84727); Whole blood-Hannum (656 normal, GEO: GSE40279); Whole blood-Johansson (729 normal, GEO: GSE87571); Whole blood-Lehne (2707 normal, GEO: GSE55763); Whole blood-LiuMS (139 normal and 140 Multiple Scleroris, GEO: GSE106648); Whole blood-LiuRA (335 normal and 354 Rheumatoid arthritis, GEO: GSE42861); Whole blood-TZH (85 normal and 620 Different health conditions, GEO: Requested Access); Whole blood-Ventham (101 normal and 279 Inflammatory bowel disease, GEO: GSE87648); Whole blood-Zannas (144 normal and 278 Trauma, GEO: GSE72680). The DNAm-atlas encompassing DNAm reference matrices for 13 tissue-types encompassing over 40 cell-types is freely available from the EpiSCORE R-package https://github.com/aet21/EpiSCORE

Declarations

Ethics approval and consent to participate Not applicable here since we have only analyzed existing publicly available data.

Consent for publication: Not applicable.

Competing interests: The authors declare no competing interests.

Funding: This work was supported by NSFC (National Science Foundation of China) grants, grant numbers 31970632 and 32370699.

Authors contributions: TH, GX and AET performed the statistical and bioinformatic analyses. AET conceived the study, and wrote the manuscript with contributions from TH and GX. QL contributed to the data-analyses. MJ and NE contributed pointers to data and helped revise the manuscript.

Acknowledgements: We thank everyone who supports open-access data.

References

1.   Baker, G.T., 3rd & Sprott, R.L. Biomarkers of aging. *Exp Gerontol* **23**, 223-39 (1988).

2.   Hadley, E.C., Lakatta, E.G., Morrison-Bogorad, M., Warner, H.R. & Hodes, R.J. The future of

aging therapies. *Cell* **120**, 557-67 (2005).

3. Levine, M.E. *et al.* An epigenetic biomarker of aging for lifespan and healthspan. *Aging (Albany NY)* **10**, 573-591 (2018).

4. Moqri, M. *et al.* Biomarkers of aging for the identification and evaluation of longevity interventions. *Cell* **186**, 3758-3775 (2023).

5. Beck, S. Taking the measure of the methylome. *Nat Biotechnol* **28**, 1026-8 (2010).

6. Bell, C.G. Epigenomic insights into common human disease pathology. *Cell Mol Life Sci* **81**, 178 (2024).

7. Seale, K., Horvath, S., Teschendorff, A., Eynon, N. & Voisin, S. Making sense of the ageing methylome. *Nat Rev Genet* **23**, 585-605 (2022).

8. Horvath, S. & Raj, K. DNA methylation-based biomarkers and the epigenetic clock theory of ageing. *Nat Rev Genet* **19**, 371-384 (2018).

9. Field, A.E., Robertson, N.A., Wang, T., Havas, A., Ideker, T. & Adams, P.D. DNA Methylation Clocks in Aging: Categories, Causes, and Consequences. *Mol Cell* **71**, 882-895 (2018).

10. Bell, C.G. *et al.* DNA methylation aging clocks: challenges and recommendations. *Genome Biol* **20**, 249 (2019).

11. Zhang, Q. *et al.* Improved precision of epigenetic clock estimates across tissues and its implication for biological ageing. *Genome Med* **11**, 54 (2019).

12. Lu, A.T. *et al.* DNA methylation GrimAge strongly predicts lifespan and healthspan. *Aging (Albany NY)* **11**, 303-327 (2019).

13. Levine, M.E., Hosgood, H.D., Chen, B., Absher, D., Assimes, T. & Horvath, S. DNA methylation age of blood predicts future onset of lung cancer in the women's health initiative. *Aging (Albany NY)* **7**, 690-700 (2015).

14. Levine, M.E. *et al.* Menopause accelerates biological aging. *Proc Natl Acad Sci U S A* **113**, 9327-32 (2016).

15. Marioni, R.E. *et al.* DNA methylation age of blood predicts all-cause mortality in later life. *Genome Biol* **16**, 25 (2015).

16. Marioni, R.E. *et al.* The epigenetic clock is correlated with physical and cognitive fitness in the Lothian Birth Cohort 1936. *Int J Epidemiol* **44**, 1388-96 (2015).

17. de Magalhaes, J.P. Distinguishing between driver and passenger mechanisms of aging. *Nat Genet* **56**, 204-211 (2024).

18. Teschendorff, A.E. & Relton, C.L. Statistical and integrative system-level analysis of DNA methylation data. *Nat Rev Genet* (2017).

19. Jaffe, A.E. & Irizarry, R.A. Accounting for cellular heterogeneity is critical in epigenome-wide association studies. *Genome Biol* **15**, R31 (2014).

20. Jonkman, T.H. *et al.* Functional genomics analysis identifies T and NK cell activation as a driver of epigenetic clock progression. *Genome Biol* **23**, 24 (2022).

21. Luo, Q. *et al.* A meta-analysis of immune-cell fractions at high resolution reveals novel associations with common phenotypes and health outcomes. *Genome Med* **15**, 59 (2023).

22. Maity, A.K. & Teschendorff, A.E. Cell-attribute aware community detection improves differential abundance testing from single-cell RNA-Seq data. *Nat Commun* **14**, 3244 (2023).

23. Zheng, S.C. *et al.* A novel cell-type deconvolution algorithm reveals substantial contamination by immune cells in saliva, buccal and cervix. *Epigenomics* (2018).

24. Horvath, S. DNA methylation age of human tissues and cell types. *Genome Biol* **14**, R115

(2013).

25. Levine, M.E., Lu, A.T., Bennett, D.A. & Horvath, S. Epigenetic age of the pre-frontal cortex is associated with neuritic plaques, amyloid load, and Alzheimer's disease related cognitive functioning. *Aging (Albany NY)* **7**, 1198-211 (2015).

26. Levine, A.J. *et al.* Accelerated epigenetic aging in brain is associated with pre-mortem HIV-associated neurocognitive disorders. *J Neurovirol* **22**, 366-75 (2016).

27. Levine, A.J., Panos, S.E. & Horvath, S. Genetic, transcriptomic, and epigenetic studies of HIV-associated neurocognitive disorder. *J Acquir Immune Defic Syndr* **65**, 481-503 (2014).

28. Lu, A.T. *et al.* Genetic variants near MLST8 and DHX57 affect the epigenetic age of the cerebellum. *Nat Commun* **7**, 10561 (2016).

29. Lu, A.T. *et al.* Genetic architecture of epigenetic and neuronal ageing rates in human brain regions. *Nat Commun* **8**, 15353 (2017).

30. Murthy, M. *et al.* Epigenetic age acceleration is associated with oligodendrocyte proportions in MSA and control brain tissue. *Neuropathol Appl Neurobiol* **49**, e12872 (2023).

31. Shireby, G. *et al.* DNA methylation signatures of Alzheimer's disease neuropathology in the cortex are primarily driven by variation in non-neuronal cell-types. *Nat Commun* **13**, 5620 (2022).

32. Shireby, G.L. *et al.* Recalibrating the epigenetic clock: implications for assessing biological age in the human cortex. *Brain* **143**, 3763-3775 (2020).

33. Brunet, A., Goodell, M.A. & Rando, T.A. Ageing and rejuvenation of tissue stem cells and their niches. *Nat Rev Mol Cell Biol* **24**, 45-62 (2023).

34. Chondronasiou, D. *et al.* Multi-omic rejuvenation of naturally aged tissues by a single cycle of transient reprogramming. *Aging Cell* **21**, e13578 (2022).

35. de Magalhaes, J.P. & Ocampo, A. Cellular reprogramming and the rise of rejuvenation biotech. *Trends Biotechnol* **40**, 639-642 (2022).

36. Gill, D. *et al.* Multi-omic rejuvenation of human cells by maturation phase transient reprogramming. *Elife* **11**(2022).

37. Cole, J.J. *et al.* Diverse interventions that extend mouse lifespan suppress shared age-associated epigenetic changes at critical gene regulatory regions. *Genome Biol* **18**, 58 (2017).

38. Wang, T. *et al.* Epigenetic aging signatures in mice livers are slowed by dwarfism, calorie restriction and rapamycin treatment. *Genome Biol* **18**, 57 (2017).

39. Trapp, A., Kerepesi, C. & Gladyshev, V.N. Profiling epigenetic age in single cells. *Nat Aging* **1**, 1189-1201 (2021).

40. Zhu, T. *et al.* A pan-tissue DNA methylation atlas enables in silico decomposition of human tissue methylomes at cell-type resolution. *Nat Methods* **19**, 296-306 (2022).

41. Zheng, S.C., Breeze, C.E., Beck, S. & Teschendorff, A.E. Identification of differentially methylated cell types in epigenome-wide association studies. *Nat Methods* **15**, 1059-1066 (2018).

42. Zhang, Z. *et al.* Hierarchical deconvolution for extensive cell type resolution in the human brain using DNA methylation. *Front Neurosci* **17**, 1198243 (2023).

43. Chien, J.F. Cell type-specific effects of age and sex on human cortical neurons. *bioRxiv* (2024).

44. Cai, M., Zhou, J., McKennan, C. & Wang, J. scMD facilitates cell type deconvolution using single-cell DNA methylation references. *Commun Biol* **7**, 1 (2024).

45. Friedman, J., Hastie, T. & Tibshirani, R. Regularization Paths for Generalized Linear Models via Coordinate Descent. *J Stat Softw* **33**, 1-22 (2010).

46. Jaffe, A.E. *et al.* Mapping DNA methylation across development, genotype and schizophrenia in the human frontal cortex. *Nat Neurosci* **19**, 40-7 (2016).

47. You, C. *et al.* A cell-type deconvolution meta-analysis of whole blood EWAS reveals lineage-specific smoking-associated DNA methylation changes. *Nat Commun* **11**, 4779 (2020).

48. Pihlstrom, L. *et al.* Epigenome-wide association study of human frontal cortex identifies differential methylation in Lewy body pathology. *Nat Commun* **13**, 4932 (2022).

49. Teschendorff, A.E., Zhu, T., Breeze, C.E. & Beck, S. EPISCORE: cell type deconvolution of bulk tissue DNA methylomes from single-cell RNA-Seq data. *Genome Biol* **21**, 221 (2020).

50. Johnson, N.D. *et al.* Differential DNA methylation and changing cell-type proportions as fibrotic stage progresses in NAFLD. *Clin Epigenetics* **13**, 152 (2021).

51. Teschendorff, A.E. Avoiding common pitfalls in machine learning omic data science. *Nat Mater* **18**, 422-427 (2019).

52. Bonder, M.J. *et al.* Genetic and epigenetic regulation of gene expression in fetal and adult human livers. *BMC Genomics* **15**, 860 (2014).

53. Wang, L. *et al.* Alpha-1 Antitrypsin Deficiency Liver Disease, Mutational Homogeneity Modulated by Epigenetic Heterogeneity With Links to Obesity. *Hepatology* **70**, 51-66 (2019).

54. Horvath, S. *et al.* Obesity accelerates epigenetic aging of human liver. *Proc Natl Acad Sci U S A* **111**, 15538-43 (2014).

55. Ahrens, M. *et al.* DNA methylation analysis in nonalcoholic fatty liver disease suggests distinct disease-specific and remodeling signatures after bariatric surgery. *Cell Metab* **18**, 296-302 (2013).

56. Ying, K. *et al.* Causality-enriched epigenetic age uncouples damage and adaptation. *Nat Aging* **4**, 231-246 (2024).

57. Battram, T. *et al.* The EWAS Catalog: a database of epigenome-wide association studies. *Wellcome Open Res* **7**, 41 (2022).

58. Xiong, Z. *et al.* EWAS Open Platform: integrated data, knowledge and toolkit for epigenome-wide association study. *Nucleic Acids Res* **50**, D1004-D1009 (2022).

59. Gaunt, T.R. *et al.* Systematic identification of genetic influences on methylation across the human life course. *Genome Biol* **17**, 61 (2016).

60. Bellenguez, C. *et al.* New insights into the genetic etiology of Alzheimer's disease and related dementias. *Nat Genet* **54**, 412-436 (2022).

61. Andrade-Guerrero, J. *et al.* Alzheimer's Disease: An Updated Overview of Its Genetics. *Int J Mol Sci* **24**(2023).

62. Hassan, M., Shahzadi, S., Alashwal, H., Zaki, N., Seo, S.Y. & Moustafa, A.A. Exploring the mechanistic insights of Cas scaffolding protein family member 4 with protein tyrosine kinase 2 in Alzheimer's disease by evaluating protein interactions through molecular docking and dynamic simulations. *Neurol Sci* **39**, 1361-1374 (2018).

63. Barreby, E., Chen, P. & Aouadi, M. Macrophage functional diversity in NAFLD - more than inflammation. *Nat Rev Endocrinol* **18**, 461-472 (2022).

64. Zhu, T., Zheng, S.C., Paul, D.S., Horvath, S. & Teschendorff, A.E. Cell and tissue type independent age-associated DNA methylation changes are not rare but common. *Aging (Albany NY)* **10**, 3541-3557 (2018).

65. Aryee, M.J. *et al.* Minfi: a flexible and comprehensive Bioconductor package for the analysis of Infinium DNA methylation microarrays. *Bioinformatics* **30**, 1363-9 (2014).

66. Troyanskaya, O. *et al.* Missing value estimation methods for DNA microarrays. *Bioinformatics* **17**, 520-5 (2001).

67. Teschendorff, A.E. *et al.* A beta-mixture quantile normalization method for correcting probe design bias in Illumina Infinium 450 k DNA methylation data. *Bioinformatics* **29**, 189-96 (2013).

68. Park, C.S. *et al.* Hepatocyte gene expression and DNA methylation as ancestry-dependent mechanisms in African Americans. *NPJ Genom Med* **4**, 29 (2019).

69. Bacalini, M.G. *et al.* Molecular Aging of Human Liver: An Epigenetic/Transcriptomic Signature. *J Gerontol A Biol Sci Med Sci* **74**, 1-8 (2019).

70. Diez-Villanueva, A. *et al.* DNA methylation events in transcription factors and gene expression changes in colon cancer. *Epigenomics* **12**, 1593-1610 (2020).

71. Kirby, M.K. *et al.* Genome-wide DNA methylation measurements in prostate tissues uncovers novel prostate cancer diagnostic biomarkers and transcription factor binding patterns. *BMC Cancer* **17**, 273 (2017).

72. Huan, T. *et al.* Epigenome-wide association study of DNA methylation and microRNA expression highlights novel pathways for human complex traits. *Epigenetics* **15**, 183-198 (2020).

73. Roos, L. *et al.* Higher Nevus Count Exhibits a Distinct DNA Methylation Signature in Healthy Human Skin: Implications for Melanoma. *J Invest Dermatol* **137**, 910-920 (2017).

74. Robinson, O. *et al.* Determinants of accelerated metabolomic and epigenetic aging in a UK cohort. *Aging Cell* **19**, e13149 (2020).

75. Lehne, B. *et al.* A coherent approach for analysis of the Illumina HumanMethylation450 BeadChip improves data quality and performance in epigenome-wide association studies. *Genome Biol* **16**, 37 (2015).

76. Teschendorff, A.E., Breeze, C.E., Zheng, S.C. & Beck, S. A comparison of reference-based algorithms for correcting cell-type heterogeneity in Epigenome-Wide Association Studies. *BMC Bioinformatics* **18**, 105 (2017).

77. Hannum, G. *et al.* Genome-wide methylation profiles reveal quantitative views of human aging rates. *Mol Cell* **49**, 359-67 (2013).

78. Chen, L. *et al.* Genetic Drivers of Epigenetic and Transcriptional Variation in Human Immune Cells. *Cell* **167**, 1398-1414 e24 (2016).

79. Borenstein, M., Hedges, L.V., Higgins, J.P. & Rothstein, H.R. A basic introduction to fixed-effect and random-effects models for meta-analysis. *Res Synthesis Meth* **1**, 97-111 (2010).

80. Balduzzi, S., Rucker, G. & Schwarzer, G. How to perform a meta-analysis with R: a practical tutorial. *Evid Based Ment Health* **22**, 153-160 (2019).

81. Pulit, S.L. *et al.* Meta-analysis of genome-wide association studies for body fat distribution in 694 649 individuals of European ancestry. *Hum Mol Genet* **28**, 166-174 (2019).

Figure Legends

Figure-1: Quantification of intrinsic and extrinsic aging. a) Graphical depiction of how the predictive accuracy of an epigenetic clock derived from bulk-tissue can be decomposed in terms of an extrinsic and intrinsic aging process, as shown. The accuracy, measured by the R^2 value, of a clock reflecting the extrinsic process can be estimated from the given formula, if one can estimate the accuracies of the full unadjusted clock and the one fully adjusted for cell-type heterogeneity (CTH), the latter reflecting the intrinsic aging of the individual cell-types. **b)** Barplots compare the R^2 values of the unadjusted clock, and clocks adjusted for 12 and 9 immune cell types, in each of 11 whole blood cohorts. Right panel displays the average and standard deviation of the two indicated ratios across the 11 cohorts. These ratios measure the fraction of the unadjusted clock's accuracy that can be attributed to intrinsic aging. **c)** As b) but for brain DNAm datasets, where the adjusted clock was adjusted for 7 brain cell-types.

Figure-2: Meta-analysis of brain cell-type fractions with age. a) Forest plots of associations between brain cell-type fractions and age, for 7 brain cell-types (ExN=excitatory neurons, InN=inhibitory neurons, Oligo/OPC=oligodendrocytes/oligo precursor cells, Micro=microglia, Astro=astrocytes, Stromal, Endo=endothelial) across 13 independent DNAm datasets. Number of samples in each cohort is given in brackets. P-values from a Fixed and Random Effects models are also given. Heterogeneity statistics and P-values of heterogeneity are given above each plot. **b)** Effect sizes and P-values from the Random Effects meta-analysis model.

Figure-3: Computational strategy for building cell-type specific epigenetic clocks. a) Given a tissue-type of interest with DNAm profiles measured over a relatively large number of samples, we first estimate the underlying cell-type proportions in each sample using an existing tissue-specific DNAm reference matrix. Density plots depict the distribution of cell-type fractions across all samples. Next, using a sufficiently large training set of samples encompassing a relatively wide age-range, we apply the CellDMC algorithm to infer age-associated DNAm changes in each of the underlying cell-types (age-DMCTs). Barplot depicts the number of age-DMCTs in each cell-type. **b)** The construction of cell-type specific clocks then proceeds by restricting to age-DMCTs of one cell-type: an intrinsic clock is built by adjusting the DNAm training dataset for variations in cell-type fractions (CTFs), defining a matrix of DNAm residuals. Alternatively, the training over the age-DMCTs can be done on

the DNAm data matrix without adjustment for CTFs which will result in a "semi-intrinsic" clock. In either case, Elastic Net models are learned for each choice of a penalty parameter L, and the optimal model $L*$ is selected based on the best generalization performance (smallest root mean square error (RMSE)) obtained in a blinded model selection set. This optimal model then defines the corresponding cell-type specific DNAm-clock. This procedure can be done for each cell-type separately, assuming that sufficient numbers of age-DMCTs in that cell-type can be identified. Once the cell-type specific clocks are built, these are then validated in independent DNAm datasets. **c)** Top row: boxplots display the root mean square error (RMSE) between predicted and true ages, for semi-intrinsic (Sin) and intrinsic (In) clocks, as assessed in 50 simulated validation sets of 200 mixtures for four different scenarios. Cell-type specific clocks were constructed from 50 simulated training sets of 200 mixtures (mixing together 3 cell-types that we call granulocytes, monocytes and lymphocytes) with age-DMCTs occurring only in one cell-type (lymphocytes). In scenarios (i) and (iii), no cell-type fractions change with age, in scenarios (ii) and (iv) the lymphocyte fraction changes with age. In scenarios (i) and (ii), age-DMCTs do not discriminate immune-cell types from each other, in scenarios (iii) and (iv) age-DMCTs discriminate lymphocytes from granulocytes and monocytes. P-values from a two-tailed paired Wilcoxon test are given. Bottom row: as top-row but now displaying the PCC of predicted vs true age.

Figure-4: Construction and validation of neuron-clocks. a) The training is done on the Jaffe et al dataset encompassing 416 prefrontal cortex (PFC) samples. Fractions for 3 broad cell-types (neurons, glia and endothelial/stromal) are inferred using the HiBED algorithm. Subsequently, CellDMC is applied to infer DMCTs in each cell-type. The Fisher-test one-tailed P-value of overlap between Neu-DMCTs and Glia-DMCTs is given. Finally, elastic net predictors of chronological age are trained from the Neu-DMCTs parameterized by a penalty parameter. **b)** Scatterplots of adjusted DNAm values (adjusted for variations in cell-type fractions) of one Neu-DMCT and one Glia-DMCT against the corresponding Neuron and Glia fraction, respectively. Samples have been colored according to age-group and regression lines for each age-group have been added. **c)** Optimal model (parameter) selection is performed using a completely independent PFC dataset from Philstrom et al. Scatterplot displays the predicted vs chronological age for the optimal model. **d)** Barplot of Pearson Correlation Coefficients (PCC) for 4 clocks (Neu-In, Neu-Sin, Brain and Horvath) across sorted neuron and glia datasets. Number of samples in each dataset is given in brackets. **e)** Barplot of Pearson Correlation Coefficients (PCC) for the same 4 clocks (Neu-In, Neu-Sin, Brain and Horvath) across brain and non-brain DNAm datasets. Each cohort is labeled with the number of samples in brackets. WB=whole blood, PFC=prefrontal cortex, TL=temporal lobe, FL=frontal lobe, MTG/STG=middle/superior temporal gyrus, PL=parietal lobe, SC=sensory cortex, VC=visual cortex, MC=motor cortex, Hippo=hippocampus, EC=entorhinal cortex, Stria=striatum, CG=cingulate gyrus, NAc=nucleus accumbens, CRB=cerebellum, Mono=monocyte. **f)** For each of the 4 clocks, boxplots of PCCs comparing the values

obtained in brain vs non-brain datasets. The number of datasets in each group is given below x-axis. P-value is from a one-tailed Wilcoxon rank sum test.

Figure-5: Neuron and glia specific clocks predict age-acceleration in AD. a) Forest plot of associations between DNAmAge and Alzheimer's Disease (AD) for 6 different epigenetic clocks across 5 independent DNAm datasets from frontal lobe (FL) or prefrontal cortex (PFC). The effect sizes and P-values of a fixed and random effects models are also given. Statistics and P-values of heterogeneity are shown above each panel. Number of samples in each dataset is given to the left. **b)** As a) but for DNAm datasets profiling regions in the temporal lobe (TL), which includes middle and superior temporal gyrus (MTG and STG). **c)** As a), but combing FL and TL regions.

Figure-6: Construction and validation of hepatocyte clock. a) Overall strategy used to train a hepatocyte specific clock from bulk liver tissue samples. Plots show how the root mean square error (RMSE) changes as a function of the penalty parameter in the elastic net regression model, and scatterplot at the bottom compares the predicted and chronological ages for the optimal model defining the hepatocyte-clock. Pearson Correlation Coefficient (PCC), Median Absolute Error (MedAE) and regression P-value are given. **b)** Validation of the hepatocyte and liver clocks (latter trained on age-DMCs, not hepatocyte-DMCTs) in a primary hepatocyte culture. Also shown is the result for Horvath's clock. PCC, MedAE and P-values from a linear regression are given. **c)** Barplot comparing the PCC-values of the hepatocyte and liver clocks across several liver and non-liver tissue datasets. Boxplots compare the PCC values of the hepatocyte and liver clocks in the liver/hepatocyte dataset compared to all others. P-value is from a one-tailed Wilcoxon rank sum test. **d)** Boxplots compare the extrinsic age-acceleration (EAA) of the liver, hepatocyte and Horvath clocks in two independent liver DNAm datasets, one profiling non-alcoholic fatty liver disease (NAFLD) cases and the other profiling liver-tissue from obese (BMI>35) individuals. In each case, the P-value is from a one-tailed Wilcoxon rank sum test comparing the distribution to a mean value of 0 (red dashed line). **e)** Boxplots of EAA from the three clocks in an independent DNAm liver dataset, with samples stratified according to BMI levels, as shown. For each panel, we give the PCC and linear regression test P-value.

a) Quantification of the two components of tissue-specific epigenetic clocks

UNADJ (unadjusted for CTH)

Extrinsic Aging

Aging due to changes in cell-type composition

Intrinsic Aging

Aging due to DNAm changes in individual cell-types

ADJ (adjusted for CTH)

$$R^2_{Extrinsic} = 1 - R^2_{Intrinsic} = 1 - \frac{R^2_{ADJ}}{R^2_{UNADJ}}$$

b)

Whole Blood DNAm datasets

R^2

Barturen(101) Flanagan(184) Hannon1(304) Hannon2(405) Hannum(656) Johansson(729) LiuMS(139) LiuRA(335) TZH(85) Ventham(101) Zannas(144)

Ratio

Average

R^2
- ADJ-12CT
- ADJ-9CT
- UNADJ

Ratio
- R²(12CT)/R²(UNADJ)
- R²(9CT)/R²(UNADJ)

c)

Brain DNAm datasets

R^2

Gasparoni(26) Horvath(20) Huynh(19) Murphy(38) Pidsley(24) Pihlstrom(67) Rydbirk(37) Semick(47) Smith(68) Torabi(50) Xu(25)

Cohort

Ratio

Average

R^2
- ADJ
- UNADJ

Ratio
- R²(ADJ)/R²(UNADJ)