TimeFlies: an snRNA-seq aging clock for the fruit fly head sheds light on sex-biased aging

Nikolai Tennant[1]*, Ananya Pavuluri[2]*, Kate O'Connor-Giles[5,6], Gunjan Singh[4,5], Erica Larschan[2,4**], and Ritambhara Singh[1,2,3**]

*Equal contribution

**Co-corresponding authors

1. Data Science Institute, Brown University, Providence, RI, USA
2. Center for Computational Molecular Biology, Brown University, Providence, RI, USA
3. Department of Computer Science, Brown University, Providence, RI, USA
4. Department of Molecular Biology, Cell Biology, and Biochemistry, Brown University, Providence, RI, USA
5. Department of Neuroscience, Brown University, Providence, RI, USA
6. Carney Institute for Brain Science, Brown University, Providence, RI, USA

Abstract

Keywords: *Drosophila melanogaster, transcriptomics, deep learning, single-cell, aging clock, dosage compensation*

Although multiple high-performing epigenetic aging clocks exist, few are based directly on gene expression. Such transcriptomic aging clocks allow us to extract age-associated genes directly. However, most existing transcriptomic clocks model a subset of genes and are limited in their ability to predict novel biomarkers. With the growing popularity of single-cell sequencing, there is a need for robust single-cell transcriptomic aging clocks. Moreover, clocks have yet to be applied to investigate the elusive phenomenon of sex differences in aging. We introduce TimeFlies, a pan-cell-type scRNA-seq aging clock for the *Drosophila melanogaster* head. TimeFlies uses deep learning to classify the donor age of cells based on genome-wide gene expression profiles. Using explainability methods, we identified key marker genes contributing to the classification, with lncRNAs showing up as highly enriched among predicted biomarkers. The top biomarker gene across cell types is lncRNA:*roX1*, a regulator of X chromosome dosage compensation, a pathway previously identified as a top biomarker of aging in the mouse brain. We validated this finding experimentally, showing a decrease in survival probability in the absence of roX1 *in vivo*. Furthermore, we trained sex-specific TimeFlies clocks and noted significant differences in model predictions and explanations between male and female clocks, suggesting that different pathways drive aging in males and females.

Graphical Abstract

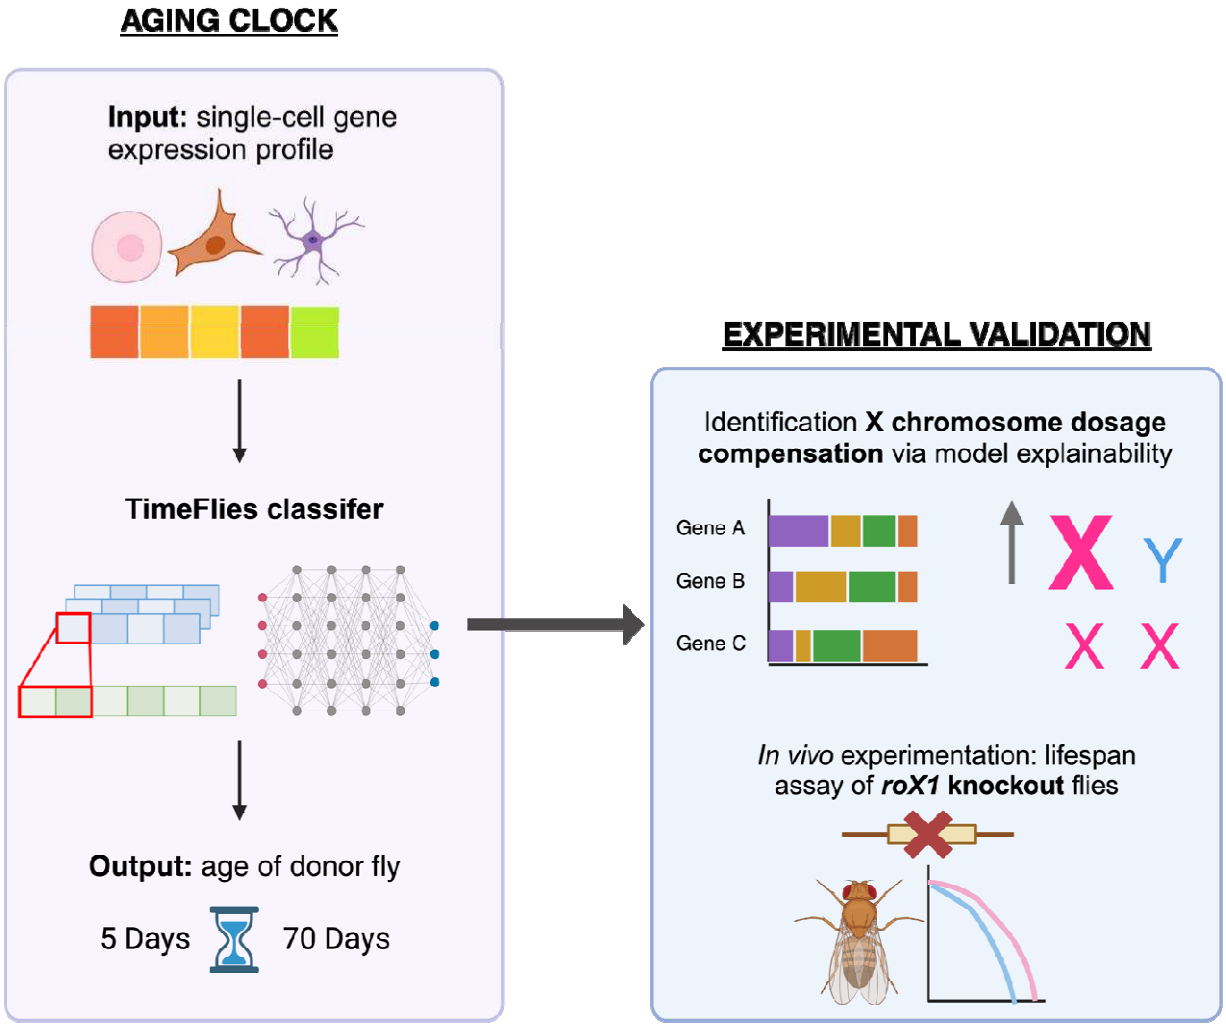

Introduction

Aging is characterized by time-related dysfunction and accrued damage in an organism. Lopez-Otín et al. have suggested twelve hallmarks of aging at the molecular, cellular, and systemic levels which underlie age-associated phenotypes [1]. A priority in the field of aging research has been the development of "aging clocks," statistical estimators that determine the donor age of a sample based on biological measurements. These clocks allow us to discover candidate biomarkers associated with the key hallmarks of aging.

The vast majority of published aging clocks are based on DNA methylation (DNAm) data. The first aging clocks were published by Hannum et al. [2] and Horvath [3]. Hannum et al. developed an ElasticNet-based model that predicts human age from whole blood samples based on bulk DNAm levels at 71 CpG sites [2]. Horvath then developed a more robust DNAm clock, generalizable across 51 human tissue types– and even to chimpanzee tissue– utilizing 353 CpG sites [3]. A few groups have since used methylation marks to augment other clinical data points of interest in aging clock development, with the goal of understanding mortality risk and disease in the context of aging [4,5]. As many methylation marks are highly conserved, there has been an increased interest in using aging clocks to study the comparative biology of aging. Recently, the Horvath group has developed a pan-Mammalian clock that generalizes to 185 mammal species [6]. In summary, DNAm clocks exhibit high performance and have proven to be generalizable across species. The associations between DNA methylation and aging phenotypes have been widely studied for the past several decades [7,8,9,10], and the aforementioned clocks allow us to deepen our understanding.

While DNAm clocks have shown reliably high performance and have myriad contributions to various avenues of geroscience, it can be difficult to validate and apply their findings. Epigenetic alterations, like DNA methylation, ultimately underlie changes in gene expression. DNAm aging clocks require considerable downstream analysis to determine which genes are proximal to CpG site biomarkers. Furthermore, many CpG sites identified by aging clocks are not explicitly associated with specific genes, making their significance to gene regulation events difficult to analyze. Thus, transcriptomic aging clocks have the potential to unveil more direct associations between genes of interest and aging phenotypes. Identifying such genes as biomarkers of aging will provide researchers with more easily manipulable targets for experiments, as modification of gene expression and disruption of gene products via small molecules is more feasible to implement [11].

Progress in bulk transcriptomic aging clocks has been limited due to the plethora of challenges that come with transcriptomic data. Gene fusion, alternative splicing, and post-transcriptional modifications add layers of complexity to the RNA-seq and microarray data that are difficult to disentangle. One of the first transcriptomic aging clocks fit to human peripheral blood samples

obtained significant correlations between predicted and actual age, although there was a high variability across cohorts [12]. Furthermore, these clocks were trained on microarray data, a technique that has become outdated since the advent of RNA-seq due to limited dynamic range. Fleischer et al. developed a suite of regression models for an internally collected dataset of human dermal fibroblasts, which achieved noteworthy performance (r=0.81). However, this clock was not tested on external data [13]. Meyer and Schumacher found that simply binarizing RNA-seq data—that is, assigning expression values of either 0 or 1—significantly improved the performance of their *Caenorhabditis elegans* aging clock [14]. However, as gene expression exists on a continuum and is highly variable in nature, binarizing the data results in the loss of information—a binarized dataset does not properly reflect the nuances of gene expression dynamics with aging. Furthermore, the authors had to perform feature selection for an optimal set of clock genes rather than using all features in the dataset. A recent methylation clock paper showed that using all available CpG sites rather than a subset both improved model performance and created a more robust model [15], which may translate to similar results in transcriptomic clocks. Holzscheck et al. published a novel gene set-based, knowledge-primed transcriptomic aging clock using deep neural networks. This methodology yields successful performance and is highly interpretable at the pathway-level [16] but requires significant feature engineering. Genes with unknown functions would also be omitted, limiting the potential to discover new age-associated genes. Overall, while there has been progress in the development of high-performance transcriptomic aging clocks, we have yet to fully harness their potential for transcriptome-wide analysis and discovery of novel biomarkers.

Recently, there has been a rise in the popularity of single-cell sequencing, because single-cell resolution unmasks the heterogeneity within biological signals, most of which are highly cell-type-specific. From a statistical and machine learning perspective, single-cell datasets have thousands of samples, thus eliminating the need to integrate several independent bulk RNA-seq datasets and address batch effects. Several single-cell aging atlases have been published, including the Tabula Muris Senis [17], the Cell Atlas of Worm Aging [18], and the Aging Fly Cell Atlas (AFCA) [19]. These data allow us to examine the dynamics of aging in different cell populations of interest. However, single-cell data poses a diverse array of computational challenges. Notably, single-cell RNA-seq often has very high dropout rates compared to bulk RNA-seq, in which the data only reflect a fraction of the cell's gene expression. This results in highly sparse data (a high percentage of zero values). Despite these challenges, Yu et al. successfully created a single-nuclei transcriptomic clock pipeline for the aging female mouse hypothalamus. The most efficient and interpretable model, ElasticNet, which was the focus of the paper, reported an AUPRC of 0.967. However, the authors binarized the input data and subset the features to only highly variable genes rather than using the whole transcriptome [20]. Mao et al. also developed SCALE, a framework to assign a tissue-specific relative aging score at single-cell resolution to samples from the Tabula Muris Senis. However, this pipeline requires users to identify tissue-specific aging-related gene sets as input features, thus limiting the scope of novel

biomarker discovery [21]. Thus, there is a gap in the field because we currently lack an interpretable single-cell transcriptomic aging clock that allows for a comprehensive transcriptome-wide analysis of aging signatures for biomarker discovery.

It is known that lifespan and healthspan are sexually dimorphic across diverse species in the animal kingdom. Yet, the innate biological mechanisms that underlie sex differences in aging remain poorly understood [22]. Sex differences in aging are seldom considered in aging research, with studies often treating sex as a confounding variable rather than a source of relevant biological variation. Understanding why aging affects males and females differently across species is fundamental to the comparative biology of aging and, on a translational level, to the development of better interventions for an aging population. As aging clocks are a framework to discover, develop, and validate hypotheses for aging biology, they can help provide insights into sex differences in aging. To our knowledge, aging clocks have yet to be used for a comprehensive study of potential genes and pathways that contribute to sex-biased aging phenotypes in any species. This leaves a crucial gap for us to begin to bridge with our investigation.

We present TimeFlies, a highly robust and accurate aging clock at single-cell resolution. We chose to develop our model based on the AFCA [19] data from the fruit fly *Drosophila melanogaster* because it is a well-studied model organism in genetics and genomics with a short lifespan which allows us to rapidly test our predictions *in vivo*. It is an ideal candidate for studying aging and sex differences in aging due to its relatively short lifespan, extensively annotated reference transcriptome, and a plethora of widely documented genetic manipulation techniques. Furthermore, the *Drosophila* brain is arguably the most well understood across species because all the connections between individual neurons have been mapped, and individual neural circuits can be genetically manipulated [23,24]. Despite these advantages, there is a notable lack of aging clocks in the fruit fly. Thus, we chose to develop an aging clock for the *Drosophila* head as brain aging is such an important area for translational work, especially regarding sex differences.

TimeFlies uses a 1D convolutional neural network to predict age, across the four AFCA [19] time points, from the single-cell gene expression profile. It learns from the genome-wide gene expression signals and does not require *any* feature engineering or noise reduction prior to model training. Our model generalizes across all cell types in the fly head. We have done an in-depth feature explainability analysis of TimeFlies for the discovery of potential aging marker genes. Our model identifies a strong role of X-chromosome dosage compensation in aging dynamics, which we demonstrate is a conserved feature between *Drosophila* and mice despite their long evolutionary distance from each other. Furthermore, we have performed sex-specific aging clock modeling to identify sex-differential transcriptomic aging signatures at single-cell resolution. Our analysis showed stark differences between male-specific and female-specific clocks, identifying

pathways and functions such as vision and synaptic transmission that are affected by aging in a sex-biased manner. Following the computational analysis, we experimentally tested the role of the top TimeFlies clock gene— lncRNA:*roX1* — on fruit fly aging *in vivo* and found that the knockout of this gene has an effect on survival in both male and female flies, validating a transcriptomic biomarker found by TimeFlies. Overall, TimeFlies is a reliable aging clock based on explainable deep learning that has offered valuable insights into transcriptomic aging marker discovery and has opened many new avenues for future study related to sex-specific brain aging.

Results

TimeFlies is a pan-cell-type aging clock for biomarker discovery

TimeFlies generalizes across all cell types in the fly head with state-of-the-art performance (F1 score=0.9416, AUC=0.9515) in age classification (timepoints: Day 5, Day 30, Day 50, Day 70) despite very high variability between cell types (Fig 1a, "ALL"). However, as the dataset is not uniformly distributed across cell types, we also trained cell-type-specific clocks. We chose the five most populous broad cell types to analyze. TimeFlies maintained very high performance across all five of the cell types of interest, with the lowest performance on muscle cells, although the F1 score was still above 0.88 (Fig 1a). Using broader cell-type categories ensured that even the cell-type-specific clocks were highly robust and could generalize across many specific subtypes. Next, we performed model explainability on the pan-cell-type clock. We generated a SHAP summary plot (Fig 1b) of the top 20 features from the pan-cell-type model. We noticed a drop-off in SHAP value magnitude after the 20th feature, hence, we limited our analysis to the top 20. Subsequently, we ran a gene set enrichment analysis on those 20 genes, which indicated that many of the genes are involved in synaptic activity (Fig. 1c). Furthermore, we determined the top features of each of the cell-type-specific clocks. TimeFlies was able to learn unique features for different cell types that were representative of each cell's biological function (Supplementary Fig 2), with some overlapping features (Fig 1d-e). Notably, the three features that were among the most significant for every cell type were three different long non-coding RNAs: lncRNA:*roX1*, lncRNA:*noe*, and lncRNA:*roX2* (Fig 1d).

This result is interesting because the *roX1* and *roX2* genes are long noncoding RNAs encoded on the X chromosome and involved in dosage compensation. This highly conserved process equalizes the levels of X-linked genes between male (XY) and female (XX) organisms. In *Drosophila melanogaster*, this process is achieved by upregulation of the male X chromosome, while in humans, rodents, and other mammals, one of the two female X chromosomes is silenced, which is referred to as X chromosome inactivation (XCI) [25]. *roX1* and *roX2* are essential components of the male-specific lethal (MSL) complex, which facilitates hyperacetylation of H4K16 along the X chromosome targets in males, a modification associated with increased transcriptional activity. The *roX* RNAs help localize the MSL complex to the X

chromosome [26,27]. Despite the differences in structure and length of *roX1* and *roX2*, they have redundant functions due to the presence of a similar stem loop region [28]. X chromosome upregulation is a highly conserved process across species, including in mammals, to tune X-linked gene expression levels throughout development [29]. Fascinatingly, in a single-nuclei RNA-seq study of the aging female mouse hypothalamus, *Xist*, the master regulator of XCI and the mouse analog of the roX genes, was the top feature in an X chromosome-based aging clock of neurons. The authors also showed that *Xist* expression is upregulated with age in some neuronal populations [30]. This suggests that, despite the evolutionary distance between mice and fruit flies, dosage compensation appears to be conserved as a significant component of the aging process. Equally intriguing is the selection of lncRNA:*noe* by the analysis. This gene was discovered by Kim et al. in 1998 [31]. It is abundantly expressed in the central nervous system and encodes a small peptide of 74 amino acid length [31]. However, the function of the noncoding RNA and its peptide product has remained unknown since its discovery. *noe* is located within an intron of the *blot* gene [32], which is a sodium/chloride-dependent neurotransmitter transporter [33]. Notably, expression of *noe* is highly enriched in adult males, with moderate expression in pupae and adult females [34].

Of the top 20 clock genes (Figure 1C), six are lncRNAs. lncRNAs were among the top five clock genes in each cell-type-specific clock (Figure 1E). In the AFCA dataset, there are 15992 expressed genes, of which 2100 are lncRNAs (13.13%). This enrichment of lncRNAs in TimeFlies feature explanations is significant, with 30% of the top features being lncRNAs. Previous studies have identified several age-associated lncRNAs [35] and their evolutionary conservation across species [36]. It has also been shown that, in the fruit fly, lncRNAs and their known targets are differentially expressed during dietary restriction, a well-studied aging intervention [37]. Our feature explainability analysis reflects the age-associated enrichment of lncRNAs and makes the case for further evaluation of lncRNA-based gene regulation during aging.

Remarkably, most features selected by TimeFlies are not among the top 1000 highly variable genes. Of the top 20 SHAP features, only 4 are included in the top 1000 highly variable genes (Supplementary Figure 2). Thus, many genes may have been overlooked with feature selection approaches during preprocessing if using classical differential expression analysis and not performing modeling with TimeFlies. For example, expression patterns of *roX1*, *noe*, and *roX2* do not show significant linear associations with age due to high variability in expression levels at each time point (Supplementary Figure 2). This suggests that TimeFlies detects complex age-associated patterns of expression. Thus, the explainability analysis of TimeFlies offers a more comprehensive biomarker discovery strategy than simple linear models or differential expression analyses.

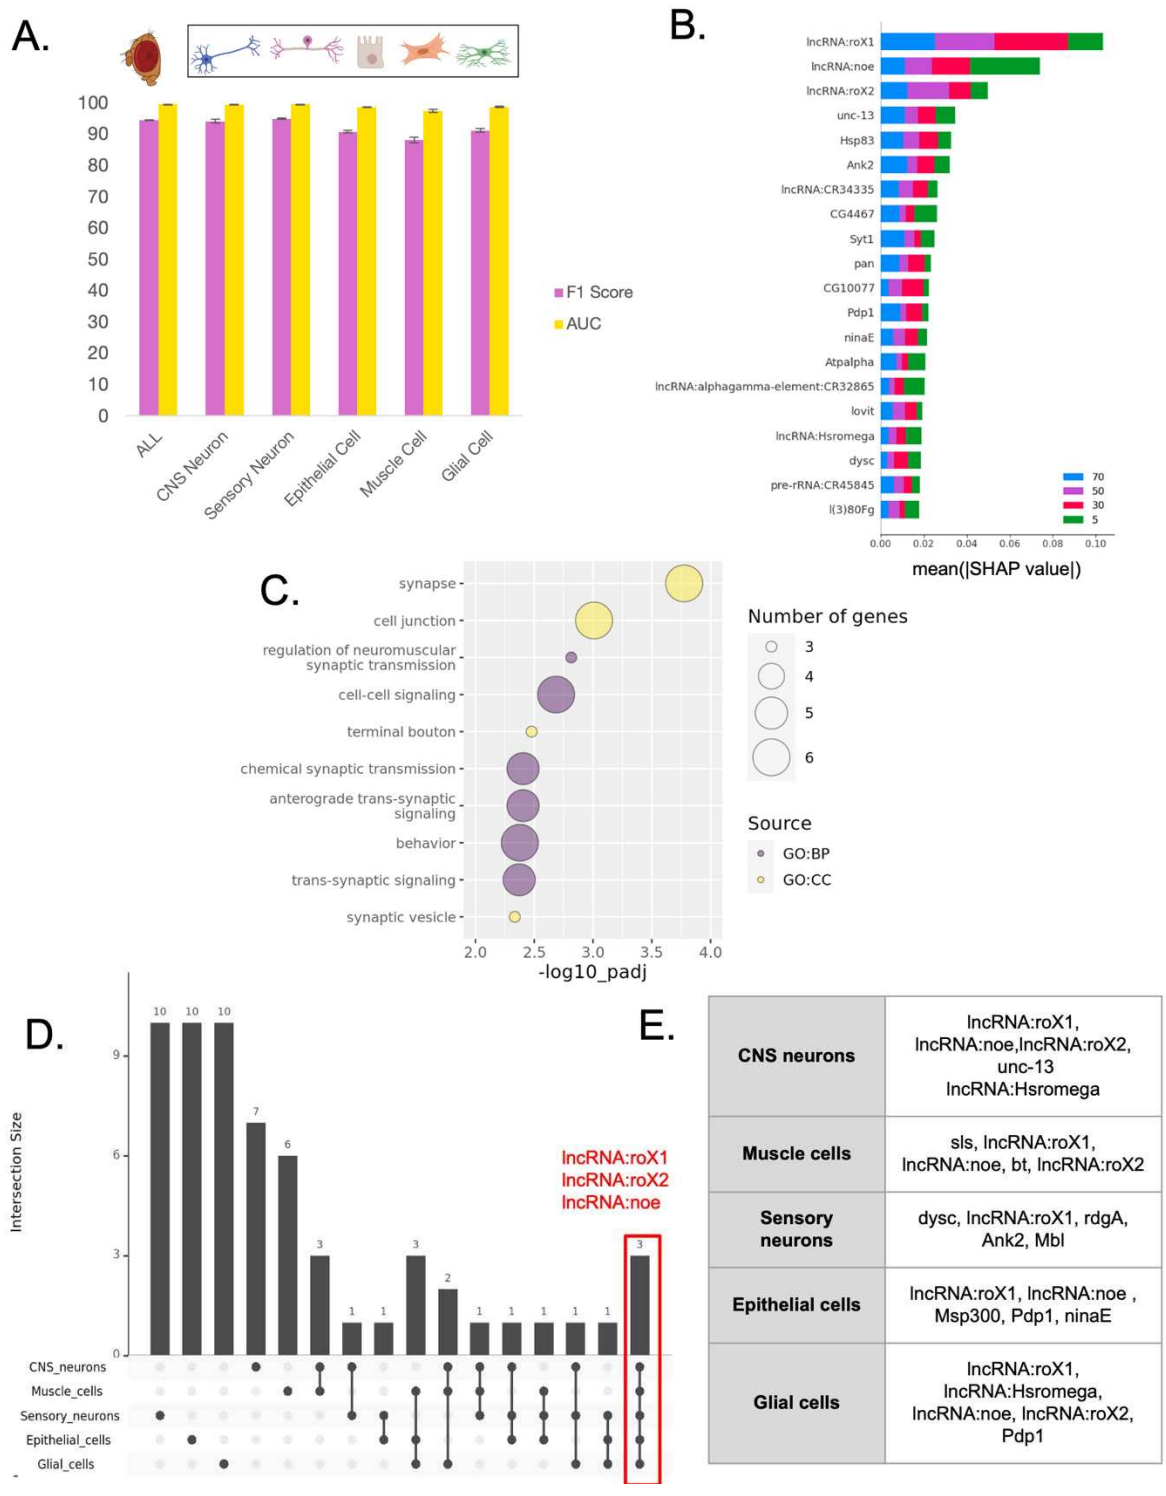

CNS neurons	lncRNA:roX1, lncRNA:noe,lncRNA:roX2, unc-13 lncRNA:Hsromega
Muscle cells	sls, lncRNA:roX1, lncRNA:noe, bt, lncRNA:roX2
Sensory neurons	dysc, lncRNA:roX1, rdgA, Ank2, Mbl
Epithelial cells	lncRNA:roX1, lncRNA:noe , Msp300, Pdp1, ninaE
Glial cells	lncRNA:roX1, lncRNA:Hsromega, lncRNA:noe, lncRNA:roX2, Pdp1

Figure 1. A) TimeFlies age classification performance on test data across five cell types. **B)** SHAP summary plot showing the list of top 20 features used by TimeFlies in the classification task. Bars signify the average impact of the feature on model output magnitude. **C)** Gene set enrichment analysis performed on the list of genes in panel B. Yellow represents GO: cellular component and purple represents GO: biological process. **D)** UpSet plot generated from sets of top 20 SHAP features for each cell type. Chart displays the number of features that are unique to a cell type-specific clock and the number of overlapping genes in each clock combination. Red rectangle highlights the three genes that occur in the SHAP list of every cell type-specific clock: lncRNA:*roX1*, lncRNA:*noe*, and lncRNA:*roX2*. **E)** The top 5 genes of each cell type-specific clock, in descending order of mean SHAP value magnitude.

Sex differences in predictive aging genes

Female fruit flies, on average, have longer lifespans than male fruit flies. Thus, we developed male-specific and female-specific TimeFlies clocks to investigate differences in sex-specific aging biomarkers. We obtained the top 20 genes, following previous analysis, of each sex-specific clock and performed gene set enrichment analysis (GSEA) (Fig 2 a-b, Supplementary Figures 3-4). Interestingly, the female-specific clock used several top features related to the eye (Fig 2b, Supplementary Figure 3). This is particularly noteworthy as many age-associated eye diseases in humans are more prevalent in females [38]. There were nine overlapping genes used by both the male-specific clock and the female-specific clock (Fig 2c). We further tested whether a clock trained on only male data can generalize to female data, and vice versa. As expected, clocks trained on only female samples have a very low performance on male samples, and clocks trained on only male samples have a lowered performance on female samples. The difference is more marked on female clocks tested on male samples (Fig 2d). The female-trained clock tends to classify 30-day-old males as 50-days-old while the male-trained clock tends to classify 30-day-old females as 5-days old (Fig 2e-f). These results reaffirm that aging is a highly sex-biased process, even at single-cell resolution. Thus, it is imperative that sex differences be considered when developing and analyzing aging clocks.

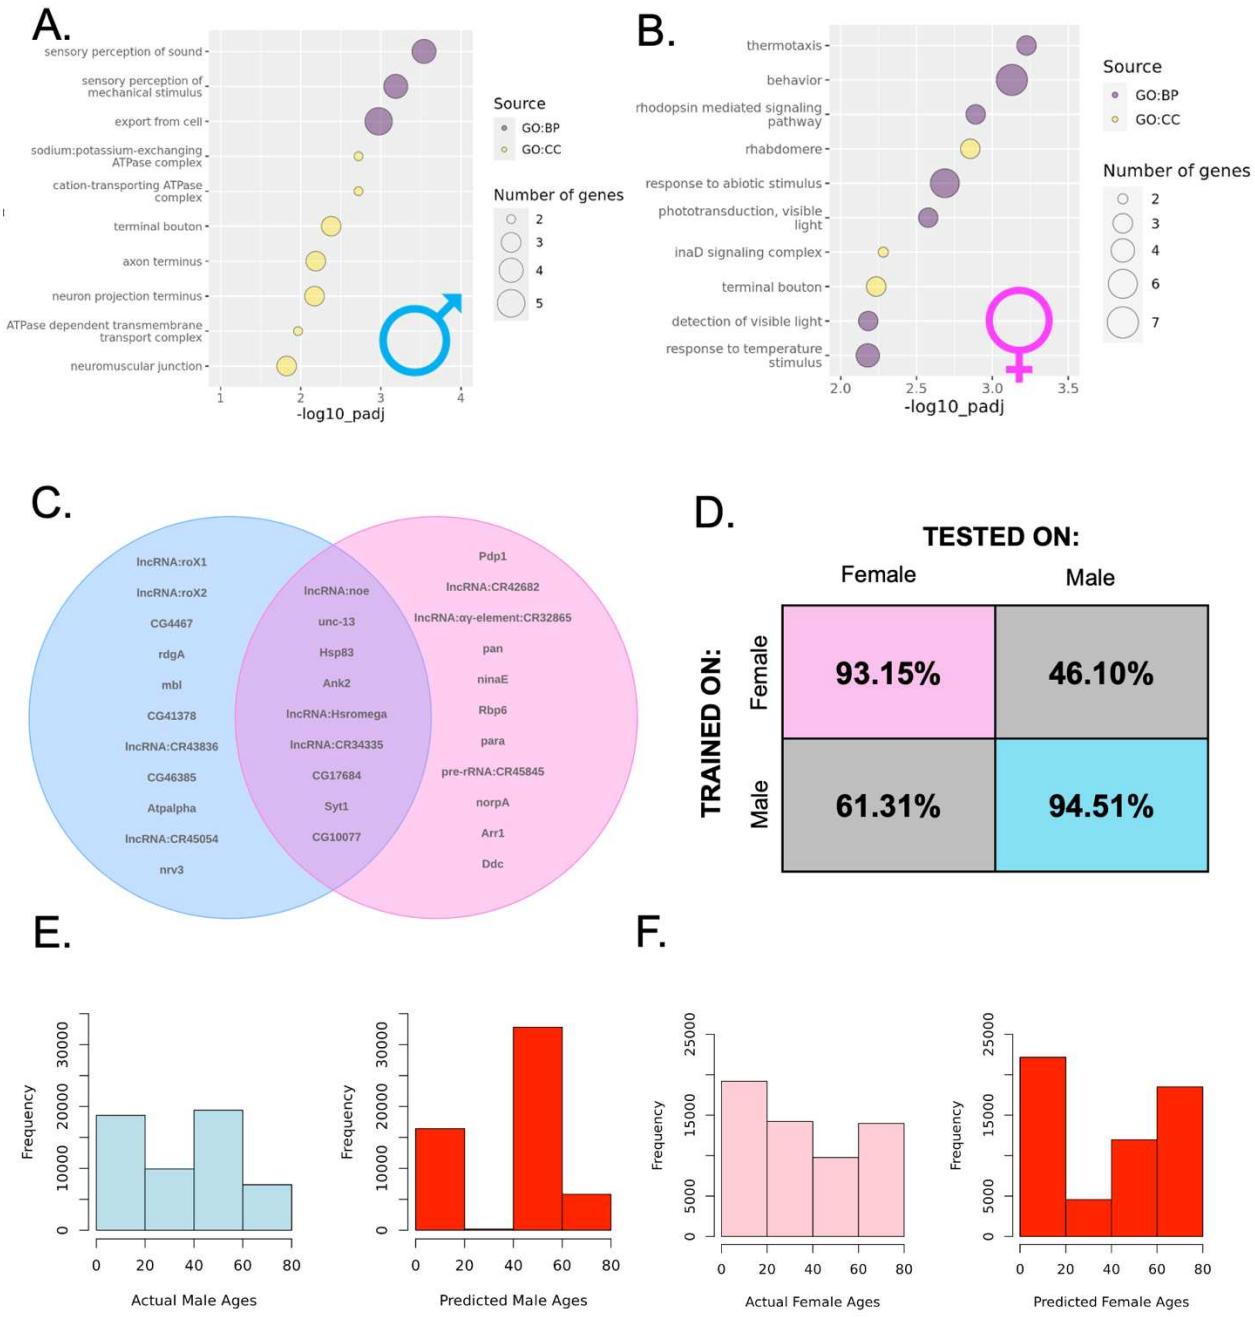

Figure 2. A) Gene set enrichment analysis for top 20 features in male-specific clock (included all cell types). **B)** Gene set enrichment analysis for top 20 features in female-specific clock. **C)** Venn Diagram of top 20 male features and top 20 female features. Blue indicates uniqueness to male-specific clock, pink indicates uniqueness to female-specific clock, and middle is the set of overlapping genes. **D)** F1 score of clocks that were either trained and tested on the one sex or trained on one sex and tested on the opposite sex**. E)** Distribution of ages in the male test set (left) and male ages predicted by the female-trained clock (right). **F)** Distribution of ages in the female test set (left) and female ages predicted by the male-trained clock (right).

roX1 knockout decreases survival probability in both sexes

There is a sizable body of literature on the role of *roX1* and *roX2* throughout embryonic and larval development in *Drosophila*. However, to our knowledge no literature exists addressing the role of the *roX* genes in aging. Considering the age-related hypothalamic upregulation of their mouse analog ncRNA, *Xist* [30] and the high feature importance of *roX1* and *roX2* in TimeFlies (Fig. 2b, d-e), we hypothesized that disrupting the *roX* RNAs would decrease lifespan in male flies. Mutations in the *roX1* gene alone do not show any developmental phenotypes, however, mutant male flies lacking both *roX1* and *roX2* do not survive past eclosion [39, 40]. Due to the early lethality of the double mutants, we performed a lifespan assay with *roX1* single-mutant flies. The $roX1^{SMC17A}$ genotype has been characterized as a null or severe hypomorph [39, 41] thus, we selected this line and compared it to a genetically matched common control.

We generated Kaplan-Meier survival curves (Fig 3) for each group and performed log-rank tests to determine whether the corresponding curves were significantly different. There was no difference in Kaplan-Meier survival probability between mutant males and mutant females. Moreover, there was no significant difference overall between *roX1* null males and control males. However, **when comparing only the time frames of days 3-40, there is a notably decreased survival probability in mutant males (X^2 = 30.2, p = 4e-08).** Fascinatingly, there is a **decreased survival probability throughout the entire lifespan in mutant females compared to control females (X^2 = 91.5, p < 2e-16).** (Fig 3). These results were unexpected, as females are the homogametic sex, do not undergo dosage compensation, and have a negligible expression of *roX1* [34]. However, the *yin* gene is located directly downstream of the *roX1* gene [42] and displays high expression in adult females and only low-to-moderate expression in adult males [34]. It is possible that the mutation of *roX1* has an impact on the *yin* gene and subsequently is detrimental to female survival. In summary, our lifespan assay validates the relevance of the top clock gene identified by TimeFlies.

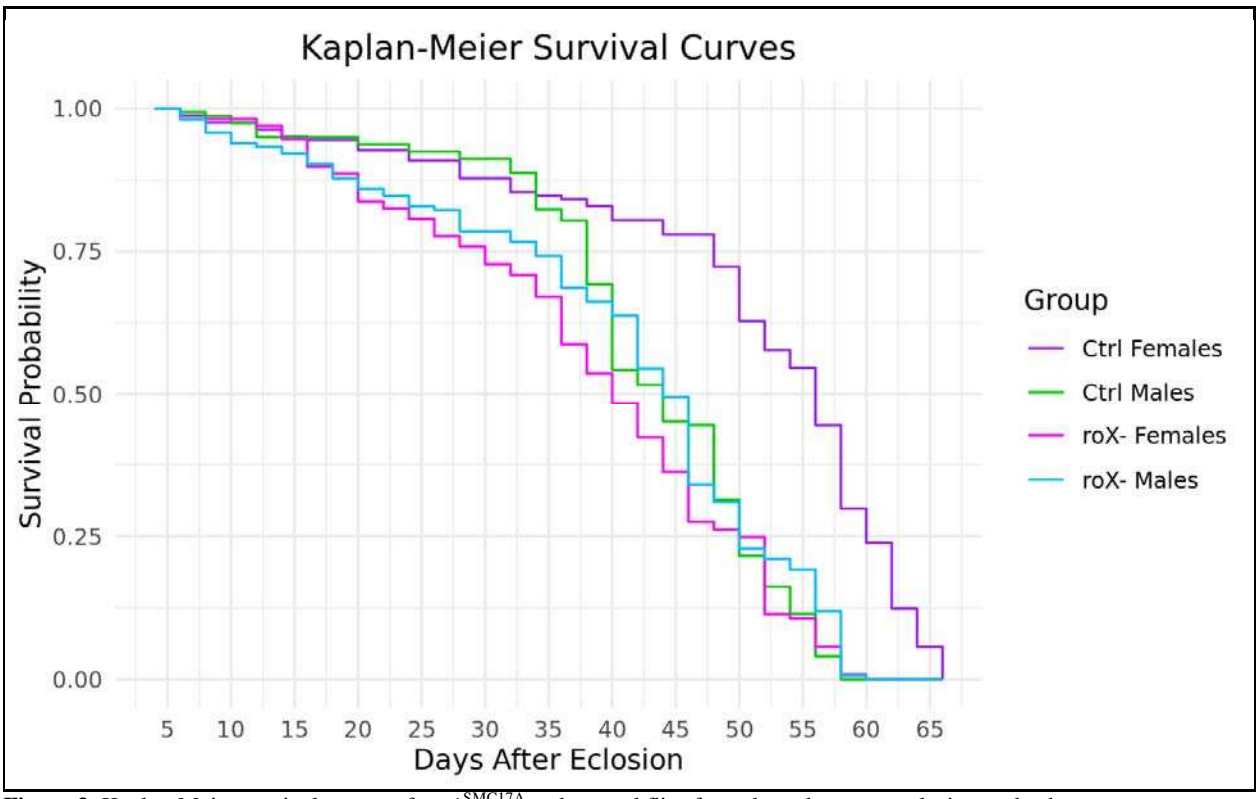

Figure 3. Kaplan-Meier survival curves of rox1^{SMC17A} and control flies from three days post-eclosion to death.

Discussion

We have developed an aging clock based on explainable deep neural networks that classifies *Drosophila melanogaster* head age at single-cell resolution with high accuracy. We have used the Aging Fly Cell Atlas [19], a diverse atlas of gene expression dynamics in hundreds of cell types in the fly head at four time points across the lifespan. Our clock, TimeFlies, requires no feature engineering prior to training, unlike its predecessors. Remarkably, it generalizes to all cell types despite the significant sparsity and high variability of single-cell RNA-seq data. Following the training and testing of our clock, we performed a thorough feature explanation using Shapley values to discover potential transcriptomic signatures of aging. We performed gene set enrichment analysis on the top 20 genes with the highest Shapley value magnitudes and found a significant enrichment of genes related to synaptic transmission. While synapse formation and function in the fruit fly have been thoroughly studied during development [43,44] how synapses change during normal and accelerated aging remains incompletely understood. lncRNA:*roX1* and lncRNA:*roX2*, noncoding RNAs on the X chromosome involved in the process of dosage compensation, were universal top features across all cell types, along with lncRNA:*noe*, which has unknown function and very limited associated literature. It is imperative to investigate lncRNA:*noe*, which may be relevant to the aging brain. Remarkably, long noncoding RNAs were enriched in TimeFlies feature explanations, consistent with previous studies that suggest an important role for lncRNAs in aging [35,36,37]. Analysis of our model reflects that lncRNA-

mediated gene regulation events may be significant to brain aging processes, calling for further investigation.

Further feature explainability analysis revealed noteworthy differences in female-specific and male-specific clocks, implying that aging is highly sex-specific, even at single-cell resolution. The female clock was unable to generalize to male test data, and the male clock was unable to generalize to the female data. There were also significant differences in gene set enrichment analysis following the feature explainability analysis of each clock. Notably, the explanation of the female-specific clock yielded many vision-related genes, suggesting a need to study sex differences in the aging eye in *Drosophila*. This finding is especially relevant given sex differences in age-associated eye diseases in humans [38].

The identification, via Shapley analysis, of the *roX* noncoding RNAs as top clock genes, is especially interesting due to a similar finding in the mouse brain. An snRNA-seq of the aging female mouse hypothalamus showed age-differential expression of *Xist*, the mouse analog of *roX* genes. Furthermore, corresponding clocks on that dataset identified *Xist* expression as a predictive factor in neuronal aging [30]. We set out to assess the role of dosage compensation in fruit fly aging and further validate the findings of TimeFlies. We performed a lifespan assay on a *roX1* null mutant line to gauge the effect of *roX1* on aging. *roX1* and *roX2*, another top TimeFlies gene, are thought to play a redundant role in dosage compensation, however, the knockout of both genes simultaneously leads to male lethality during early development, thus requiring us to first consider one gene individually. Observation of the *roX1* null population across lifespan showed that the knockout significantly reduced survival probability in males from day 3-40, with no significant difference relative to controls after day 40, and remarkably, a decreased survival rate for females throughout the entire lifespan. To fully understand the role of the roX RNAs in aging, we must develop a system in which both *roX1* and *roX2* are knocked out in a temporally controlled manner following eclosion.

To further test the robustness of TimeFlies, we will apply it to other single-cell aging and development atlases of *Drosophila melanogaster* as they are published, with specific interest in applying our clock to datasets of age-associated disease models such as the Alzheimer's Disease Fly Cell Atlas [45].

Methods

Dataset

The Aging Fly Cell Atlas (AFCA) [19] is a publicly available dataset documenting the single-cell transcriptomic profiles of fruit flies at ages 5, 30, 50, and 70 days. It includes both fly head samples and body samples. Here, we focus on the head data, with the objective of better understanding the aging fruit fly brain. The fly head dataset contains 289981 cells across 16 broad cell types and 15992 genes. The AFCA includes a near-equal distribution of the male and female samples, unlike many other published aging atlases. This allows for the investigation of sex-specific aging dynamics. The authors of the AFCA have published their own aging clocks on the dataset. However, these clocks are trained on very specific cellular subtypes and do not generalize to the whole dataset, nor do they specifically address sex differences. The authors also performed feature interpretation on their clocks, but the analysis of these features was limited to ribosomal protein-coding genes, with very limited discussion of other relevant genes [19]. This leaves the door open for a new clock that generalizes across all cell types for the AFCA dataset and a comprehensive analysis of sex-differential transcriptomic patterns in aging, which are especially prevalent in the brain.

To ensure that our model was learning genuine biological signals rather than batch effects, we generated a batch-corrected dataset to test against and found no significant differences (Supplementary Figure 5).

Model development and interpretation

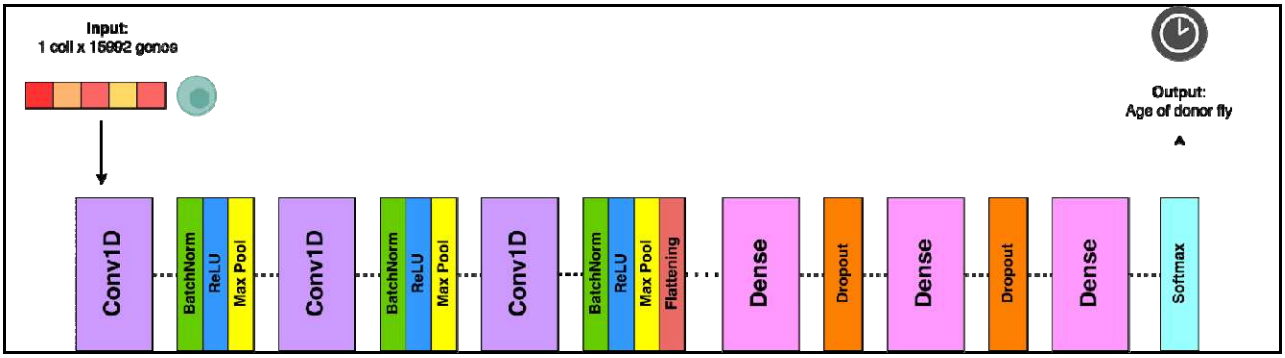

Figure 4. A detailed model architecture of TimeFlies framework. TimeFlies consists of a stack of three 1D convolution layers (separated by batch normalization, nonlinear activation, and max pooling), a flattening operation, and a stack of three dense layers with dropout layers in between. Softmax activation is applied to the output, which is the age of the donor fly.

The input data is initially a sparse matrix, meaning a matrix format that does not explicitly store zero-valued data to conserve space and memory. Specifically, it is in a coordinate format (COO), consisting of the coordinates of the non-zero values. Prior to providing this data as input to TimeFlies, we converted this matrix to a dense matrix using the numpy [46] and scipy [47] libraries in Python.

The TimeFlies aging clock (detailed in Figure 4) utilizes a convolutional neural network (CNN) consisting of three 1D convolution blocks, pooling/flattening operations, and two fully connected layers with dropout layers in between. CNNs have been used for genomic applications to predict regulatory activity from sequential genomic data like DNA sequences, etc. [48, 49, 50, 51]. We selected a CNN-based model for TimeFlies architecture due to its comparatively high performance and efficiency (Figure 5).

TimeFlies is implemented in Python with the Tensorflow library [52]. We do not perform any feature selection of genes and input the transcriptome-wide gene expression profile. An input sample is a vector of gene expression for a single cell. All the samples are split into training, validation, and test sets of 80%, 10%, and 10%, respectively. Due to the high dimensionality of the dataset, GPU acceleration was used to speed up training time for TimeFlies and its benchmark models. Feature explanation of TimeFlies was performed by obtaining Shapley values from GradientExplainer [53] and observing the features ranked highest. Gene set enrichment analysis was performed in R with g:Profiler [54].

Baseline evaluation

Performance of the 1D CNN framework was benchmarked against ElasticNet logistic regression and RandomForest, which were implemented with the scikit-learn library [55], XGBoost, which was implemented with the xgboost library [56], and a simple multilayer perceptron (MLP) neural network which was also implemented with Tensorflow [52]. Accuracy, F1 Score, and area under the curve (AUC) were used as evaluation metrics for the multiclass classification task. The 1D CNN outperformed every model except MLP (Fig 5). It exhibits on-par performance with the MLP model (no significant difference) despite increased computational efficiency due to a reduced number of trainable parameters. The features—genes—in the AFCA dataset were originally organized solely in alphabetical order without obvious spatial significance. However, shuffling the gene order with several different seeds has no significant impact on the TimeFlies performance (Supplementary Table 1) or model interpretation. Hence, we chose to proceed with the 1D CNN framework for the final architecture of TimeFlies.

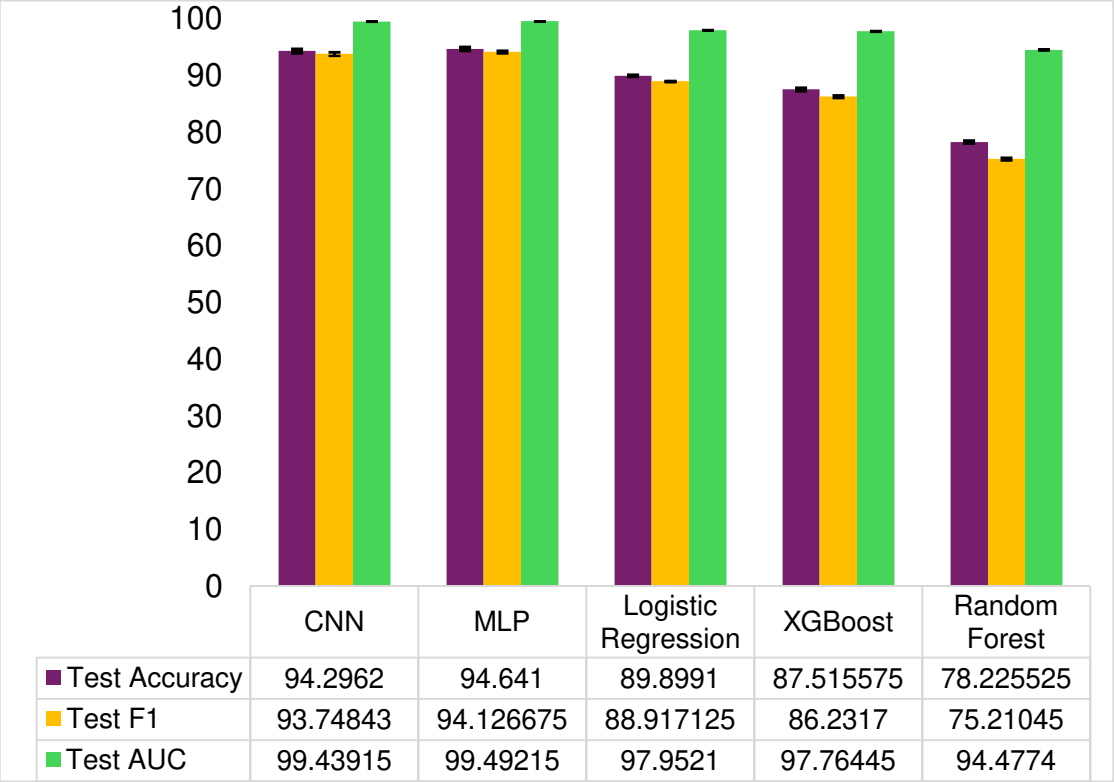

	CNN	MLP	Logistic Regression	XGBoost	Random Forest
Test Accuracy	94.2962	94.641	89.8991	87.515575	78.225525
Test F1	93.74843	94.126675	88.917125	86.2317	75.21045
Test AUC	99.43915	99.49215	97.9521	97.76445	94.4774

Figure 5. A bar chart of accuracy, area under the curve (AUC) and F1 score of all tested models when tasked with classifying donor age (5-day, 30-day, 50-day, or 70-day) from the cell's gene expression profile. Each model was run with 5 different seeds, average metrics are plotted with error bars representing standard deviation.

Lifespan assay

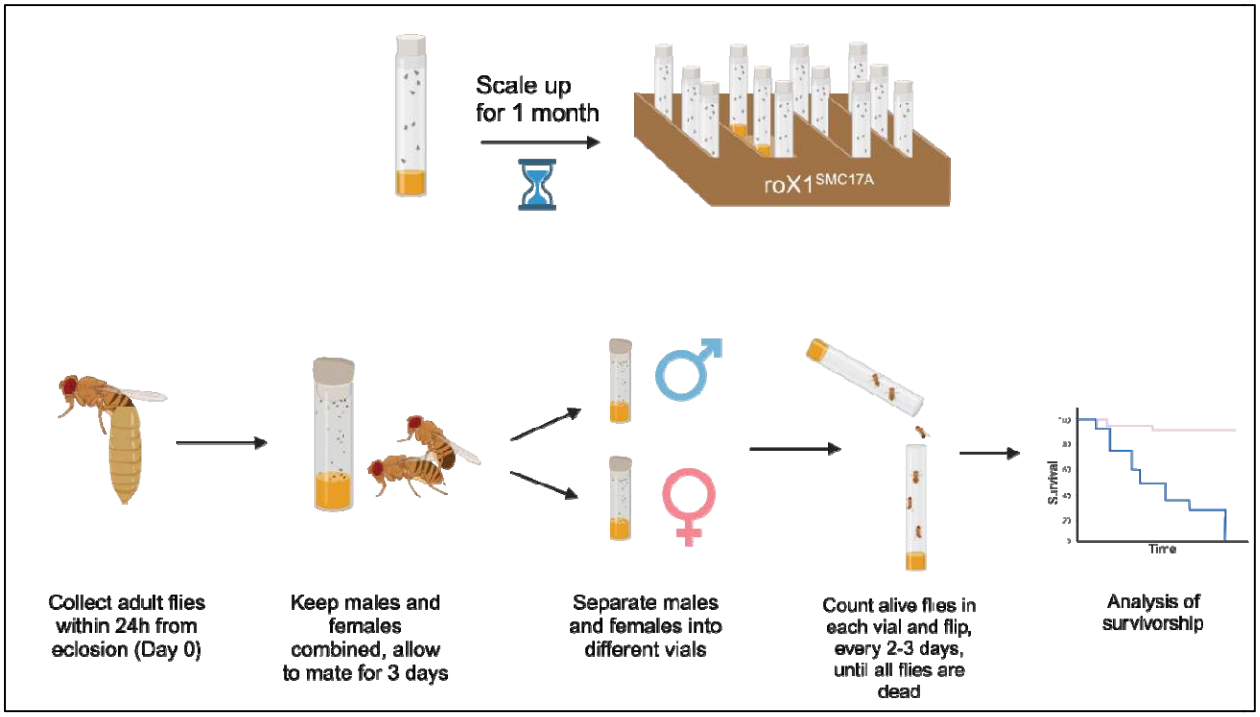

Figure 6. A schematic of the lifespan assay. Knockout and control flies were obtained and scaled up to a sizable population. Subsequently, a sex-specific lifespan assay was conducted.

roX1^{SMC17A} flies were gifted by Victoria Meller's lab (Wayne State University, Detroit, MI, USA) [41]. Flies were stored in vials with standard corn agar media at 25 C. The fly population was scaled up for one month. To validate that the received roX1^{SMC17A} line showed the proper marker, standard PCR and gel electrophoresis were done to check for the presence of the *LacZ* gene. The *w1118* fly line was used as a white-eyed control (RRID:BDSC_3605). To initiate the lifespan assay, 20-25 male flies and 20-25 female flies were collected within 24 hours following eclosion, placed in one vial, and allowed to mate for three days. Several such vials were prepared. On Day 3, flies were separated by sex and placed in different vials. Every 2-3 days, the number of alive and dead flies in each vial was recorded, and living flies were flipped into a new vial (Figure 6). Escaped flies and flies that faced death for non-natural reasons (i.e., crushed by the vial cap or stuck to food) were censored from the analysis. A total of 166 mutant males, 168 mutant females, 160 control males, and 165 control females were used in the experiment. A Kaplan-Meier Survival Analysis [57] was performed at the end of the experiment to gauge the potential effects of *roX1* knockout on lifespan.

Author Contributions

NT developed the TimeFlies aging clock model and the baseline models for evaluation. NT performed all iterations of the model runs. NT uploaded and maintained all code on GitHub. AP conducted feature explainability analysis and further downstream analysis including gene set enrichment. AP performed the lifespan assay and subsequent Kaplan-Meier analysis. AP wrote the manuscript and created the figures with aid from all other authors. KOCG advised experimental design and feature explainability analysis throughout the project tenure. EL and RS supervised the project and assisted with the experimental design, analysis of results, and manuscript writing.

Acknowledgments

This study was made possible by NSF Biology Integration Institute for Sex, Genomics, Aging, and Evolution (NSF DBI 2213824) and the NIH Institutional Training Program for Interactionist Cognitive Neuroscience (T32-MH115895).

Data Availability

The Aging Fly Cell Atlas is accessible at https://hongjielilab.shinyapps.io/AFCA/. It is downloadable in h5ad file format.

Code Availability

All code is available on GitHub at https://github.com/rsinghlab/TimeFlies.

References

1. López-Otín C, Blasco MA, Partridge L, Serrano M, Kroemer G. Hallmarks of aging: An expanding universe. *Cell*. 2023;186(2):243-278. doi:10.1016/j.cell.2022.11.001
2. Hannum G, Guinney J, Zhao L, et al. Genome-wide Methylation Profiles Reveal Quantitative Views of Human Aging Rates. *Molecular cell*. 2012;49(2):359. doi:10.1016/j.molcel.2012.10.016
3. Horvath S. DNA methylation age of human tissues and cell types. *Genome Biology*. 2013;14(10):3156. doi:10.1186/gb-2013-14-10-r115
4. Levine ME, Lu AT, Quach A, et al. An epigenetic biomarker of aging for lifespan and healthspan. *Aging (Albany NY)*. 2018;10(4):573-591. doi:10.18632/aging.101414
5. Lu AT, Quach A, Wilson JG, et al. DNA methylation GrimAge strongly predicts lifespan and healthspan. *Aging (Albany NY)*. 2019;11(2):303-327. doi:10.18632/aging.101684
6. Lu AT, Fei Z, Haghani A, et al. Universal DNA methylation age across mammalian tissues. *Nat Aging*. 2023;3(9):1144-1166. doi:10.1038/s43587-023-00462-6
7. Vanyushin BF, Nemirovsky LE, Klimenko VV, Vasiliev VK, Belozersky AN. The 5-Methylcytosine in DNA of Rats: Tissue and Age Specificity and the Changes Induced by

Hydrocortisone and other Agents. *Gerontologia*. 2009;19(3):138-152. doi:10.1159/000211967

8. Wilson VL, Smith RA, Ma S, Cutler RG. Genomic 5-methyldeoxycytidine decreases with age. *J Biol Chem*. 1987;262(21):9948-9951.
9. Romanov GA, Vanyushin BF. Methylation of reiterated sequences in mammalian DNAs. Effects of the tissue type, age, malignancy and hormonal induction. *Biochim Biophys Acta*. 1981;653(2):204-218. doi:10.1016/0005-2787(81)90156-8
10. Christensen BC, Houseman EA, Marsit CJ, et al. Aging and Environmental Exposures Alter Tissue-Specific DNA Methylation Dependent upon CpG Island Context. *PLOS Genetics*. 2009;5(8):e1000602. doi:10.1371/journal.pgen.1000602
11. Rutledge J, Oh H, Wyss-Coray T. Measuring biological age using omics data. *Nat Rev Genet*. 2022;23(12):715-727. doi:10.1038/s41576-022-00511-7
12. Peters MJ, Joehanes R, Pilling LC, et al. The transcriptional landscape of age in human peripheral blood. *Nat Commun*. 2015;6(1):8570. doi:10.1038/ncomms9570
13. Fleischer JG, Schulte R, Tsai HH, et al. Predicting age from the transcriptome of human dermal fibroblasts. *Genome Biology*. 2018;19:221. doi:10.1186/s13059-018-1599-6
14. Meyer DH, Schumacher B. BiT age: A transcriptome□based aging clock near the theoretical limit of accuracy. *Aging Cell*. 2021;20(3):e13320. doi:10.1111/acel.13320
15. de Lima Camillo LP, Lapierre LR, Singh R. A pan-tissue DNA-methylation epigenetic clock based on deep learning. *npj Aging*. 2022;8(1):1-15. doi:10.1038/s41514-022-00085-y
16. Holzscheck N, Falckenhayn C, Söhle J, et al. Modeling transcriptomic age using knowledge-primed artificial neural networks. *npj Aging Mech Dis*. 2021;7(1):1-13. doi:10.1038/s41514-021-00068-5
17. Almanzar N, Antony J, Baghel AS, et al. A single-cell transcriptomic atlas characterizes ageing tissues in the mouse. *Nature*. 2020;583(7817):590-595. doi:10.1038/s41586-020-2496-1
18. Gao SM, Qi Y, Zhang Q, et al. Aging atlas reveals cell-type-specific effects of pro-longevity strategies. *Nat Aging*. 2024;4(7):998-1013. doi:10.1038/s43587-024-00631-1
19. Lu TC, Brbić M, Park YJ, et al. Aging Fly Cell Atlas identifies exhaustive aging features at cellular resolution. *Science*. 2023;380(6650):eadg0934. doi:10.1126/science.adg0934
20. Yu D, Li M, Linghu G, et al. CellBiAge: Improved single-cell age classification using data binarization. *Cell Reports*. 2023;42(12):113500. doi:10.1016/j.celrep.2023.113500
21. Mao S, Su J, Wang L, Bo X, Li C, Chen H. A transcriptome-based single-cell biological age model and resource for tissue-specific aging measures. doi:10.1101/gr.277491.122
22. Bronikowski AM, Meisel RP, Biga PR, et al. Sex-specific aging in animals: Perspective and future directions. *Aging Cell*. 2022;21(2):e13542. doi:10.1111/acel.13542
23. Dorkenwald S, Matsliah A, Sterling AR, et al. Neuronal wiring diagram of an adult brain. *Nature*. 2024;634(8032):124-138. doi:10.1038/s41586-024-07558-y

24. Schlegel P, Yin Y, Bates AS, et al. Whole-brain annotation and multi-connectome cell typing of Drosophila. *Nature*. 2024;634(8032):139-152. doi:10.1038/s41586-024-07686-5

25. Paro PDR, Grossniklaus PDU, Santoro DR, Wutz PDA. Dosage Compensation Systems. In: *Introduction to Epigenetics [Internet]*. Springer; 2021. doi:10.1007/978-3-030-68670-3_4

26. Franke A, Baker BS. The rox1 and rox2 RNAs Are Essential Components of the Compensasome, which Mediates Dosage Compensation in Drosophila. *Molecular Cell*. 1999;4(1):117-122. doi:10.1016/S1097-2765(00)80193-8

27. Lucchesi JC, Kuroda MI. Dosage Compensation in Drosophila. *Cold Spring Harb Perspect Biol*. 2015;7(5):a019398. doi:10.1101/cshperspect.a019398

28. Meller VH, Rattner BP. The roX genes encode redundant male-specific lethal transcripts required for targeting of the MSL complex. *The EMBO Journal*. 2002;21(5):1084. doi:10.1093/emboj/21.5.1084

29. Lentini A, Cheng H, Noble JC, et al. Elastic dosage compensation by X-chromosome upregulation. *Nat Commun*. 2022;13(1):1854. doi:10.1038/s41467-022-29414-1

30. Hajdarovic KH, Yu D, Hassell LA, et al. Single-cell analysis of the aging female mouse hypothalamus. *Nat Aging*. 2022;2(7):662-678. doi:10.1038/s43587-022-00246-4

31. Kim B, Shortridge RD, Seong C, Oh Y, Baek K, Yoon J. Molecular Characterization of a Novel *Drosophila* Gene Which Is Expressed in the Central Nervous System. *Molecules and Cells*. 1998;8(6):750-757. doi:10.1016/S1016-8478(23)13493-5

32. Perez G, Barber GP, Benet-Pages A, et al. The UCSC Genome Browser database: 2025 update. *Nucleic Acids Res*. Published online October 26, 2024:gkae974. doi:10.1093/nar/gkae974

33. Johnson K, Knust E, Skaer H. *bloated tubules (blot)* Encodes a *Drosophila* Member of the Neurotransmitter Transporter Family Required for Organisation of the Apical Cytocortex. *Developmental Biology*. 1999;212(2):440-454. doi:10.1006/dbio.1999.9351

34. Brown JB, Boley N, Eisman R, et al. Diversity and dynamics of the Drosophila transcriptome. *Nature*. 2014;512(7515):393-399. doi:10.1038/nature12962

35. Grammatikakis I, Panda AC, Abdelmohsen K, Gorospe M. Long noncoding RNAs (lncRNAs) and the molecular hallmarks of aging. *Aging (Albany NY)*. 2014;6(12):992. doi:10.18632/aging.100710

36. Cai D, Han JDJ. Aging-associated lncRNAs are evolutionarily conserved and participate in NFκB signaling. *Nat Aging*. 2021;1(5):438-453. doi:10.1038/s43587-021-00056-0

37. Yang D, Lian T, Tu J, et al. LncRNA mediated regulation of aging pathways in Drosophila melanogaster during dietary restriction. *Aging (Albany NY)*. 2016;8(9):2182. doi:10.18632/aging.101062

38. Aninye IO, Digre K, Hartnett ME, et al. The roles of sex and gender in women's eye health disparities in the United States. *Biology of Sex Differences*. 2021;12(1):57. doi:10.1186/s13293-021-00401-3

39. Deng X, Rattner BP, Souter S, Meller VH. The severity of *roX1* mutations is predicted by MSL localization on the X chromosome. *Mechanisms of Development*. 2005;122(10):1094-1105. doi:10.1016/j.mod.2005.06.004

40. Meller VH, Rattner BP. The roX genes encode redundant male-specific lethal transcripts required for targeting of the MSL complex. *The EMBO Journal*. 2002;21(5):1084. doi:10.1093/emboj/21.5.1084

41. Deng X, Meller V. Molecularly Severe roX1 Mutations Contribute to Dosage Compensation in Drosophila. *Genesis (New York, NY□: 2000)*. 2009;47:49-54. doi:10.1002/dvg.20463

42. dos Santos G, Schroeder AJ, Goodman JL, et al. FlyBase: introduction of the Drosophila melanogaster Release 6 reference genome assembly and large-scale migration of genome annotations. *Nucleic Acids Research*. 2014;43(Database issue):D690. doi:10.1093/nar/gku1099

43. Frank CA, James TD, Müller M. Homeostatic control of Drosophila neuromuscular junction function. *Synapse*. 2020;74(1):e22133. doi:10.1002/syn.22133

44. Harris KP, Littleton JT. Transmission, Development, and Plasticity of Synapses. *Genetics*. 2015;201(2):345-375. doi:10.1534/genetics.115.176529

45. Park YJ, Lu TC, Jackson T, et al. Whole organism snRNA-seq reveals systemic peripheral changes in Alzheimer's Disease fly models. Published online March 13, 2024:2024.03.10.584317. doi:10.1101/2024.03.10.584317

46. Harris CR, Millman KJ, van der Walt SJ, et al. Array programming with NumPy. *Nature*. 2020;585(7825):357-362. doi:10.1038/s41586-020-2649-2

47. Virtanen P, Gommers R, Oliphant TE, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat Methods*. 2020;17(3):261-272. doi:10.1038/s41592-019-0686-2

48. Kelley DR, Snoek J, Rinn JL. Basset: learning the regulatory code of the accessible genome with deep convolutional neural networks. *Genome Res*. 2016;26(7):990-999. doi:10.1101/gr.200535.115

49. Singh R, Lanchantin J, Robins G, Qi Y. DeepChrome: deep-learning for predicting gene expression from histone modifications. *Bioinformatics*. 2016;32(17):i639-i648. doi:10.1093/bioinformatics/btw427

50. Kelley DR, Reshef YA, Bileschi M, Belanger D, McLean CY, Snoek J. Sequential regulatory activity prediction across chromosomes with convolutional neural networks. *Genome Research*. 2018;28(5):739. doi:10.1101/gr.227819.117

51. Alipanahi B, Delong A, Weirauch MT, Frey BJ. Predicting the sequence specificities of DNA- and RNA-binding proteins by deep learning. *Nat Biotechnol*. 2015;33(8):831-838. doi:10.1038/nbt.3300

52. Abadi M, Agarwal A, Barham P, et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. Published online March 16, 2016. doi:10.48550/arXiv.1603.04467

53. Lundberg S, Lee SI. A Unified Approach to Interpreting Model Predictions. Published online November 25, 2017. doi:10.48550/arXiv.1705.07874

54. Kolberg L, Raudvere U, Kuzmin I, Adler P, Vilo J, Peterson H. g:Profiler—interoperable web service for functional enrichment analysis and gene identifier mapping (2023 update). *Nucleic Acids Research*. 2023;51(W1):W207-W212. doi:10.1093/nar/gkad347

55. Buitinck L, Louppe G, Blondel M, et al. API design for machine learning software: experiences from the scikit-learn project. Published online September 1, 2013. doi:10.48550/arXiv.1309.0238

56. Chen T, Guestrin C. XGBoost: A Scalable Tree Boosting System. Published online June 10, 2016. doi:10.48550/arXiv.1603.02754

57. Kaplan EL, Meier P. Nonparametric Estimation from Incomplete Observations. *Journal of the American Statistical Association*. 1958;53(282):457-481. doi:10.1080/01621459.1958.10501452