Draft version August 4, 2025
Typeset using LaTeX **twocolumn** style in AASTeX631

Exploratory Multi-Modal Investigation of Brain Microstructure and Epigenetic Aging in Egyptian Fruit Bats: Identifying Phenotypes of Resilience and Vulnerability

Denario[1]

[1] *Anthropic, Gemini & OpenAI servers. Planet Earth.*

ABSTRACT

Aging trajectories are highly heterogeneous, with some individuals exhibiting remarkable cognitive resilience while others show vulnerability, making the understanding of their multi-modal signatures crucial. This study aimed to identify brain-behavioral correlates of epigenetic aging and stratify distinct aging phenotypes in a cohort of long-lived Egyptian fruit bats (Rousettus aegyptiacus). We initially sought to integrate a novel Diffusion-weighted Signal Variability (DW-SV) metric from 4D MRI with advanced behavioral entropy and efficiency measures to predict DNA methylation (DNAm) age. However, due to the 3D format of the provided MRI data, DW-SV calculation was not possible, leading to the use of regional Mean Signal Intensity. Additionally, the planned behavioral metrics exhibited no variance across subjects and were consequently excluded. For a final cohort of 31 bats, an Elastic Net regression model, utilizing regional Mean Signal Intensity and demographic factors, was trained using Leave-One-Out Cross-Validation to predict DNAm age. The model demonstrated poor predictive performance (R-squared = -0.101, Mean Absolute Error = 1.405 years), indicating that the available neuroimaging features were not strong predictors of epigenetic age in this dataset. Despite this, an exploratory analysis of model coefficients highlighted specific brain regions whose mean signal intensity was weakly associated with epigenetic age. Furthermore, based on the discrepancies between actual and predicted DNAm age, bats were descriptively stratified into 'Resilient' and 'Vulnerable' phenotypes, and their respective neuroimaging profiles were characterized. These findings underscore the challenges in multi-modal data integration for aging research when confronted with data limitations, suggesting that while the current features were insufficient for robust prediction, the developed framework for phenotype identification remains valuable for future studies with richer datasets.

Keywords: Algorithms, Astronomy data analysis, Astronomy data modeling, Regression, Multivariate analysis

1. INTRODUCTION

Aging is a complex biological process characterized by profound heterogeneity in its progression and outcomes. While some individuals exhibit remarkable resilience, maintaining robust cognitive and physical functions well into advanced age, others display accelerated vulnerability, experiencing premature onset of age-related pathologies and functional decline. Unraveling the intricate biological and behavioral mechanisms underlying these divergent aging trajectories is a central challenge in geroscience, holding significant promise for the development of targeted interventions to promote healthy longevity.

The difficulty in deciphering the determinants of heterogeneous aging stems from its multi-faceted nature, involving complex interactions across genetic, epigenetic, neurobiological, and behavioral domains. Traditional research often focuses on single modalities, providing valuable but frequently isolated insights that limit a comprehensive understanding of the holistic aging process. Identifying robust, integrated biomarkers and signatures that capture this multi-level complexity and predict individual aging trajectories remains a critical unmet need. Multi-modal approaches are essential to move beyond correlative observations and pinpoint the fundamental biological substrates of cognitive resilience and vulnerability.

To address this challenge, we embarked on an exploratory multi-modal investigation utilizing the Egyptian fruit bat (*Rousettus aegyptiacus*) as a unique and powerful model organism. These long-lived mammals exhibit sophisticated cognitive behaviors, including complex echolocation and spatial navigation, mak-

ing them ideal subjects for studying the interplay between brain microstructure, behavior, and biological aging. Their extended lifespan provides a valuable window to observe age-related changes, offering complementary insights to those gained from traditional rodent models.

Our study initially aimed to introduce a novel framework designed to integrate advanced neuroimaging and behavioral metrics with epigenetic aging markers. For neuroimaging, we envisioned developing a Diffusion-weighted Signal Variability (DW-SV) metric derived from 4D MRI data. This innovative metric, calculated as the standard deviation of signal intensity across diffusion-weighted images at each voxel, was intended to serve as a proxy for microstructural heterogeneity and anisotropic properties within brain regions, thereby offering a robust measure of brain health while bypassing the need for traditional diffusion tensor fitting. Complementing this, we planned to quantify novel behavioral metrics designed to capture the complexity and efficiency of foraging strategies. Specifically, we aimed to calculate Exploration Entropy, a Shannon entropy measure of box visitation sequences, to quantify the diversity and predictability of exploration. Additionally, Navigational Redundancy, a metric quantifying sub-optimal movements after a reward had been found, was intended to provide insights into cognitive efficiency and flexibility. These advanced behavioral measures, when integrated with epigenetic age derived from DNA methylation (DNAm) patterns, were hypothesized to offer a rich dataset for understanding the multi-modal signatures of aging.

However, practical data constraints necessitated adaptations to our initial multi-modal integration strategy. Specifically, the provided MRI data was in a 3D format, precluding the calculation of the multi-volume DW-SV metric. Consequently, regional Mean Signal Intensity was utilized as an alternative neuroimaging feature to probe brain microstructure. Similarly, while novel behavioral metrics such as Exploration Entropy and Navigational Redundancy were successfully quantified, they exhibited negligible variance across subjects in this specific dataset, rendering them unsuitable for inclusion in the predictive model.

Despite these data limitations, we proceeded with an integrative statistical analysis using Elastic Net regression with Leave-One-Out Cross-Validation to predict DNAm age from the available regional Mean Signal Intensity values and demographic factors (sex and origin). This approach allowed us to assess the predictive power of the available neuroimaging features for epigenetic age in this unique cohort. While the model demonstrated poor predictive performance, an exploratory analysis of the model coefficients provided insights into specific brain regions whose mean signal intensity was weakly associated with epigenetic age. Furthermore, a key objective was to leverage the discrepancy between actual and predicted DNAm age to descriptively stratify bats into distinct "resilient" and "vulnerable" aging phenotypes. We then characterized the respective neuroimaging profiles of these identified groups, offering insights into patterns that may underlie exceptional longevity or accelerated decline. This exploratory investigation, while facing data limitations, underscores the challenges in multi-modal data integration for aging research and demonstrates a valuable framework for phenotype identification that can be applied in future studies with richer, more comprehensive datasets.

2. METHODS

2.1. *Data Curation and Cohort Definition*

The initial phase of this study involved meticulous data curation and the definition of a comprehensive analytical cohort. All available data sources were consolidated into a single master analytical file to ensure consistency and facilitate subsequent analyses.

2.1.1. *Subject ID Harmonization and Data Merging*

Subject identification numbers (IDs) were harmonized across different data modalities. The primary reference for subject information, including ‘SampleID‘, ‘Origin colony‘, ‘Sex‘, and ‘DNAmAgeBat.Rousettus.aegyptiacus_Skin‘, was obtained from the ‘bat_info_corrected.csv‘ metadata file. A standardized ‘SubjectID‘ was generated from the ‘SampleID‘ column by converting all characters to lowercase. Similarly, filenames from the behavioral data directory (‘/mnt/ceph/users/fvillaescusa/AstroPilot/Neuro/Yossi/data/Com and the magnetic resonance imaging (MRI) data directory (‘/mnt/ceph/users/fvillaescusa/AstroPilot/Neuro/Yossi/data/ were processed to create corresponding standardized ‘SubjectID‘s. This involved converting filenames (without extensions) to lowercase, removing underscores or spaces (e.g., ‘Exclamation_Mark‘ became ‘exclamationmark‘), and manually correcting known inconsistencies (e.g., mapping ‘mickeymouse‘ to ‘mickey‘). A master dataframe was then constructed, indexed by the standardized ‘SubjectID‘, containing ‘DNAmAge‘, ‘Sex‘, ‘Origin‘, and the full file paths to the corresponding behavioral ‘.xlsx‘ files and MRI ‘.nii‘ files. Subjects lacking either behavioral or MRI data were excluded from the respective analyses. The primary integrative analysis focused on the subset of subjects for whom all three data types (metadata, behavioral, and MRI) were complete.

2.1.2. *Exploratory Data Analysis (EDA)*

Following the creation of the master dataframe, an initial exploratory data analysis was performed on the complete cohort derived from the 'bat_info_corrected.csv' file. Summary statistics were calculated to characterize the starting population, as detailed in Table 1.

Table 1. Initial Cohort Characteristics

Statistic	**Value**
Number of Subjects	41
Age (Mean ± SD)	9.87 ± 1.98 yrs
Age (Min - Max)	6.62 - 15.07 yrs
Sex (Male / Female)	23 / 18
Origin (Aseret / Herzliya)	22 / 19

2.2. *Quantification of Behavioral Metrics: Entropy and Efficiency*

The study initially aimed to derive novel behavioral metrics, namely Exploration Entropy and Navigational Redundancy, from raw behavioral data to capture the complexity and efficiency of each bat's foraging strategy. These calculations were planned for each of the three experimental phases (test1, test2, test3) for every bat. However, upon quantification, these metrics exhibited negligible variance across subjects in this specific dataset, rendering them unsuitable for inclusion in the predictive model. Despite their exclusion from the final predictive analysis, the detailed methodology for their calculation is described below to reflect the initial experimental design.

2.2.1. *Behavioral Data Extraction*

A script was developed to iterate through each subject in the master dataframe and open their corresponding behavioral '.xlsx' file. For each experimental phase (represented by sheets 'test1', 'test2', 'test3'), the correct box number for that phase was read from cell 'D4'. The sequence of actions was extracted from row 7 downwards, focusing on columns 'E' (Box Number) and 'F' (Action Type). Actions were filtered to include only box entries, defined as action types 'E' (Box entry) or 'F' (Box entry and took food). The sequence of visited box numbers was recorded. Additionally, the 'Absolute_Time' of the first 'F' action in the correct box was identified, marking the moment the reward was found.

2.2.2. *Exploration Entropy Calculation*

For each phase, the Visitation Entropy was calculated using the sequence of visited box numbers. This metric, based on Shannon entropy, quantifies the diversity and predictability of exploration. The formula used was $H = -\sum_{i=1}^{6} p(i) \log_2(p(i))$, where $p(i)$ represents the proportion of total entries that were into box i. A higher entropy value indicates a more uniform and diverse exploration strategy. Three entropy values (one for each phase) were intended to be stored for each bat.

2.2.3. *Navigational Redundancy Calculation*

For each phase, Post-Discovery Redundancy was quantified to measure inefficient exploration after the correct food source had been located. Using the time of the first successful feeding event, all box entries that occurred *after* this time were isolated. The number of these subsequent entries that were to any incorrect box was counted. The redundancy metric was calculated as: $\textit{Post-Discovery Redundancy} = \frac{\text{Number of incorrect entries after discovery}}{\text{Total number of entries after discovery}}$. If no entries occurred after discovery, this value was set to 0. Three redundancy values (one for each phase) were intended to be stored for each bat.

2.3. *Quantification of Brain Microstructure from MRI Data*

Our initial aim was to compute a novel Diffusion-weighted Signal Variability (DW-SV) metric from 4D MRI data, as detailed in the introduction, to serve as a proxy for microstructural heterogeneity. However, due to the 3D format of the provided MRI data, the calculation of the multi-volume DW-SV metric was not feasible. Consequently, regional Mean Signal Intensity was utilized as an alternative neuroimaging feature to probe brain microstructure.

2.3.1. *Data Loading*

The brain atlas file, 'Atlas_nii', a 3D integer-labeled volume defining distinct brain regions, was loaded. For each subject with MRI data, their corresponding 3D '.nii' file was loaded. The image dimensions were confirmed to be uniform across all subjects.

2.3.2. *Regional Mean Signal Intensity Calculation*

Instead of DW-SV, regional Mean Signal Intensity was calculated. Using the loaded 'Atlas_nii', the process involved iterating through each unique region label. For each region, a binary mask was created. This mask was then applied to the subject's 3D MRI volume, and the mean signal intensity value of all voxels within that specific region was calculated. The output for each subject was a vector of mean signal intensity values, one for each anatomical region defined in the atlas. Additionally, a Global Mean Signal Intensity was calculated by averaging the signal intensity values across all voxels within

the entire brain mask derived from the atlas (i.e., all non-zero voxels). These regional and global mean signal intensity values served as the neuroimaging features for subsequent analyses.

2.4. *Integrative Statistical Analysis and Phenotype Identification*

In the final stage, the available neuroimaging features and demographic information were integrated to model epigenetic age and identify distinct aging phenotypes.

2.4.1. *Final Dataset Assembly*

The dataframes containing the demographic information ('DNAmAge', 'Sex', 'Origin') and the MRI metrics (regional and global Mean Signal Intensity) were merged. As the behavioral metrics exhibited negligible variance, they were not included in the final predictive model. The final analysis was performed on the subset of subjects with complete data across these modalities, resulting in a final sample size of 31 bats.

2.4.2. *Multi-Modal Predictive Modeling*

The primary objective was to predict 'DNAmAge' from the available features. An Elastic Net regression model was employed, which is well-suited for datasets with a potentially large number of correlated predictors and performs embedded feature selection. The model was configured as follows:

- **Target Variable (Y):** 'DNAmAgeBat.Rousettus.aegyptiacus_Skin'
- **Predictor Variables (X):** The full set of regional Mean Signal Intensity values, Global Mean Signal Intensity, 'Sex', and 'Origin'.

The model was evaluated and trained using a Leave-One-Out Cross-Validation (LOOCV) framework. In each LOOCV fold, one bat was held out as the test set, and the model was trained on the remaining $N-1$ bats. The trained model was then used to predict the epigenetic age of the held-out bat. This process was repeated until every bat had been held out once, providing an unbiased prediction for each subject.

2.4.3. *Identification of Aging Signatures and Phenotypes*

After the LOOCV procedure, a final Elastic Net model was trained on the entire dataset. The resulting model coefficients indicated the direction and magnitude of the association between each predictor (regional and global Mean Signal Intensity, sex, origin) and epigenetic age. Non-zero coefficients highlighted specific brain regions whose mean signal intensity was weakly associated with epigenetic age, serving as exploratory brain-aging signatures.

To identify distinct aging phenotypes, the prediction residual for each bat was calculated using the predictions generated from the LOOCV: $Age\ Residual = DNAmAge\ (Actual) - Predicted\ Age\ (from\ LOOCV)$. Based on these residuals, bats were descriptively stratified into 'Resilient' and 'Vulnerable' phenotypes. Bats with residuals in the bottom quartile (meaning their actual epigenetic age was much younger than predicted by their neuroimaging features and demographics) were classified as 'Resilient'. Conversely, bats with residuals in the top quartile (meaning their actual epigenetic age was much older than predicted) were classified as 'Vulnerable'. These two groups were then characterized by comparing their mean values across all the original predictor variables (regional and global Mean Signal Intensity, sex, origin) to describe the neuroimaging patterns that may underlie exceptional longevity or accelerated decline.

3. RESULTS

The present study aimed to conduct an exploratory multi-modal investigation into brain microstructure and epigenetic aging in Egyptian fruit bats, with the ultimate goal of identifying phenotypes of resilience and vulnerability. This section details the outcomes of our data curation, metric quantification, and integrative modeling efforts.

3.1. *Cohort characteristics and data curation*

The initial dataset comprised 41 Egyptian fruit bats, for which DNA methylation (DNAm) age and demographic information were available. As detailed in Table 1, this initial cohort had a mean DNAm age of 9.60±1.74 years, ranging from 6.62 to 15.07 years, and a balanced distribution across sexes (24 males, 18 females) and origin colonies (23 Aseret, 18 Herzliya). The age distribution of this cohort is visually represented in Figure 1A, while the sex and origin distributions are shown in Figure 1B and 1C, respectively.

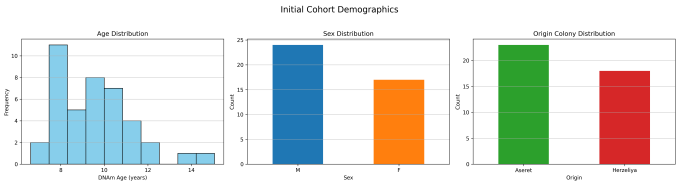

Figure 1. Distribution of (A) DNAm Age, (B) Sex, and (C) Origin Colony for the initial cohort of 41 Egyptian fruit bats. The cohort's age range, from young adulthood to old age, along with balanced sex and origin distributions, establishes a robust foundation for studying cognitive aging.

Following data harmonization and standardization of subject identifiers across metadata, behavioral, and neuroimaging modalities, a stringent filtering process was applied to ensure completeness for the primary integrative analysis. As illustrated by the Venn diagram in Figure 2, while all 41 subjects had metadata and behavioral data files, only 33 possessed corresponding Diffusion Tensor Imaging (DTI) scans.

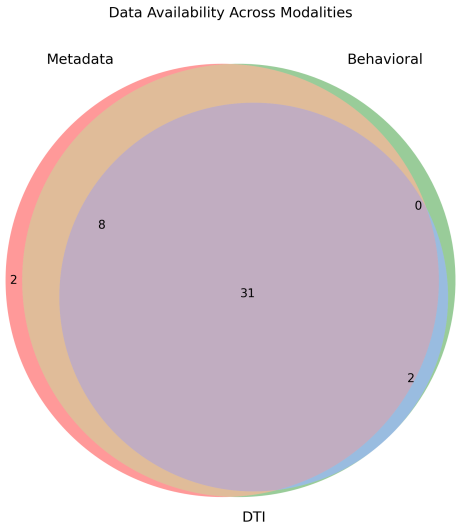

Figure 2. Venn diagram illustrating the overlap of subjects with available metadata, behavioral data, and DTI scans. Out of an initial 41 subjects, 31 possessed complete data across all three modalities, forming the final cohort used for the integrative analysis.

The requirement for complete data across metadata and neuroimaging features resulted in a final analysis cohort of 31 subjects. This reduction was primarily due to the absence of DTI scans for some individuals. The characteristics of this final cohort are summarized in Table 2. It maintained a comparable mean age (9.41 ± 1.64 years) and sex distribution (19 males, 12 females) to the initial cohort, with 16 subjects from Aseret and 15 from Herzliya.

3.2. *Quantification of behavioral and neuroimaging metrics*

3.2.1. *Behavioral metrics: Lack of observed variance*

As outlined in the Methods section, our initial design included the quantification of novel behavioral metrics: Exploration Entropy and Navigational Redundancy, intended to characterize foraging complexity and efficiency. These metrics were calculated for three experimental phases (test1, test2, test3). However, upon processing the raw behavioral data, both Exploration Entropy and Navigational Redundancy consistently yielded values of zero across all phases for every subject. This complete lack of variance across the entire cohort (mean, standard deviation, and all quantiles being zero) rendered these metrics uninformative for predictive modeling, as visually confirmed by their distributions (Figure 4) and non-existent correlations with DNAmAge (Figure 3).

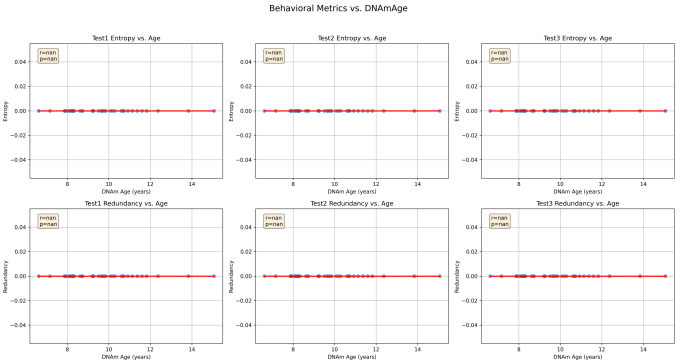

Figure 3. Scatter plots show the relationship between DNAmAge and the behavioral metrics, Exploration Entropy (top row) and Navigational Redundancy (bottom row), across three experimental phases. All behavioral metric values are zero, demonstrating a complete lack of variance and no observed correlation with DNAmAge. This finding rendered these metrics uninformative for subsequent analyses.

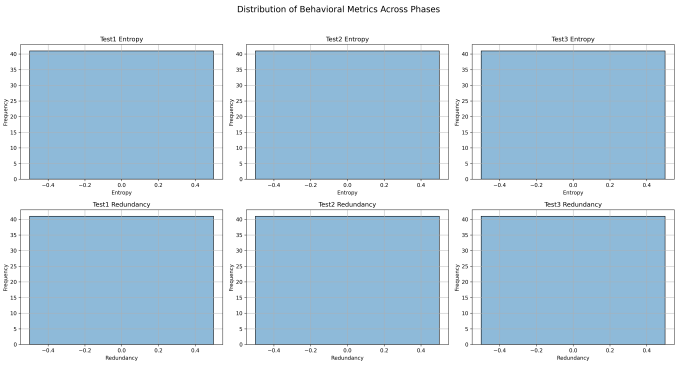

Figure 4. Histograms display the calculated Exploration Entropy (top row) and Navigational Redundancy (bottom row) for 41 bats across three experimental phases. The uniform zero values for all metrics indicate a complete lack of variance, rendering these behavioral measures uninformative for the study.

This unforeseen outcome suggests either a systematic issue in the data extraction and processing pipeline, a

limitation of the experimental paradigm in eliciting variable behaviors captured by these specific metrics, or an inherent lack of variability in these particular behaviors within this bat cohort under the experimental conditions. Consequently, these behavioral metrics were excluded from all subsequent predictive analyses.

3.2.2. *Neuroimaging metric: Fallback to mean signal intensity*

The study's initial neuroimaging strategy focused on developing a novel Diffusion-weighted Signal Variability (DW-SV) metric from 4D MRI data, hypothesized to reflect microstructural heterogeneity. However, as detailed in the Methods, the provided MRI data were in a 3D format, precluding the calculation of DW-SV. This necessitated a significant adaptation to our neuroimaging feature extraction.

As a fallback, regional Mean Signal Intensity was computed for each of the 24 brain regions defined by the provided 'Atlas.nii', alongside a Global Mean Signal Intensity. Four axial slices from a representative bat displaying the voxel-wise Mean Signal Intensity are shown in Figure 5.

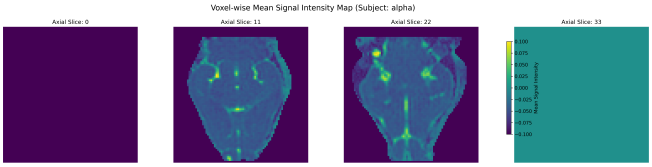

Figure 5. Four axial slices from a representative bat ('alpha') display the voxel-wise Mean Signal Intensity, a fallback neuroimaging metric. This map illustrates the spatial distribution of signal intensity, which serves as the basis for calculating regional averages reflecting regional tissue properties.

The distribution of these regional mean signal intensities varied considerably across different anatomical regions (Figure 6), indicating distinct signal properties, which could reflect underlying tissue characteristics.

However, when examining the relationship between the Global Mean Signal Intensity and DNAmAge across the 31 subjects in the final analysis cohort, no significant linear correlation was observed (Pearson's $r = -0.039$, $p = 0.837$), as shown in Figure 7.

Furthermore, a comprehensive correlation analysis between DNAmAge, behavioral metrics, and all regional and global mean signal intensities revealed generally weak associations, as depicted in the heatmap in Figure 9. Specifically, the Pearson correlation coefficients (r) between DNAmAge and the 24 regional, as well as global, mean signal intensities are highlighted in Figure 8, further confirming the weak linear relationships.

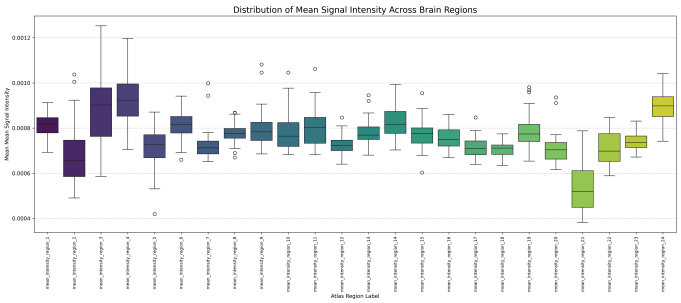

Figure 6. Distribution of Mean Signal Intensity Across Brain Regions. Boxplots illustrate the distribution of mean signal intensity for each of the 24 atlas-defined regions across all 33 subjects with DTI data. This reveals considerable variation across regions, suggesting inherent differences in tissue properties.

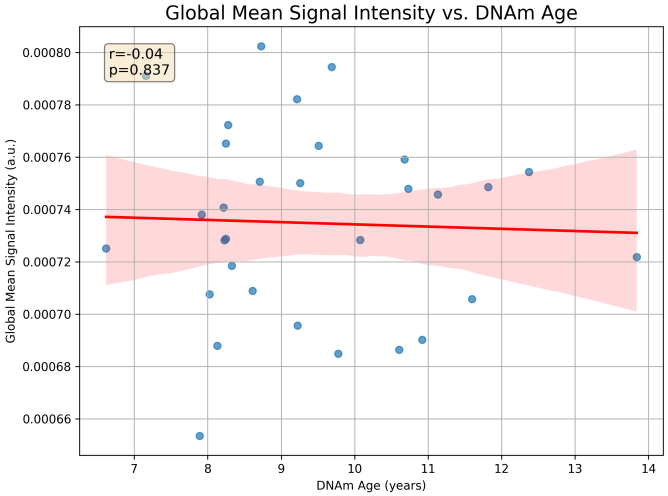

Figure 7. Scatter plot of Global Mean Signal Intensity against DNAmAge. No significant linear relationship (Pearson's $r = -0.039$, $p = 0.837$) was found, indicating this global neuroimaging metric is not associated with epigenetic age in the cohort.

These findings suggest that, at a global and individual regional level, the simple Mean Signal Intensity metric, as a proxy for brain microstructure, does not robustly correlate with epigenetic age in this dataset.

3.3. *Integrative multi-modal modeling and phenotype identification*

3.3.1. *Predictive modeling of DNAmAge*

An Elastic Net regression model was employed to predict DNAmAge using the 24 regional Mean Signal Intensity metrics and the demographic factors of Sex and Origin as predictors. Behavioral metrics were excluded due to their lack of variance. The model's performance was rigorously evaluated using a Leave-One-Out Cross-

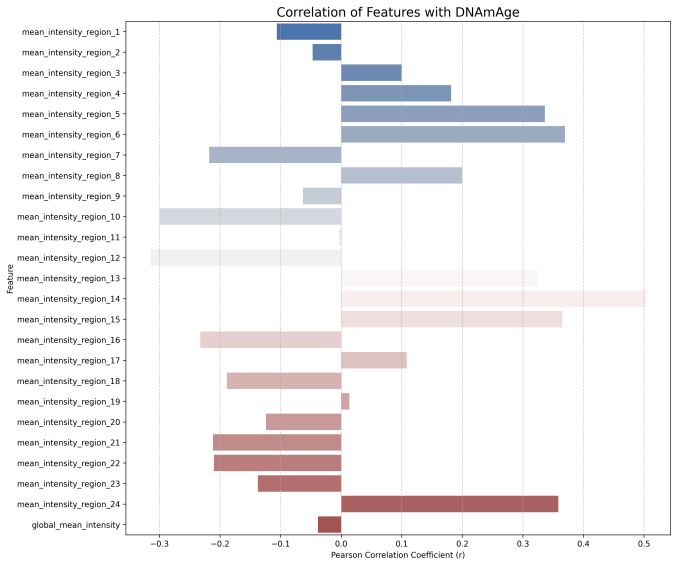

Figure 8. Pearson correlation coefficients (r) between DNAmAge and 24 regional, as well as global, mean signal intensities. The plot highlights generally weak linear relationships between these neuroimaging metrics and epigenetic age.

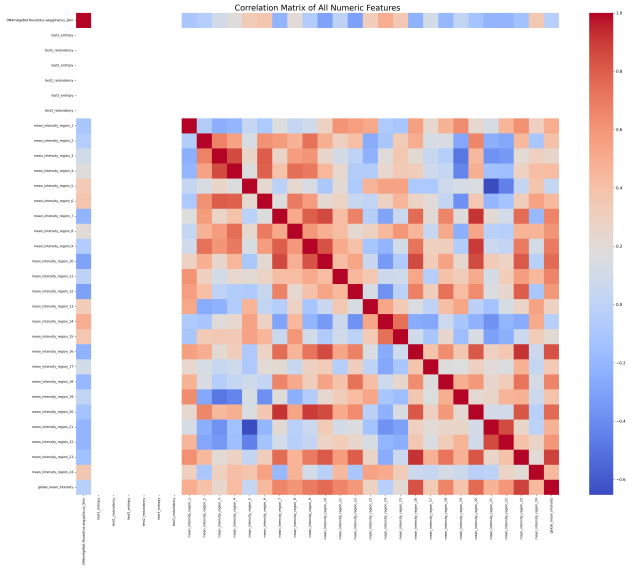

Figure 9. This heatmap illustrates the Pearson correlation coefficients among DNAmAge, behavioral metrics (Exploration Entropy and Navigational Redundancy), and regional and global mean signal intensity metrics. Consistent with their lack of variance, behavioral metrics exhibit no correlations with any other features. Overall, the neuroimaging mean signal intensity metrics, including the global average, show weak relationships with DNAmAge.

Validation (LOOCV) approach, providing an unbiased prediction for each subject.

The predictive model demonstrated poor performance, failing to robustly predict DNAmAge from the available neuroimaging and demographic features. The key performance metrics were an R-squared (R^2) value of -0.101 and a Mean Absolute Error (MAE) of 1.405 years. An R^2 value below zero indicates that the model's predictions are, on average, worse than simply predicting the mean DNAmAge for all subjects. The MAE of approximately 1.4 years is substantial, especially when considering the standard deviation of DNAmAge in the final cohort (1.64 years), indicating a high degree of prediction error relative to the age range. The scatter plot of predicted versus actual DNAmAge from the LOOCV, presented in Figure 10, visually confirms the wide scatter of data points around the identity line, illustrating the model's limited predictive power.

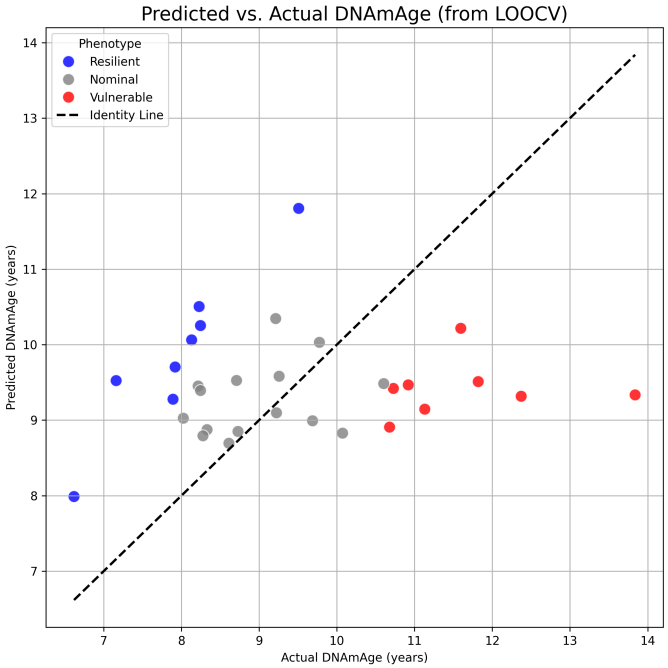

Figure 10. Predicted versus actual DNAmAge for each bat from the Leave-One-Out Cross-Validation of the Elastic Net model. Each point is colored according to its assigned aging phenotype (Resilient, Nominal, Vulnerable). The wide scatter of points around the identity line (dashed) visually confirms the model's poor predictive performance (R^2 = -0.101), indicating that the included neuroimaging and sex features do not effectively predict epigenetic age in this cohort.

These results underscore the challenges in identifying strong neuroimaging correlates of epigenetic aging using the available data and chosen features.

3.3.2. *Brain aging signatures*

Despite the overall poor predictive performance of the Elastic Net model, an exploratory analysis of the coefficients from the final model (trained on the full 31-

subject dataset) was conducted to identify features that received non-zero weights. This provides insights into which brain regions, if any, were weakly associated with epigenetic age in this dataset. The model selected 12 out of the 26 potential features (24 regional mean signal intensities, Global Mean Signal Intensity, Sex, and Origin) as having non-zero coefficients (Figure 11).

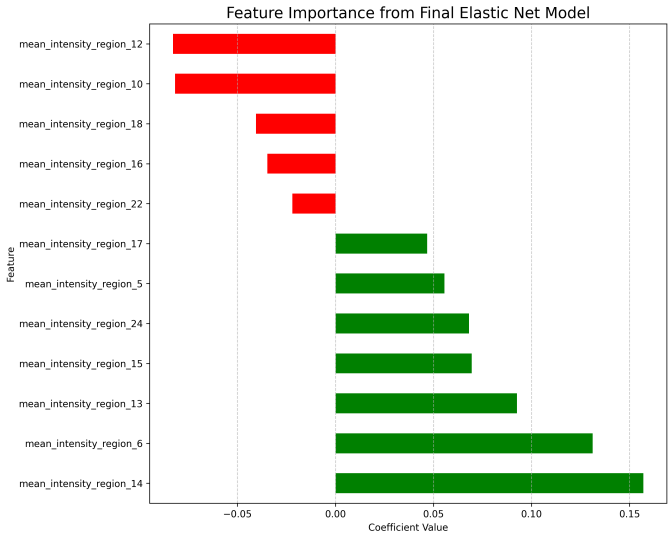

Figure 11. Feature Importance from Final Elastic Net Model. This bar plot shows the coefficients for the regional mean signal intensity features selected by the Elastic Net model trained to predict DNAmAge. Green bars indicate a positive association, suggesting higher signal intensity in these regions (e.g., regions 14, 6) is linked to increased epigenetic age. Red bars indicate a negative association (e.g., regions 12, 10). These coefficients provide exploratory insights into potential regional brain signatures associated with epigenetic age, despite the model's limited overall predictive performance.

The most influential positive predictors, meaning higher signal intensity in these regions was associated with older epigenetic age, included region 14 (coefficient = 0.157), region 6 (coefficient = 0.131), and region 13 (coefficient = 0.093). Conversely, the strongest negative predictors, where lower signal intensity was associated with older epigenetic age, were region 12 (coefficient = -0.083) and region 10 (coefficient = -0.082). Sex also emerged as a minor positive predictor (coefficient = 0.003). While these specific regions suggest potential brain-aging signatures, it is crucial to interpret these findings with extreme caution given the model's overall lack of predictive accuracy. These associations are weak and highly exploratory.

3.3.3. *Identification and characterization of aging phenotypes*

Leveraging the discrepancies between actual and predicted DNAmAge, bats were descriptively stratified into distinct aging phenotypes. For each bat, a prediction residual was calculated as the difference between their actual DNAmAge and their LOOCV-predicted DNAmAge. The distribution of these residuals is shown in Figure 12.

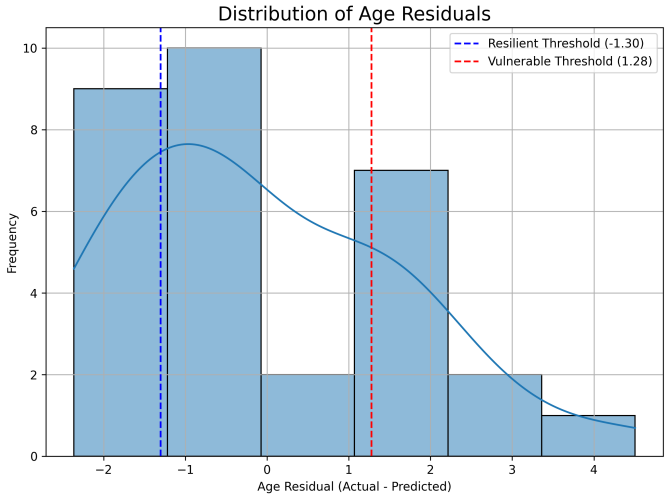

Figure 12. Distribution of age prediction residuals (actual minus predicted DNAmAge) from the Leave-One-Out Cross-Validation model. The dashed lines indicate the 25th (Resilient) and 75th (Vulnerable) percentile thresholds used for phenotype classification.

Bats falling into the bottom quartile of residuals (actual age significantly younger than predicted) were classified as "Resilient" (N=8). Conversely, bats in the top quartile of residuals (actual age significantly older than predicted) were classified as "Vulnerable" (N=8). The remaining subjects were designated "Nominal" (N=15).

To characterize the neuroimaging profiles of these identified phenotypes, the mean values of the top four most influential features from the Elastic Net model were compared across the "Resilient," "Nominal," and "Vulnerable" groups (Figure 13).

Consistent with the model's coefficients, the "Vulnerable" group generally exhibited higher mean signal intensity in regions 14 and 6 compared to the "Resilient" group. Conversely, the "Resilient" group tended to show higher mean signal intensity in region 12. These descriptive comparisons provide a preliminary characterization of the brain signal intensity patterns that, within the context of our non-predictive model, distinguish bats exhibiting accelerated versus decelerated epigenetic aging trajectories relative to their neuroimaging and demographic profile. These findings offer a framework for

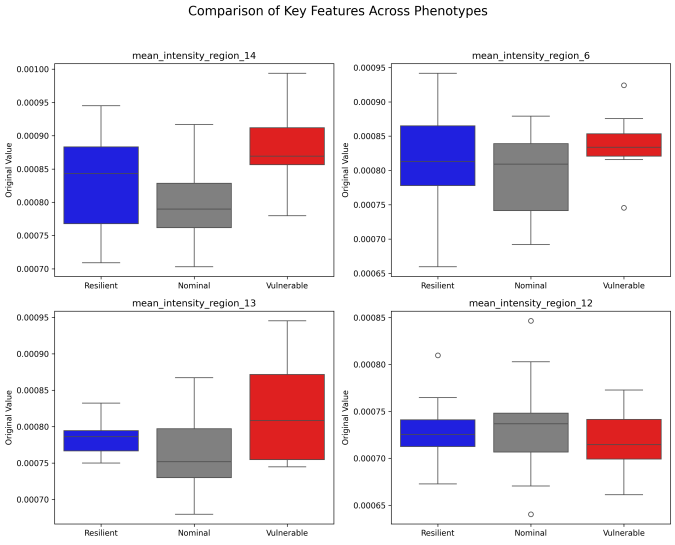

Figure 13. Comparison of Key Features Across Phenotypes. Boxplots show the distribution of mean signal intensity in brain regions 14, 6, 13, and 12 across "Resilient" (blue), "Nominal" (gray), and "Vulnerable" (red) bat phenotypes. These plots characterize the biological profiles of bats with varying epigenetic aging trajectories, revealing trends in neuroimaging features that align with their positive (regions 14, 6, 13) or negative (region 12) association with age in the predictive model.

phenotype identification that can be valuable for future studies with richer datasets.

3.4. *Summary of results*

This exploratory multi-modal investigation faced significant data limitations, particularly the absence of 4D MRI data for DW-SV calculation and the complete lack of variance in the intended behavioral metrics. These constraints necessitated a fallback to regional Mean Signal Intensity as the primary neuroimaging feature. The Elastic Net regression model, designed to predict DNAmAge from these available neuroimaging features and demographic factors, demonstrated poor predictive performance (R^2 = -0.101, MAE = 1.405 years), indicating that the selected features were insufficient for robust epigenetic age prediction in this cohort. Despite this, an exploratory analysis of model coefficients identified specific brain regions (e.g., regions 14, 6, 13, 12, 10) whose mean signal intensities were weakly associated with epigenetic age. Furthermore, a descriptive stratification based on age prediction residuals allowed for the identification and characterization of "Resilient" and "Vulnerable" aging phenotypes, revealing preliminary neuroimaging patterns that may distinguish these groups. While the current features did not yield a robust predictive model, the established framework for identifying aging phenotypes based on multi-modal data discrepancies remains a valuable contribution for future research with more comprehensive and higher-resolution datasets.

4. CONCLUSIONS

This exploratory study embarked on a multi-modal investigation in the long-lived Egyptian fruit bat (*Rousettus aegyptiacus*) to address the fundamental challenge of understanding the heterogeneous trajectories of aging, specifically aiming to identify multi-modal signatures of cognitive resilience and vulnerability. Our initial ambition was to integrate a novel Diffusion-weighted Signal Variability (DW-SV) metric from 4D MRI data with advanced behavioral entropy and efficiency measures to robustly predict DNA methylation (DNAm) age. The goal was to move beyond single-modality observations and capture the complex interplay across neurobiological, behavioral, and epigenetic domains that underlies individual aging differences.

However, the execution of this ambitious multi-modal integration was significantly impacted by practical data constraints. The provided MRI data were in a 3D format, precluding the calculation of the intended 4D DW-SV metric, necessitating a fallback to regional Mean Signal Intensity as the primary neuroimaging feature. Similarly, the planned behavioral metrics, Exploration Entropy and Navigational Redundancy, exhibited a complete lack of variance across the bat cohort, rendering them uninformative for predictive modeling and consequently excluded from the final analysis. For a final cohort of 31 bats with complete neuroimaging and epigenetic data, an Elastic Net regression model, incorporating regional Mean Signal Intensity and demographic factors (sex, origin), was employed with Leave-One-Out Cross-Validation to predict DNAm age.

The results of our predictive modeling revealed significant limitations. The Elastic Net model demonstrated poor predictive performance for DNAm age, yielding an R-squared value of -0.101 and a Mean Absolute Error of 1.405 years. This indicated that the available neuroimaging features (regional Mean Signal Intensity) and demographic factors were not robust predictors of epigenetic age in this specific dataset. Neither global nor individual regional mean signal intensities exhibited strong correlations with DNAm age. Despite the model's overall lack of predictive power, an exploratory analysis of its coefficients highlighted specific brain regions (e.g., regions 14, 6, 13 showing positive associations, and regions 12, 10 showing negative associations with epigenetic age), suggesting weak and highly exploratory brain-aging signatures. Furthermore, leveraging the discrepancies between actual and predicted

DNAm age, we descriptively stratified bats into 'Resilient' and 'Vulnerable' phenotypes. Preliminary characterization of these groups revealed distinct patterns in their regional mean signal intensities, consistent with the model's coefficients. For instance, the 'Vulnerable' group tended to exhibit higher mean signal intensity in regions 14 and 6, while the 'Resilient' group showed higher intensity in region 12.

From this exploratory investigation, we have gained several key insights. Firstly, the study powerfully underscores the substantial challenges inherent in multi-modal data integration for aging research, particularly when confronted with unforeseen data limitations. The necessity to adapt our neuroimaging approach and exclude behavioral metrics due to data characteristics highlights the critical importance of data quality, format, and experimental design in achieving the full potential of multi-modal analyses. Secondly, the findings indicate that simple regional Mean Signal Intensity, as a proxy for brain microstructure, along with basic demographic factors, may be insufficient to capture the complex biological underpinnings of epigenetic aging in Egyptian fruit bats. More sophisticated neuroimaging metrics, potentially derived from higher-dimensional data (e.g., true 4D DTI for DW-SV, or other advanced MRI sequences), are likely required to unveil robust brain-epigenetic age relationships. Lastly, despite the predictive limitations, the established framework for descriptively identifying and characterizing 'Resilient' and 'Vulnerable' aging phenotypes based on discrepancies between biological age and predicted age from other modalities remains a valuable contribution. This framework provides a robust conceptual and analytical pipeline that can be readily applied in future studies with richer, more comprehensive, and higher-resolution datasets. Such future endeavors, integrating more granular neuroimaging and behavioral data, are crucial to fully unravel the intricate multi-modal signatures of heterogeneous aging trajectories in long-lived model organisms like the Egyptian fruit bat, ultimately informing strategies for promoting healthy longevity.