

# MethylGPT: a foundation model for the DNA methylome

Kejun Ying<sup>1,2,†,\*</sup>, Jinyeop Song<sup>3,†</sup>, Haotian Cui<sup>4,5,6,†</sup>, Yikun Zhang<sup>1,†</sup>, Siyuan Li<sup>1</sup>, Xingyu Chen<sup>5,6</sup>, Hanna Liu<sup>1</sup>, Alec Eames<sup>1</sup>, Daniel L McCartney<sup>7</sup>, Riccardo E. Marioni<sup>7</sup>, Jesse R. Poganik<sup>1</sup>, Mahdi Moqri<sup>1,\*</sup>, Bo Wang<sup>5,6,8,\*</sup>, Vadim N. Gladyshev<sup>1,\*</sup>

<sup>1</sup> Division of Genetics, Department of Medicine, Brigham and Women's Hospital and Harvard Medical School, Boston, MA, USA

<sup>2</sup> T. H. Chan School of Public Health, Harvard University, Boston, MA, USA

<sup>3</sup> Department of Physics, MIT, Cambridge, MA, USA

<sup>4</sup> Peter Munk Cardiac Centre, University Health Network, Toronto, ON, Canada

<sup>5</sup> Department of Computer Science, University of Toronto, Toronto, ON, Canada

<sup>6</sup> Vector Institute, Toronto, ON, Canada

<sup>7</sup> Centre for Genomic and Experimental Medicine, Institute of Genetics and Molecular Medicine, University of Edinburgh, Edinburgh, Scotland, United Kingdom

<sup>8</sup> Department of Laboratory Medicine and Pathobiology, University of Toronto, Toronto, ON, Canada

† Equal contribution

\* Correspondence Email: KY ([kying@g.harvard.edu](mailto:kying@g.harvard.edu)), MM ([mmoqri@bwh.harvard.edu](mailto:mymoqri@bwh.harvard.edu)), BW ([bowang@vectorinstitute.ai](mailto:bowang@vectorinstitute.ai)), and VNG ([vgladyshev@bwh.harvard.edu](mailto:vgladyshev@bwh.harvard.edu))

## Abstract

DNA methylation serves as a powerful biomarker for disease diagnosis and biological age assessment. However, current analytical approaches often rely on linear models that cannot capture the complex, context-dependent nature of methylation regulation. Here we present MethylGPT, a transformer-based foundation model trained on 226,555 (154,063 after QC and deduplication)

24 human methylation profiles spanning diverse tissue types from 5,281 datasets, curated 49,156  
25 CpG sites, and 7.6 billion training tokens. MethylGPT learns biologically meaningful representa-  
26 tions of CpG sites, capturing both local genomic context and higher-order chromosomal features  
27 without external supervision. The model demonstrates robust methylation value prediction (Pear-  
28 son  $R=0.929$ ) and maintains stable performance in downstream tasks with up to 70% missing  
29 data. Applied to age prediction across multiple tissue types, MethylGPT achieves superior accu-  
30 racy compared to existing methods. Analysis of the model's attention patterns reveals distinct  
31 methylation signatures between young and old samples, with differential enrichment of devel-  
32 opmental and aging-associated pathways. When finetuned to mortality and disease prediction  
33 across 60 major conditions using 18,859 samples from Generation Scotland, MethylGPT  
34 achieves robust predictive performance and enables systematic evaluation of intervention effects  
35 on disease risks, demonstrating potential for clinical applications. Our results demonstrate that  
36 transformer architectures can effectively model DNA methylation patterns while preserving bio-  
37 logical interpretability, suggesting broad utility for epigenetic analysis and clinical applications.

## 38 **Introduction**

39 DNA methylation is an epigenetic modification where methyl groups are added to cytosine resi-  
40 dues at CpG dinucleotides. This modification regulates gene expression by recruiting methyl-  
41 CpG binding proteins and modifying chromatin accessibility<sup>1</sup>. DNA methylation regulates mul-  
42 tiple biological processes through distinct mechanisms. During development, dynamic methyla-  
43 tion changes guide cellular differentiation by silencing lineage-inappropriate genes and activat-  
44 ing cell-type-specific programs<sup>2</sup>. Methylation also maintains genomic stability through the re-  
45 pression of transposable elements<sup>3</sup>.

46 Beyond its fundamental role in gene regulation, DNA methylation exhibits key characteristics of  
47 an ideal biomarker: stability in the resting state, but with dynamic response to environmental fac-  
48 tors, accessibility in various biological specimens, and early alterations preceding clinical mani-  
49 festations<sup>4</sup>. Genome-wide methylation profiling has revealed distinctive signatures across nu-  
50 merous pathological states, enabling molecular diagnostics, particularly in cancer detection and  
51 cardiovascular risk assessment<sup>5</sup>.

52 Alongside disease prediction, age-associated methylation patterns also enable the development  
53 of highly accurate “epigenetic aging clocks”<sup>6</sup>. These clocks have evolved from simple age pre-  
54 dictors to sophisticated biomarkers of biological aging, with recent advances such as  
55 DunedinPACE<sup>7</sup>, GrimAge<sup>8</sup>, causality-enriched clocks<sup>9</sup>, and the high-dimensional ageome<sup>10</sup>,  
56 demonstrating strong associations with health outcomes and mortality risk. Notably, these meth-  
57 ylation-based aging indices often outperform conventional clinical measures in predicting age-  
58 related diseases and longevity<sup>11,12</sup>, highlighting their potential for monitoring therapeutic inter-  
59 ventions targeting the aging process.

60 However, several analytical challenges impede the clinical implementation of methylation-based  
61 diagnostics. Current computational approaches predominantly rely on linear models and simple  
62 statistical methods, which are fundamentally limited in their ability to capture complex, non-  
63 linear relationships in methylation data. These linear models assume independence between CpG  
64 sites, failing to account for the regulatory networks and higher-order interactions that character-  
65 ize methylation patterns. Moreover, the same DNA methylation pattern may have different bio-  
66 logical implications depending on the cellular and tissue context: a complexity that linear models  
67 are unable to capture<sup>13–15</sup>. The limitations of linear models become even more apparent when  
68 dealing with technical artifacts, including batch effects and missing data, which introduce sub-  
69 stantial non-linear variability in methylation measurements<sup>16</sup>. The field urgently needs a unified  
70 analytical framework capable of modeling complex, non-linear patterns, accounting for context-  
71 dependent effects, and performing robust pattern analysis across diverse clinical contexts.

72 Recent advances in artificial intelligence, particularly transformer architectures and foundation  
73 models<sup>17</sup>, have revolutionized the analysis of complex biological sequences. Foundation models  
74 have emerged across multiple omics layers: for proteomics, ESM-2/ESM-3<sup>18,19</sup> and  
75 AlphaFold2/AlphaFold3<sup>20,21</sup> have achieved unprecedented accuracy in structure prediction and  
76 function annotation; for genomics, Enformer<sup>22</sup> and Evo<sup>23</sup> have demonstrated capability in pre-  
77 dicting gene regulation and variant effects. In the single-cell domain, models like Geneformer<sup>24</sup>,  
78 scGPT<sup>25</sup>, and scFoundation<sup>26</sup> have enabled zero-shot cell-type classification and in-silico per-  
79 turbation. And more recently, the Precious3GPT has emerged as a multimodal transformer model  
80 integrating multi-omics data for aging research and drug discovery<sup>27</sup>

81 These foundation models demonstrate remarkable capability in learning comprehensive biologi-  
82 cal patterns that generalize across tasks. However, despite the success of foundation models  
83 across various omics layers, DNA methylation analysis lacks such a unified approach, relying  
84 instead on task-specific models that fail to capture the full complexity of methylation patterns.  
85 The achievements of foundation models in related domains suggest that a similar approach could  
86 transform methylation analysis by providing a unified framework that preserves biological con-  
87 text while enabling adaptations to diverse specific tasks.

88 Here, we introduce MethylGPT (Fig. 1a), a transformer-based foundation model for DNA meth-  
89 ylation. Trained on methylation profiles from over 150,000 human samples spanning diverse tis-  
90 sue types, MethylGPT implements a novel embedding strategy to capture methylation patterns at  
91 physiologically relevant CpG sites. This approach enables unified analysis of DNA methylation  
92 data across multiple experimental contexts and downstream applications, including age predic-  
93 tion and disease association detection.

## 94 **Results**

### 95 **Development and validation of MethylGPT**

96 To enable the pretraining of large-scale model, we collected 226,555 human DNA methylation  
97 profiles from 5,281 datasets through the EWAS Data Hub and Clockbase<sup>28,29</sup>. After quality con-  
98 trol and deduplication, we used 154,063 samples to pretrain MethylGPT. The model focuses on  
99 49,156 physiologically-relevant CpG sites, selected based on association with EWAS traits  
100 (Methods)<sup>30</sup>. These methylation profiles, representing samples from over 20 different tissue  
101 types, were processed to generate 7.6 billion training tokens (CpG sites), enabling comprehen-  
102 sive coverage of methylation patterns across the human epigenome.

103 The core architecture of MethylGPT consists of a methylation embedding layer followed by 12  
104 transformer blocks (Fig. 1a). Our methylation embedding process captures both the CpG site to-  
105 kens and their methylation states through an element-wise attention mechanism. This design en-  
106 ables the model to learn complex dependencies between distant CpG sites while maintaining lo-  
107 cal methylation context. The model was pre-trained using two complementary loss functions: a  
108 masked language modeling (MLM) loss where the model predicts methylation levels for 30%

109 randomly masked CpG sites and a reconstruction loss where the Classify token (CLS) embed-  
110 ding is used to reconstruct the complete DNA methylation profile.

111 To evaluate the model's performance, we first assessed its ability to predict DNA methylation  
112 values at masked CpG sites in the test set. During training, the model achieved rapid conver-  
113 gence with minimal overfitting, reaching a best model test mean squared error (MSE) of 0.014 at  
114 epoch 10 (Fig. 1b). The model demonstrated robust prediction accuracy across different methyla-  
115 tion levels, achieving an overall mean absolute error (MAE) of 0.074 and a Pearson correlation  
116 coefficient of 0.929 between predicted and actual methylation values (Fig. 1c-f).

### 117 **MethylGPT learns biologically meaningful CpG representations**

118 To investigate whether MethylGPT captures biologically relevant DNA methylation features, we  
119 analyzed the learned representations of 49K CpG sites in the embedding space (Fig. 2a). Dimen-  
120 sionality reduction using UMAP revealed distinct patterns in the contextualized CpG embedding  
121 space (Fig. 2b). CpG sites clustered according to their genomic contexts, with clear separation  
122 based on CpG island relationships (island, shore, shelf, and other regions), suggesting that our  
123 model learned underlying regulatory features of the methylome without explicit supervision.

124 The embedding space organization reflected known biological properties of DNA methylation  
125 regulation. CpG sites within enhancer regions showed distinct clustering patterns (Fig. 2c), con-  
126 sistent with their specialized regulatory roles. Furthermore, the embeddings demonstrated a clear  
127 separation of sex chromosomes from autosomes (Fig. 2d). This organization indicates that  
128 MethylGPT successfully captured both local sequence context and higher-order chromosomal  
129 features that influence methylation patterns.

130 The transformer architecture enabled our model to learn these complex relationships through its  
131 attention mechanism, which integrates both local CpG site features and broader genomic context  
132 (Fig. 2a) instead of treating CpG sites as independent entities as in previous methods.

### 133 **MethylGPT learns tissue-specific and sex-specific methylation patterns**

134 To evaluate whether MethylGPT captures biologically meaningful sample-level features, we  
135 analyzed the zero-shot embedding spaces of DNA methylation samples before and after model  
136 processing. The contextualized sample embeddings from MethylGPT showed clear biological  
137 organization, with distinct clustering patterns by tissue type and sex (Fig. 3a). Major tissue types,

138 including whole blood, brain, liver, and skin, formed well-defined clusters, suggesting that  
139 MethylGPT successfully learned tissue-specific methylation signatures without explicit supervi-  
140 sion. Notably, batch effects were not significant in the observed embeddings (Fig. 3b).  
141 MethylGPT embeddings also revealed strong sex-specific methylation patterns across tissues  
142 (Fig. 3c). Male and female samples showed consistent separation in the embedding space, re-  
143 flecting known sex-specific methylation differences.

144 The superiority of MethylGPT's learned representations becomes apparent when compared to  
145 the raw methylation data directly generated UMAP embeddings (Fig. 3d-f). While raw methyla-  
146 tion profiles showed some degree of tissue-specific clustering, the boundaries between different  
147 tissue types were less distinct, and the overall organization was more diffuse. The raw data  
148 embeddings exhibited less defined tissue-specific clusters (Fig. 3d), stronger batch-specific clus-  
149 tering (Fig. 3e), and weaker sex-specific separation (Fig. 3f), highlighting MethylGPT's ability  
150 to enhance biologically relevant signals through its contextualized embedding approach.

### 151 **MethylGPT enables accurate age prediction across diverse tissue types**

152 To evaluate MethylGPT's capability in downstream applications, we first assessed its perfor-  
153 mance in predicting chronological age from DNA methylation patterns. We utilized a diverse  
154 dataset of 11,453 samples spanning multiple tissue types<sup>31</sup>, with an age distribution ranging  
155 from 0 to 100 years (Fig. 4a). The majority of samples were derived from whole blood (47.2%)  
156 and brain tissue (34.5%), providing broad coverage of physiologically distinct methylation pat-  
157 terns.

158 The pre-trained MethylGPT embeddings showed inherent age-related organization even before  
159 fine-tuning (Fig. 4b), suggesting that the model captured age-associated methylation features  
160 during pre-training. After fine-tuning for age prediction, the sample embeddings demonstrated  
161 stronger age-dependent clustering (Fig. 4c) while maintaining tissue-specific patterns (Fig. 4d).

162 We compared MethylGPT's age prediction performance against existing methods, including  
163 ElasticNet<sup>32</sup>, MLP (AltumAge)<sup>31</sup>, Horvath's skin and blood clock<sup>33</sup>, and other established age  
164 predictors. MethylGPT achieved superior accuracy with a median absolute error (MedAE) of  
165 4.45 years on the validation set, outperforming other methods (Fig. 4e). This improvement was

166 consistent across both validation and test sets, demonstrating the model's robust generalization  
167 capability.

168 Notably, MethylGPT showed remarkable resilience to missing data, a common challenge in  
169 methylation analysis. We systematically evaluated prediction performance under increasing lev-  
170 els of data missingness (10-90%). MethylGPT maintained stable performance with up to 70%  
171 missing data, significantly outperforming both ElasticNet and Multi-Layer Perceptron (MLP)  
172 approaches (Fig. 4f). This robustness suggests that the model's contextualized embeddings effec-  
173 tively capture redundant age-related signals across multiple CpG sites, enabling reliable predic-  
174 tions despite incomplete methylation profiles.

175 To further validate MethylGPT's ability to capture biologically meaningful age-related patterns,  
176 we analyzed DNA methylation profiles during iPSC reprogramming<sup>34</sup>. The model's embeddings  
177 revealed a clear rejuvenation trajectory (Fig. 4g), with samples progressively shifting towards a  
178 younger methylation state over the reprogramming time course. Notably, when compared with  
179 conventional epigenetic clocks (Horvath's clock and GrimAge), MethylGPT showed consistent  
180 detection of rejuvenation effects, predicting a significant decrease in epigenetic age during re-  
181 programming (Fig. 4h). This agreement with established aging biomarkers, while accounting for  
182 the broader epigenomic context through the transformer architecture, provides independent sup-  
183 port for iPSC reprogramming as a rejuvenation method rather than merely a cell identity trans-  
184 formation. The predicted age trajectory showed a sharp decline after day 20 of reprogramming,  
185 reaching near-zero predicted ages by day 30, consistent with the restoration of a pluripotent epi-  
186 genetic state.

### 187 **Age-specific attention patterns reveal distinct methylation signatures**

188 To investigate how MethylGPT processes age-related methylation patterns, we analyzed the  
189 model's multi-head self-attention weights (Fig. 5a). By examining the attention weight matrices,  
190 we observed that the model learned distinct patterns of CpG site interactions between young (age  
191 < 20) and old (age > 60) samples, suggesting that the transformer architecture captures age-  
192 specific relationships in methylation data.

193 We further analyzed the attention weight distributions across three age groups (< 20, 20-60, and  
194 > 60 years) to understand how the model's attention mechanism adapts to different age ranges



195 (Fig. 5b). The attention patterns revealed systematic shifts in how the model weighs relationships  
196 between CpG sites across the lifespan, potentially reflecting underlying biological changes in  
197 methylation regulation during aging. Interestingly, attention weights are concentrated on a few  
198 CpG sites, suggesting that this sparse set of sites may be significantly relevant to age-specific  
199 methylation attention. To identify such statistically influential CpG sites, we extracted sites with  
200 large differential attention scores ( $>1.5$  fold change) that were statistically significant (p-value  $<$   
201  $0.05$ ) between young and old samples (Fig. 5c). We analyzed the associated EWAS traits and  
202 age-specific methylation signatures of the identified important CpGs in both young and old sam-  
203 ples (Fig. 5d). In young samples, high-attention CpG sites showed the strongest associations with  
204 non-age-associated phenotypes, including sex and autoimmune diseases. Conversely, old sam-  
205 ples showed strong attention weights at CpG sites associated with aging, as well as aging-related  
206 traits like BMI and thyroid lesions<sup>35</sup>, validating our model's biological relevance.

207 To understand the biological significance of age-specific attention patterns, we performed func-  
208 tional enrichment analysis on CpG sites with differential attention weights between young and  
209 old samples. Gene Ontology (GO) and Reactome pathway analysis revealed distinct biological  
210 processes associated with high-attention CpG sites in each age group (Fig. 5e). In young sam-  
211 ples, highly attended CpG sites were enriched for developmental processes, including cellular  
212 response to growth factor stimulus. In contrast, CpG sites receiving higher attention in older  
213 samples showed enrichment for oxidative stress and amino acid metabolism. These enrichment  
214 patterns validate that MethylGPT's attention mechanism captures biologically meaningful age-  
215 specific methylation signatures.

### 216 **Disease risk prediction and intervention analysis**

217 To evaluate MethylGPT's utility in clinical applications, we analyzed its ability to predict disease  
218 risks and assess intervention effects in the Generation Scotland cohort ( $n = 18,859$ ). We fine-  
219 tuned the pre-trained model to predict the risk of 60 major diseases across eight categories, in-  
220 cluding cardiovascular, respiratory, neurological, and autoimmune conditions, as well as overall  
221 mortality, over a 10-year window (Fig. 6a,b). Our results demonstrate that the model achieved an  
222 overall Area Under the Curve (AUC) of 0.74 on the validation set and 0.72 on the test set (Fig.  
223 6c).



224 Using this disease prediction framework, we systematically evaluated the impact of eight differ-  
225 ent interventions on predicted disease incidence (Fig. 6d). The model revealed distinct, interven-  
226 tion-specific effects across disease categories. Smoking cessation demonstrated the strongest  
227 protective effect against 10-year mortality ( $\beta = -0.13$ ) and also reduced cardiovascular disease  
228 risk. Notably, high-intensity training showed strong benefits for respiratory, neurological and  
229 autoimmune diseases. Similarly, the Mediterranean diet provided modest but consistent protec-  
230 tive effects across multiple disease categories, though with varying magnitude.

231 Interestingly, Everolimus treatment showed a significant risk increase for autoimmune diseases.  
232 Although counter-intuitive, this finding is consistent with previous studies showing that pro-  
233 longed immunosuppressant treatment is associated with an increased incidence of autoimmune  
234 diseases <sup>36</sup>.

235 Together, these findings demonstrate the potential of MethylGPT for predicting intervention-  
236 specific health outcomes and optimizing personalized intervention strategies.

## 237 Discussion

238 DNA methylation patterns have shown potential as a universal biomarker for disease stratifica-  
239 tion and monitoring. In oncology, methylation patterns enable the identification of cancer tissue  
240 of origin, achieving 81-93% accuracy in predicting primary sites of metastatic tumors and can-  
241 cers of unknown primary origin <sup>37</sup>. Methylation-based cardiovascular risk scores demonstrate  
242 superior predictive accuracy compared to conventional clinical factors <sup>38</sup>. Furthermore, methyla-  
243 tion markers can predict type 2 diabetes onset years before clinical presentation, providing criti-  
244 cal windows for preventive intervention <sup>39</sup>.

245 Our results demonstrate that a transformer-based foundation model approach can effectively  
246 model DNA methylation patterns while maintaining biological relevance. The organization of  
247 CpG sites in the embedding space based on genomic context and regulatory features suggests  
248 that MethylGPT captures fundamental aspects of methylation regulation without explicit super-  
249 vision. This capability addresses a key limitation of traditional linear models that treat CpG sites  
250 as independent entities.

251 The model's performance in age prediction across diverse tissue types, with improved accuracy  
252 over existing methods, demonstrates its potential utility. Particularly notable is the resilience to  
253 missing data, maintaining stable performance with up to 70% missingness. This robustness likely  
254 stems from the model's ability to leverage redundant biological signals across multiple CpG  
255 sites.

256 Analysis of age-specific attention patterns revealed distinct methylation signatures between  
257 young and old samples. The enrichment of development-related processes in younger samples  
258 and aging-associated pathways in older samples, which is consistent with previous studies<sup>40,41</sup>,  
259 suggests that the attention mechanism captures biologically meaningful age-dependent changes  
260 in methylation regulation. These findings provide new insights into how methylation patterns  
261 evolve across the lifespan.

262 Several directions for future research emerge from this work. Integration of additional epigenetic  
263 features beyond CpG methylation could provide a more comprehensive view of regulatory  
264 mechanisms. The development of interpretable attention visualization tools could help bridge the  
265 gap between model predictions and biological mechanisms. Additionally, exploring the model's  
266 application to single-cell methylation analysis could reveal cell-type-specific regulatory patterns.

267 In conclusion, MethylGPT demonstrates how transformer architectures can capture context-  
268 dependent methylation patterns while maintaining biological interpretability. The model's robust  
269 performance in handling missing data suggests potential utility in both research and clinical ap-  
270 plications.

## 271 **Methods**

### 272 **Pretraining data collection and preprocessing**

273 For the pretraining dataset for the methylGPT, we compose DNA methylation data from 154,063  
274 human samples through the EWAS Data Hub and Clockbase<sup>28,29</sup>. For quality control, we initial-  
275 ly collected data from approximately 300,000 patients and filtered out low-quality entries with  
276 high levels of missing data (>40% of total CpG sites). We also applied deduplication to ensure  
277 no repetitions in the training data. The cleaned dataset was randomly sampled and quality-  
278 checked, covering individuals across 20 distinct tissue types<sup>42</sup>.

279 DNA methylation data have varying numbers of CpG entries depending on the array platform  
280 (Illumina 27k, Illumina 450k, and EPIC). To address these differences and ensure biological rel-  
281 evance, we focused on 49,156 CpG sites selected based on importance by EWAS traits<sup>30</sup> and  
282 array format compatibility. In detail, these 49,156 CpG sites satisfy either (1) CpG are associated  
283 with more than 5 traits according to EWAS catalog or (2) CpG appears in more than 95% of the  
284 pretraining dataset. All methylation values were normalized using standard protocols. Missing  
285 values were marked for downstream masked prediction tasks.

286 Data is processed into a matrix of  $X \in \mathbb{R}^{N \times M}$ , where each element  $X_{i,j}$  denotes the magnitude of  
287 methylation of a CpG site  $j$  in sample  $i$ .  $N$  is the number of samples and  $M$  is the number of CpG  
288 sites (i.e. 49,156).

## 289 **Model architecture**

290 MethylGPT consists of three main components: an embedding module, a transformer module,  
291 and task-specific heads. The input data  $X$  is tokenized and fed into the modules consecutively.  
292 We depict the input tokenization and the module details as follows.

## 293 **CpG site tokenization**

294 The processed data contains methylation readings of  $M$  (49,156) CpG sites. For each site  
295  $c_{ij}$  ( $j \in \{0, 1, \dots, M\}$ ), we assign an integer identifier  $id(c_{ij})$ . The full CpG tokens for an individ-  
296 ual sample are  $t_c^{(i)} = id(c_{ij})$ .

## 297 **Embedding layers**

298 We utilize the embedding layers for the CpG site tokens to map each token to a fixed-length em-  
299 bedding vector of dimension  $D$ . We employ fully connected layers for the methylation values to  
300 encode the methylation level into vector embeddings and maintain the ordinal relation of the val-  
301 ues.

302 For each CpG site, the embedding module projects both CpG site identifiers and their methyla-  
303 tion values into separate embeddings (referred to as CpG embeddings and methylation value  
304 embeddings), which are then merged through an element-wise sum. The final embedding for  
305 sample  $i$  is defined as:

$$h^{(i)} = emb_c(t_c^{(i)}) + emb_x(x^{(i)})$$

306 The embedding dimension is set to 64. A special [CLS] token is prepended to each sequence for  
307 learning sample-level representations.

### 308 **Transformer module**

309 We employ the self-attention transformer<sup>17,43</sup> to encode the complete input embedding. The  
310 transformer module comprises 6 transformer blocks, each containing a multi-head self-attention  
311 layer (4 heads) and a standard MLP layer. Layer normalization and residual skip connections are  
312 applied after each layer. The self-attention mechanism operates on the sequence of M embedding  
313 vectors, making it particularly suitable for capturing interactions between CpG sites. The trans-  
314 former processes the sequence according to:

$$h_0^{(i)} = h^{(i)}, h_l^{(i)} = transformer\_block(h_{l-1}^{(i)}) \forall l \in [1, n]$$

315 We utilize the resulting representation  $h_n^{(i)} \in \mathbb{R}^{M,D}$  for both CpG-level and sample-level tasks.  
316 The self-attention layer leverages FlashAttention for efficient training and inference<sup>44</sup>. The  
317 model dimension is set to 64, with also an intermediate dimension of 64 in the feed-forward lay-  
318 ers. The transformer module processes a sequence of input embeddings comprising 49,157 sites  
319 with 64 dimensions and outputs “contextualized embeddings” of the same shape.

320 The input dimension M can reach tens of thousands of CpG sites, consuming huge memory and  
321 creating a significant challenge for efficient model training. We leverage the Flash-Attention<sup>44</sup>  
322 implementation as a tool to greatly accelerate the training and inference of the model while min-  
323 imizing memory footprint.

324 The task-specific heads attached to the transformer process contextualized embeddings into di-  
325 verse predictions specific to the task. In the pre-training phase, a linear layer projects output  
326 embeddings of each CpG site to predict the methylation value. In the fine-tuning phase, the MLP  
327 or convolutional layers process the complete output embeddings to predict biological age or oc-  
328 currence of disease.

## 329 **Model pretraining**

330 The model was trained on two complementary objectives. First, we randomly masked 30% of  
331 CpG sites (i.e., their value embeddings were excluded from the input embedding process) and  
332 trained the model to reduce the MSE between the predicted and original methylation values at  
333 the masked CpG sites. The Methylation Value Prediction (MVP) objective is defined as:

$$\tilde{x}^{(i)} = MLP(h_n^{(i)}), L_{MVP} = \frac{1}{|M_{mask}|} \sum_{j \in M_{mask}} (\tilde{x}_j^{(i)} - x_j^{(i)})^2$$

334 where  $\tilde{x}^{(i)} \in \mathbb{R}^M$  represents the row of predicted methylation value estimates for sample  $i$ . The  
335 MVP objective encourages the model to effectively encode relationships among the CpG sites in  
336 the dataset.

337 Second, a profile reconstruction task used the [CLS] token output embedding to reconstruct  
338 complete methylation profiles, as also described in a previous study<sup>25</sup>. The model feeds the  
339 [CLS] token's output embedding from the previous step back into the [CLS] token input, while  
340 all other tokens are masked. The objective of the profile reconstruction task is:

$$\tilde{x}^{(i)} = MLP(h_n^{(i)}), L_{MVP} = \frac{1}{|M_{all}|} \sum_{j \in M_{all}} (\tilde{x}_j^{(i)} - x_j^{(i)})^2$$

341 Training was performed using the AdamW optimizer with a learning rate of 0.001, The model  
342 was trained for 10 epochs with a batch size of 16 on NVIDIA A100 GPU. The learning rate is set  
343 to decay 10% after each epoch.

## 344 **Evaluation metrics**

345 Model performance was assessed through multiple metrics. We calculated the MSE and MAE for  
346 methylation value prediction, along with Pearson correlation coefficients between predicted and  
347 actual methylation values. For age prediction tasks, we measured accuracy using MedAE in  
348 years. For disease prediction tasks, model effectiveness was evaluated using the AUC, which  
349 measures the model's discriminative ability to differentiate between various disease states.

## 350 **Age prediction experiments**

351 For the evaluation of age prediction, we utilized a dataset comprising 13,505 samples with  
352 21,368 CpG sites<sup>31</sup>. From the accompanying metadata, we identified training (5,461 samples),  
353 validation (1,366 samples), and test (4,626 samples) sets, with a total of 49,156 CpG sites.

354 We fine-tuned pre-trained MethylGPT using the downstream prediction head ResNet1D. The  
355 ResNet1D consists of six residual blocks, where each residual block includes two 1D convolu-  
356 tional layers with a kernel size of 3, each followed by batch normalization and ReLU activation.  
357 Specifically, we input 49,156 CpG sites into the MethylGPT, generating an embedding with di-  
358 mensions (49,156, 64). To reduce dimensionality, this embedding was passed through a 3x3 1D  
359 convolutional layer, condensing the feature space to 32 channels. The reduced-dimensionality  
360 output was subsequently fed into six residual blocks, followed by an average pooling layer and a  
361 linear layer for age prediction. Both the pre-trained MethylGPT and the downstream ResNet1D  
362 prediction head were trained using the MSE loss function as the optimization objective.

363 To assess robustness, we systematically masked increasing proportions (10-90%) of CpG sites in  
364 the test set and evaluated prediction performance. Comparison methods (ElasticNet, MLP,  
365 Horvath's clock) were trained and evaluated on the same data splits.

## 366 **Disease prediction and intervention evaluation**

367 We fine-tuned the pre-trained model, maintaining consistency with the downstream prediction  
368 head architecture, ResNet1D, used in age prediction to demonstrate the generalizability of the  
369 pre-trained model across downstream tasks. By utilizing the same downstream network structure  
370 for both age and disease prediction, we aimed to confirm that the model's effectiveness was not  
371 due to meticulous architecture optimization but rather due to its inherent flexibility.

372 To evaluate this, we curated datasets from the Generation Scotland cohort (n = 18,859), compris-  
373 ing 1,378 samples for training, 295 for validation, and 296 for testing. In fine-tuning, the model  
374 was trained to simultaneously predict the risk of 60 major diseases across eight categories, lead-  
375 ing to the development of a comprehensive disease prediction model. For each disease category,  
376 a sample was labeled as '1' if the disease was present and '0' otherwise. This multi-label classi-  
377 fication task, where a sample could have one or multiple co-occurring diseases, introduced sub-

378 stantial complexity to the prediction challenge. Both the pre-trained MethylGPT and the down-  
379 stream ResNet1D prediction head were optimized using the cross-entropy loss function.

380 To further explore the impact of interventions on predictive outcomes, we applied the disease  
381 prediction model to assess data from eight types of interventions across six GEO datasets  
382 (GSE219217<sup>45</sup>, GSE268211<sup>46</sup>, GSE176325<sup>47</sup>, GSE191297<sup>48</sup>, GSE201532<sup>49</sup>, GSE276988<sup>50</sup>),  
383 encompassing a total of 183 samples. The interventions examined in this study included Mediter-  
384 ranean fiber (n=36), high-intensity training (n=5), folate supplementation (n=43), anti-TNF ther-  
385 apy (n=59), smoking cessation (n=16), glyNAC (n=8), everolimus (n=8), and metformin (n=8).  
386 Each intervention included an intra-group control as part of the trial design. For the phased inter-  
387 ventions, only the longest duration of each intervention was retained for analysis.

### 388 **Attention analysis**

389 Age-specific attention patterns were analyzed by extracting attention scores from all heads in the  
390 final transformer layer. We computed mean attention scores for each CpG site across samples  
391 within defined age groups (<20, 20-60, >60 years). CpG sites with significantly different atten-  
392 tion scores between age groups were identified using two-sided t-tests with Benjamini-Hochberg  
393 correction.

394 For CpG sites showing differential attention patterns between age groups, we performed Gene  
395 Ontology (GO) and Reactome gene set enrichment analysis using MethylGSA<sup>51</sup>.

### 396 **Statistical analysis**

397 All statistical tests were two-sided unless otherwise specified. Error bars in the figures represent  
398 standard deviation across samples. Sample sizes and statistical methods are specified in figure  
399 legends.

### 400 **Code availability**

401 The MethylGPT code and pre-trained models will be made available on github upon publication.

### 402 **Data availability**

403 All methylation data used in this study are available through EWAS Data Hub, GEO, and  
404 Clockbase. Processed datasets and analysis scripts will be deposited to github upon publication.




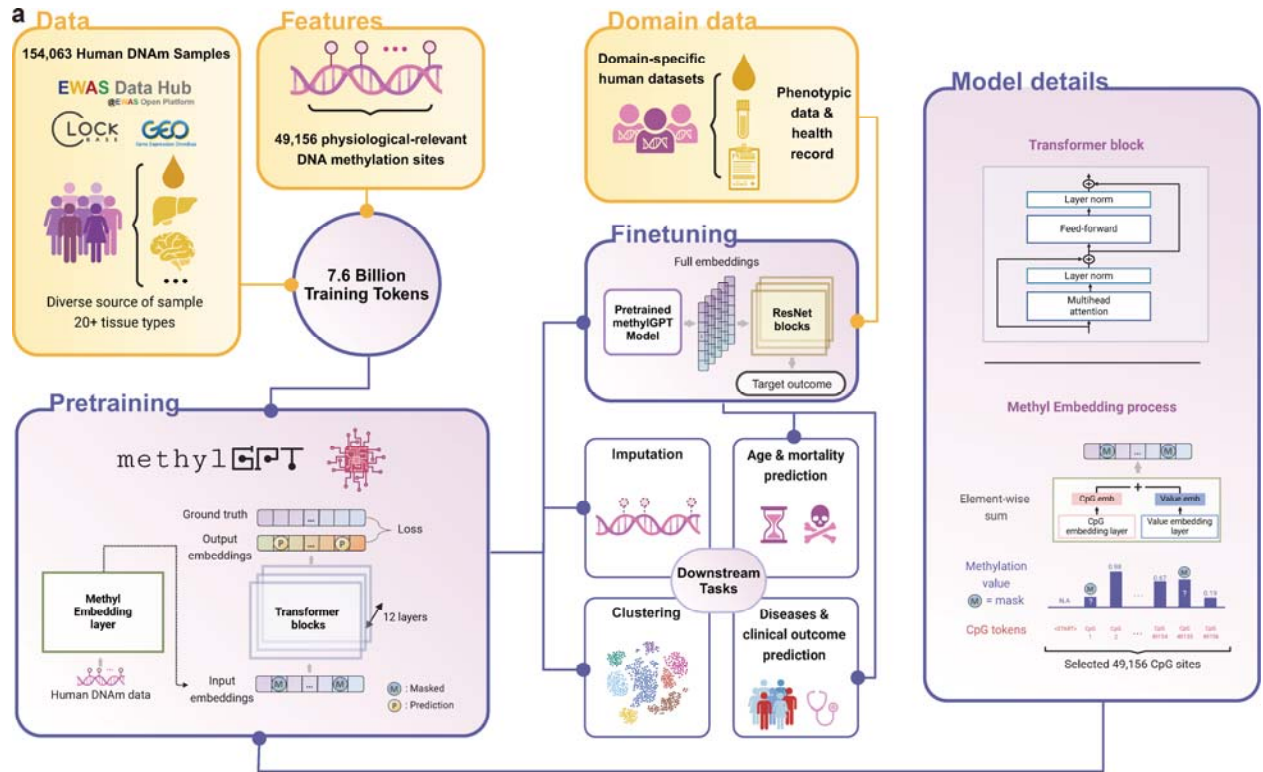
## 405 **Acknowledgments**

406 This work was supported by grants from the National Institute on Aging and Hevolution to  
407 VNG. It was also supported by the James Fickel Foundation. KY was supported by National In-  
408 stitute on Aging grant F99AG088431. We thank the Biomarkers of Aging Consortium and Gen-  
409 eration Scotland for providing access to their datasets. Generation Scotland received core support  
410 from the Chief Scientist Office of the Scottish Government Health Directorates [CZD/16/6] and  
411 the Scottish Funding Council [HR03006] and is currently supported by the Wellcome Trust  
412 [216767/Z/19/Z]. Genotyping of the Generation Scotland samples was carried out by the Genet-  
413 ics Core Laboratory at the Edinburgh Clinical Research Facility, University of Edinburgh, Scot-  
414 land and was funded by the Medical Research Council UK and the Wellcome Trust (Wellcome  
415 Trust Strategic Award “STratifying Resilience and Depression Longitudinally” (STRADL) Ref-  
416 erence 104036/Z/14/Z). The DNA methylation profiling and analysis was supported by Well-  
417 come Investigator Award 220857/Z/20/Z and Grant 104036/Z/14/Z (PI: Prof AM McIntosh) and  
418 through funding from NARSAD (Ref: 27404; awardee: Dr DM Howard) and the Royal College  
419 of Physicians of Edinburgh (Sim Fellowship; Awardee: Prof HC Whalley)

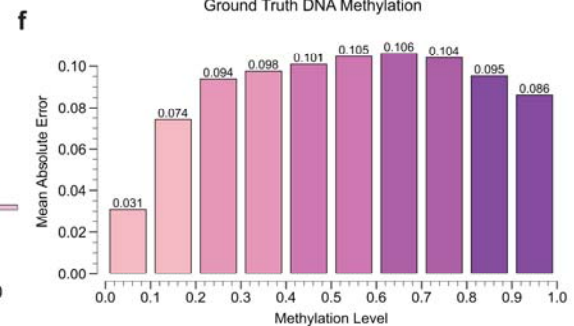
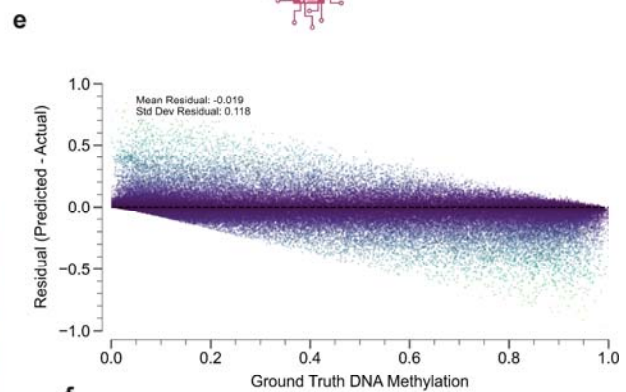
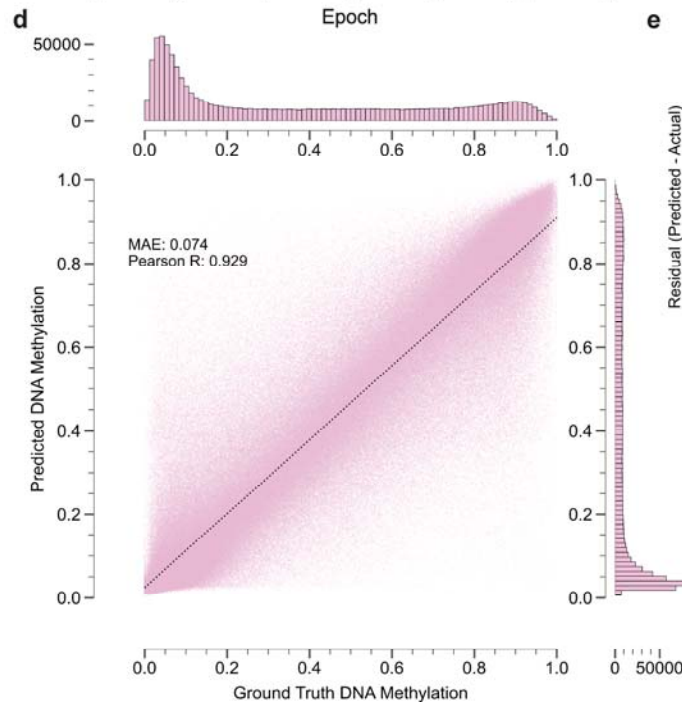

## 420 **Author contributions**

421 KY conceived the idea and designed the study. KY, SL, and HL collected initial data. KY, JS,  
422 and HC designed the model and performed pre-training. KY and YZ performed model fine-  
423 tuning and analysis. XC and AE helped with data analysis. DLM, REM, and MM helped with  
424 human cohort data curation. KY, JS, and HC wrote the manuscript. All authors edited and con-  
425 tributed to the manuscript. BW and VNG supervised the study.

426 **Figures**

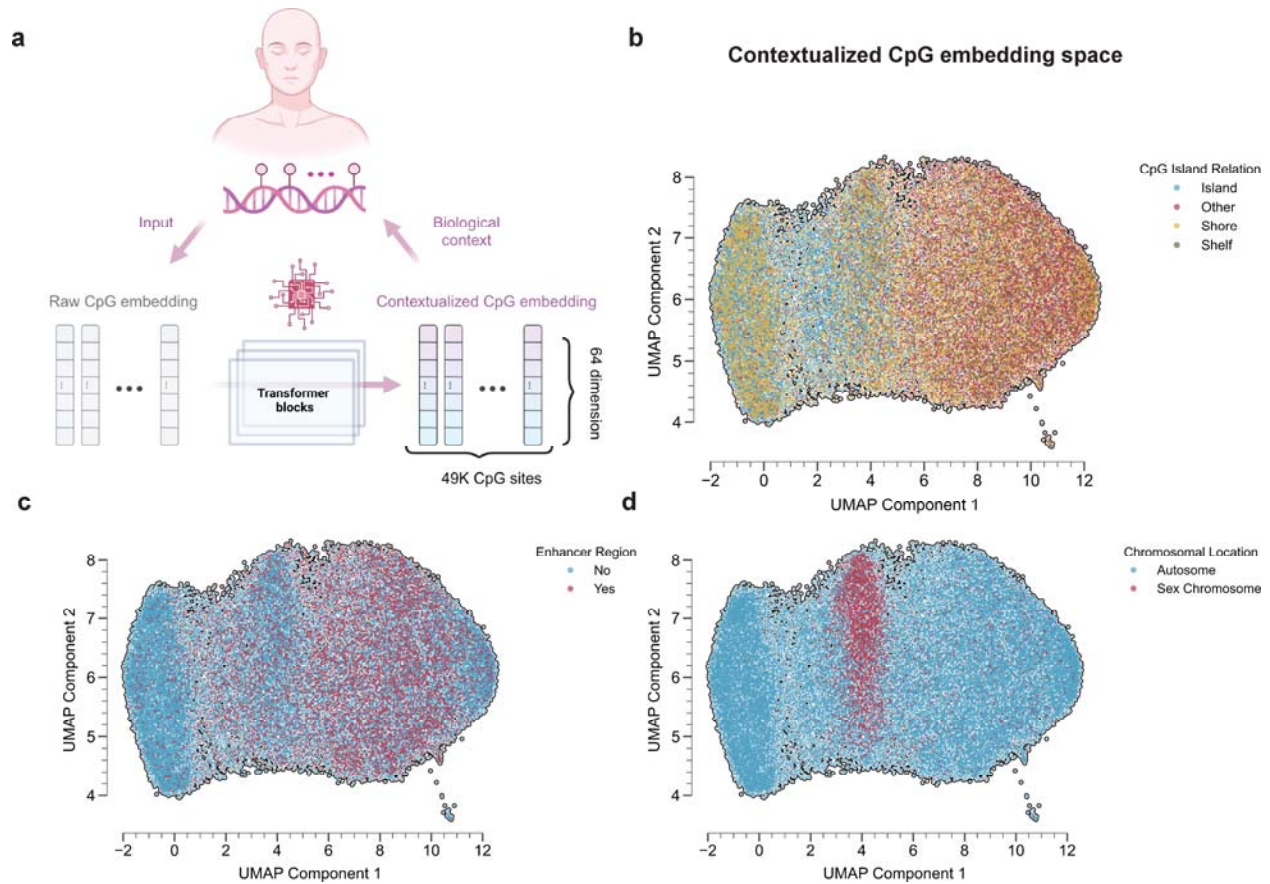


**c** Imputing missing/masked DNA methylation value



428 **Figure 1. Overview of MethylGPT architecture and performance. a.** Model architecture dia-  
429 gram showing data flow from 154,063 human DNAm samples through feature extraction (49,156  
430 CpG sites) to generate 7.6 billion training tokens. Components, including transformer block de-  
431 tails and the methyl embedding process, are highlighted. **b.** Training curve showing MLM loss  
432 over epochs, with train and validation MSE trajectories converging at epoch 10 (Best Model Test  
433 MSE: 0.014). **c.** Illustration of the imputing process for missing/masked DNA methylation val-  
434 ues using MethylGPT. **d.** Joint density plot showing the correlation between predicted and  
435 ground truth DNA methylation values (Pearson R: 0.929, MAE: 0.074). **e.** Residual plot showing  
436 prediction errors across different methylation levels. **f.** Bar plot showing mean absolute error  
437 across different methylation levels (0.0-1.0).

438

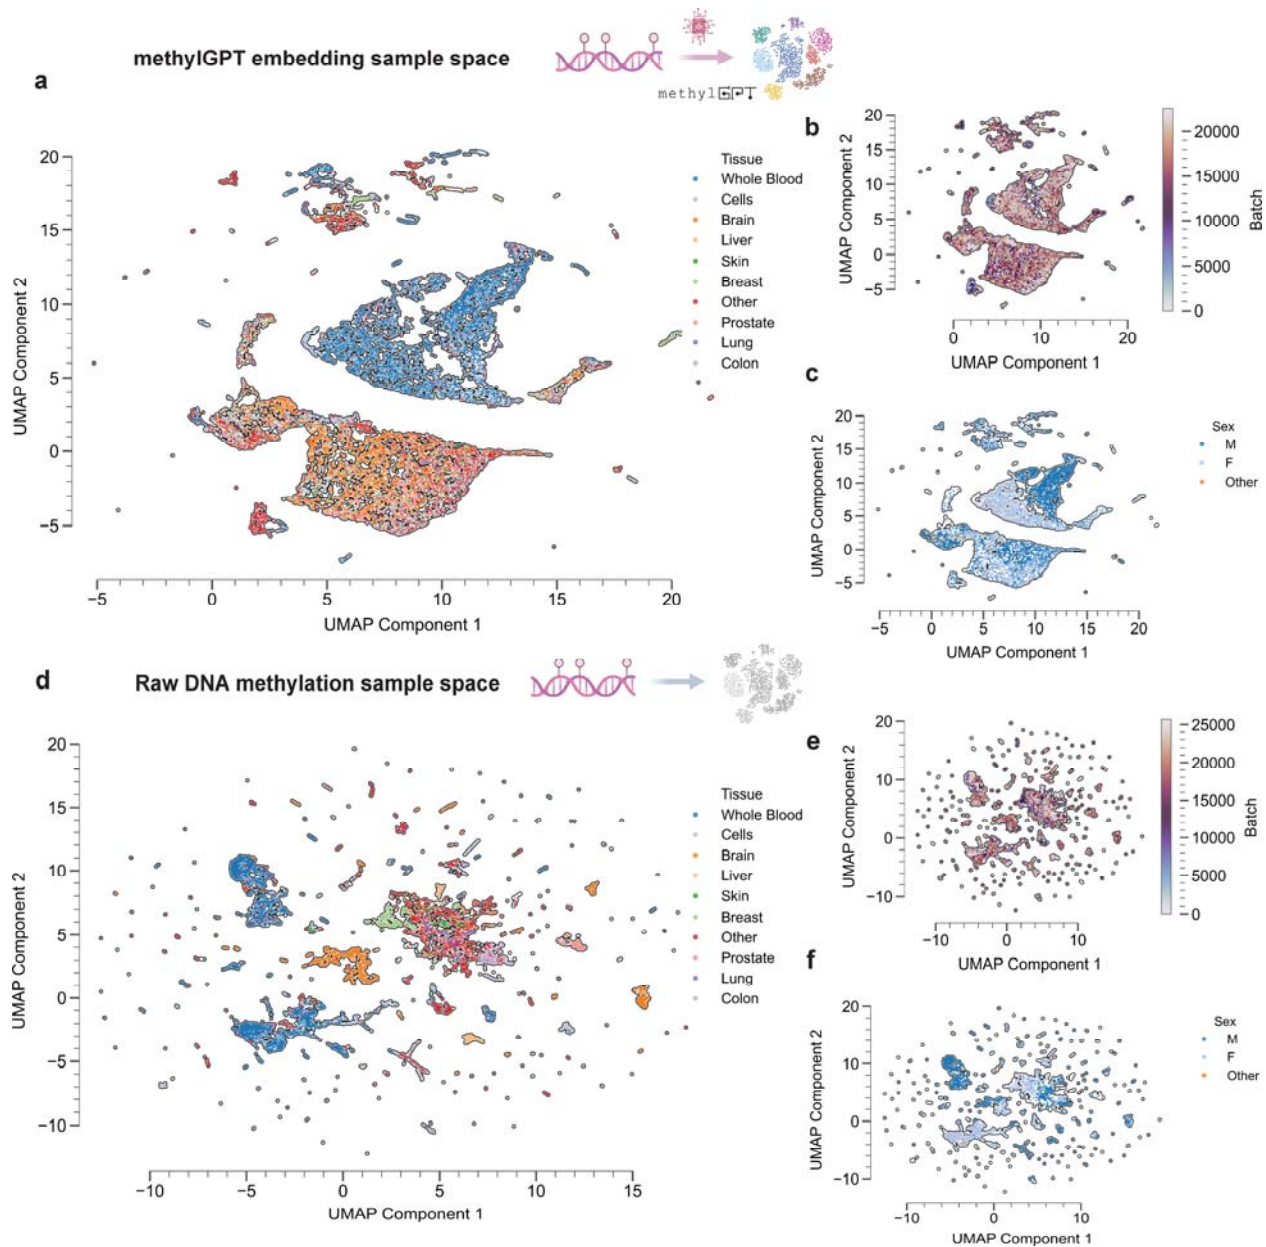


439

440 **Figure 2. Analysis of contextualized CpG embedding space.** **a.** Schematic illustration of the  
441 CpG embedding process, showing the transformation from raw CpG input to contextualized  
442 embeddings through transformer blocks. **b.** UMAP visualization of 49K CpG sites colored by  
443 CpG island relationship (Island, Shore, Shelf, Other). **c.** UMAP plot highlighting enhancer re-  
444 gions (Yes/No) in the embedding space. **d.** UMAP visualization showing the separation of CpG  
445 sites by chromosomal location, with distinct clustering of sex chromosomes and autosomes.

446

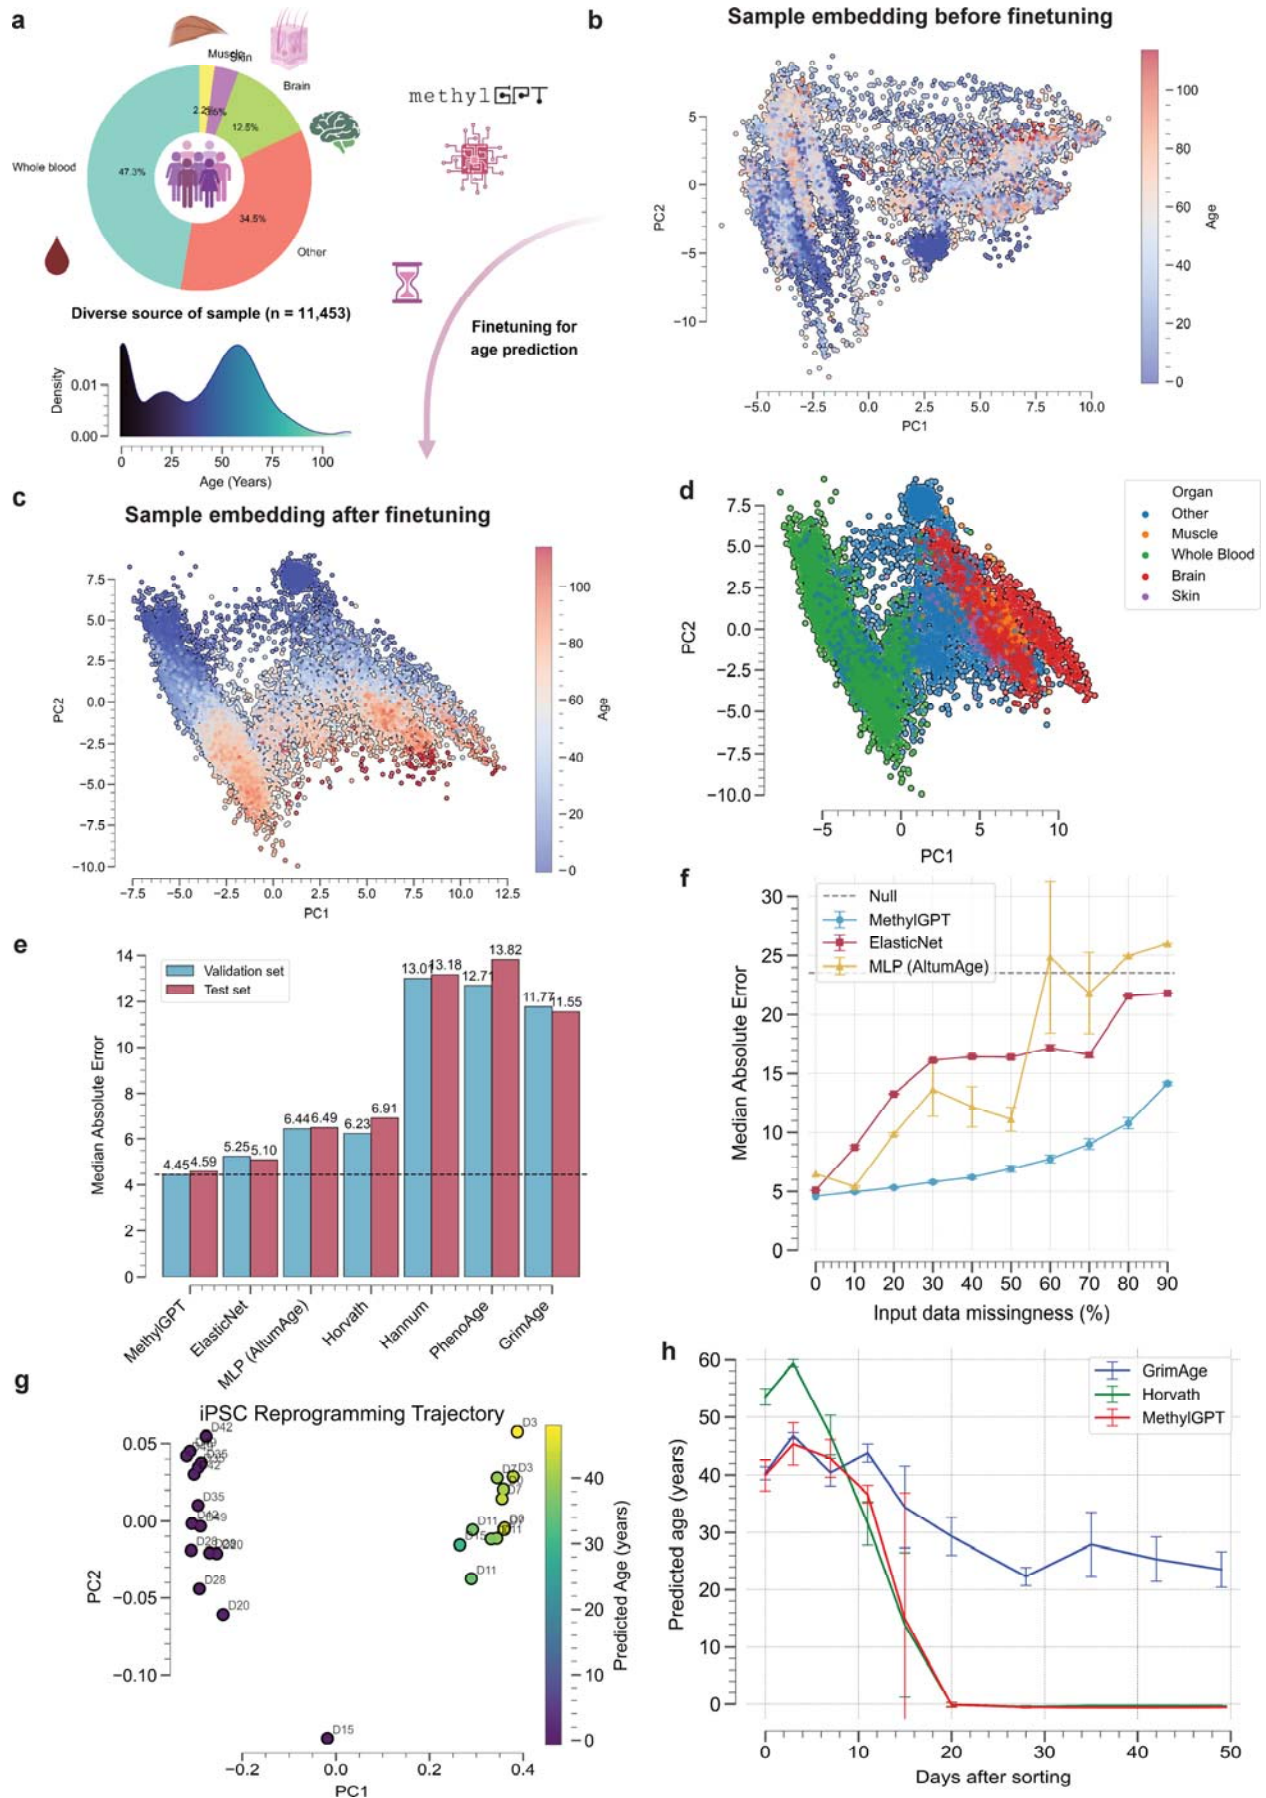




447

448 **Figure 3. Sample-level embedding analysis.** **a.** UMAP visualization of MethylGPT sample  
449 embeddings colored by tissue type, showing distinct clustering of major tissue types including  
450 whole blood, brain, liver, and skin. **b.** Sample density plot of the embedding space highlighting  
451 minimal batch effects. **c.** Sex-specific clustering in the embedding space, displaying a clear separation  
452 between male and female samples. **d-f.** Comparative analysis of raw DNA methylation  
453 sample embeddings, showing less distinct clustering by tissue type (**d**), more pronounced batch  
454 effects (**e**), and weaker separation by sex (**f**).

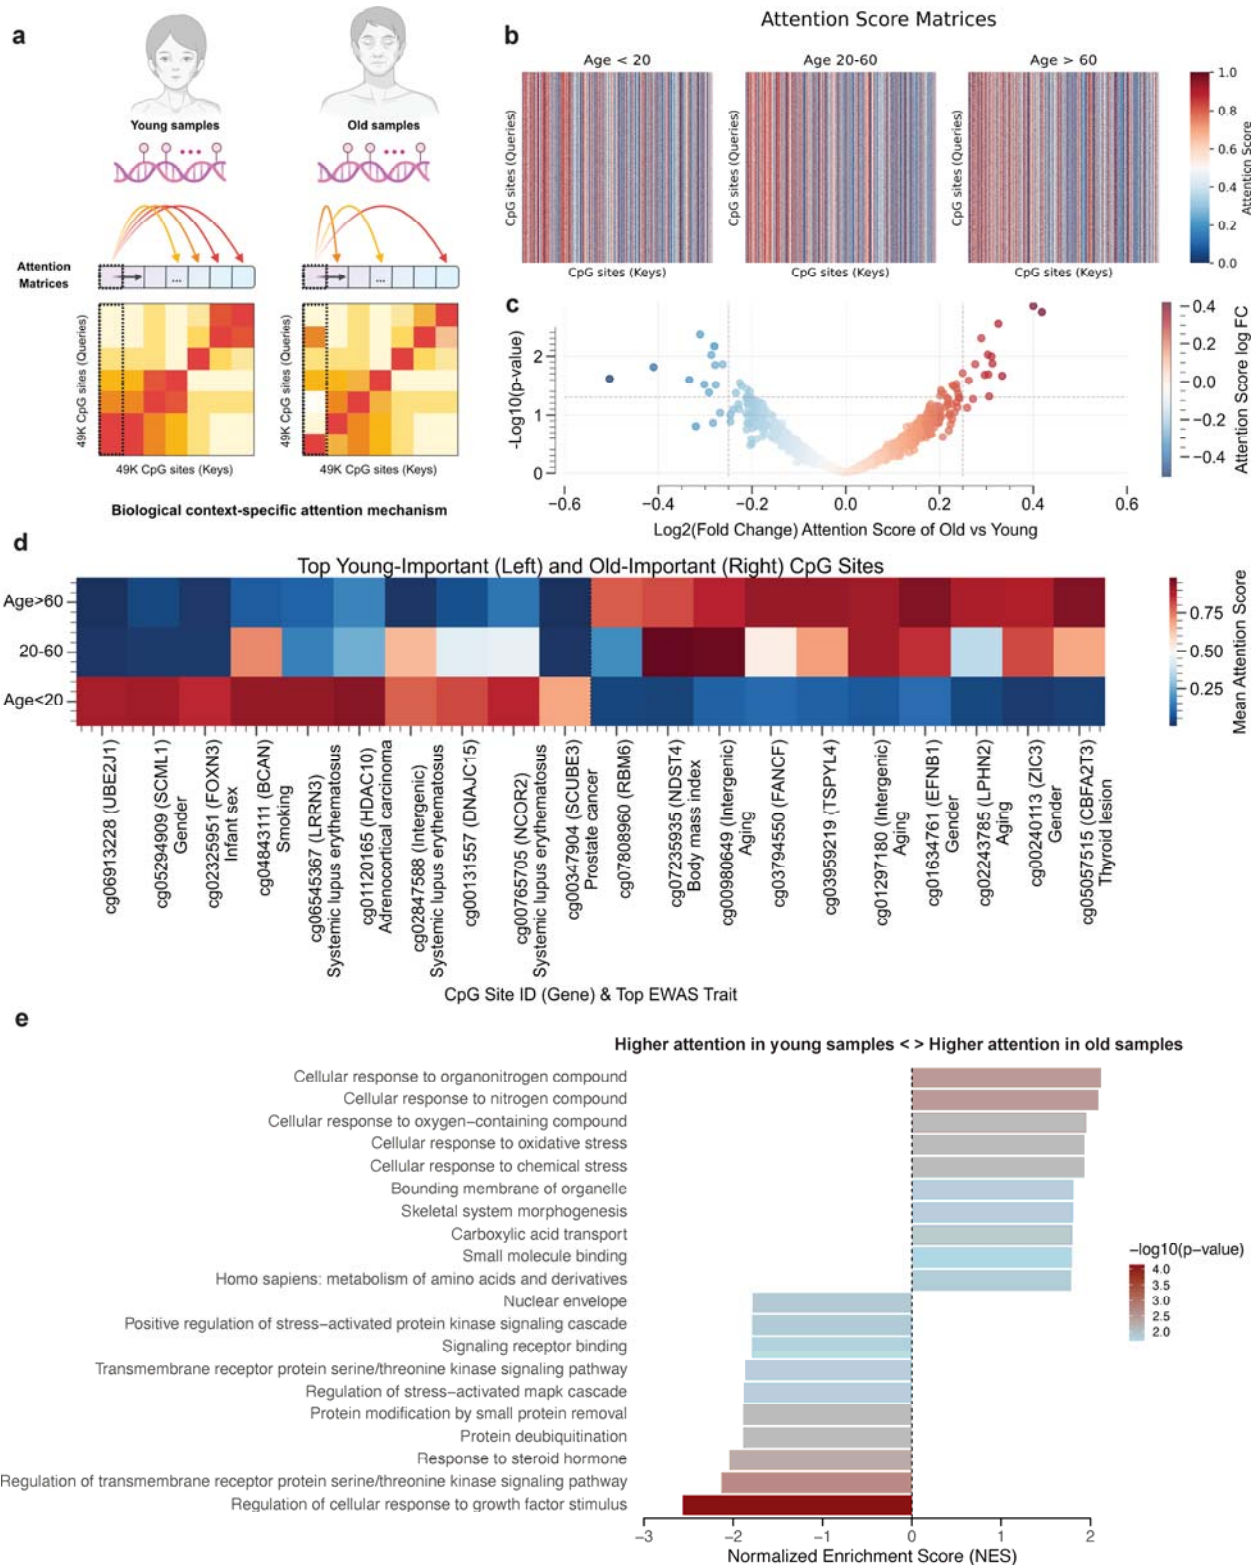
455





457 **Figure 4. Age prediction performance and robustness analysis.** **a.** Sample composition pie  
458 chart showing tissue distribution within the age finetuning dataset (n=11,453) and age distribu-  
459 tion density plot. **b.** PCA visualization of sample embeddings before fine-tuning, colored by age.  
460 **c.** Sample embeddings after fine-tuning for age prediction, showing enhanced age-related organ-  
461 ization. **d.** Tissue-specific clustering was maintained after fine-tuning. **e.** Benchmark comparison  
462 of age prediction performance across different methods on validation and test datasets. Median  
463 Absolute Errors are annotated. **f.** Robustness analysis showing prediction performance under in-  
464 creasing levels of missing data (10-90%) on test dataset for different methods. **g.** Principal com-  
465 ponent analysis of MethylGPT embeddings during iPSC reprogramming, colored by predicted  
466 age, showing progressive trajectory towards younger methylation states. **h.** Comparison of pre-  
467 dicted age trajectories during iPSC reprogramming across different epigenetic clocks (GrimAge,  
468 Horvath's clock) and MethylGPT, demonstrating consistent detection of rejuvenation effects. Er-  
469 ror bars represent standard deviation across replicate samples.

470



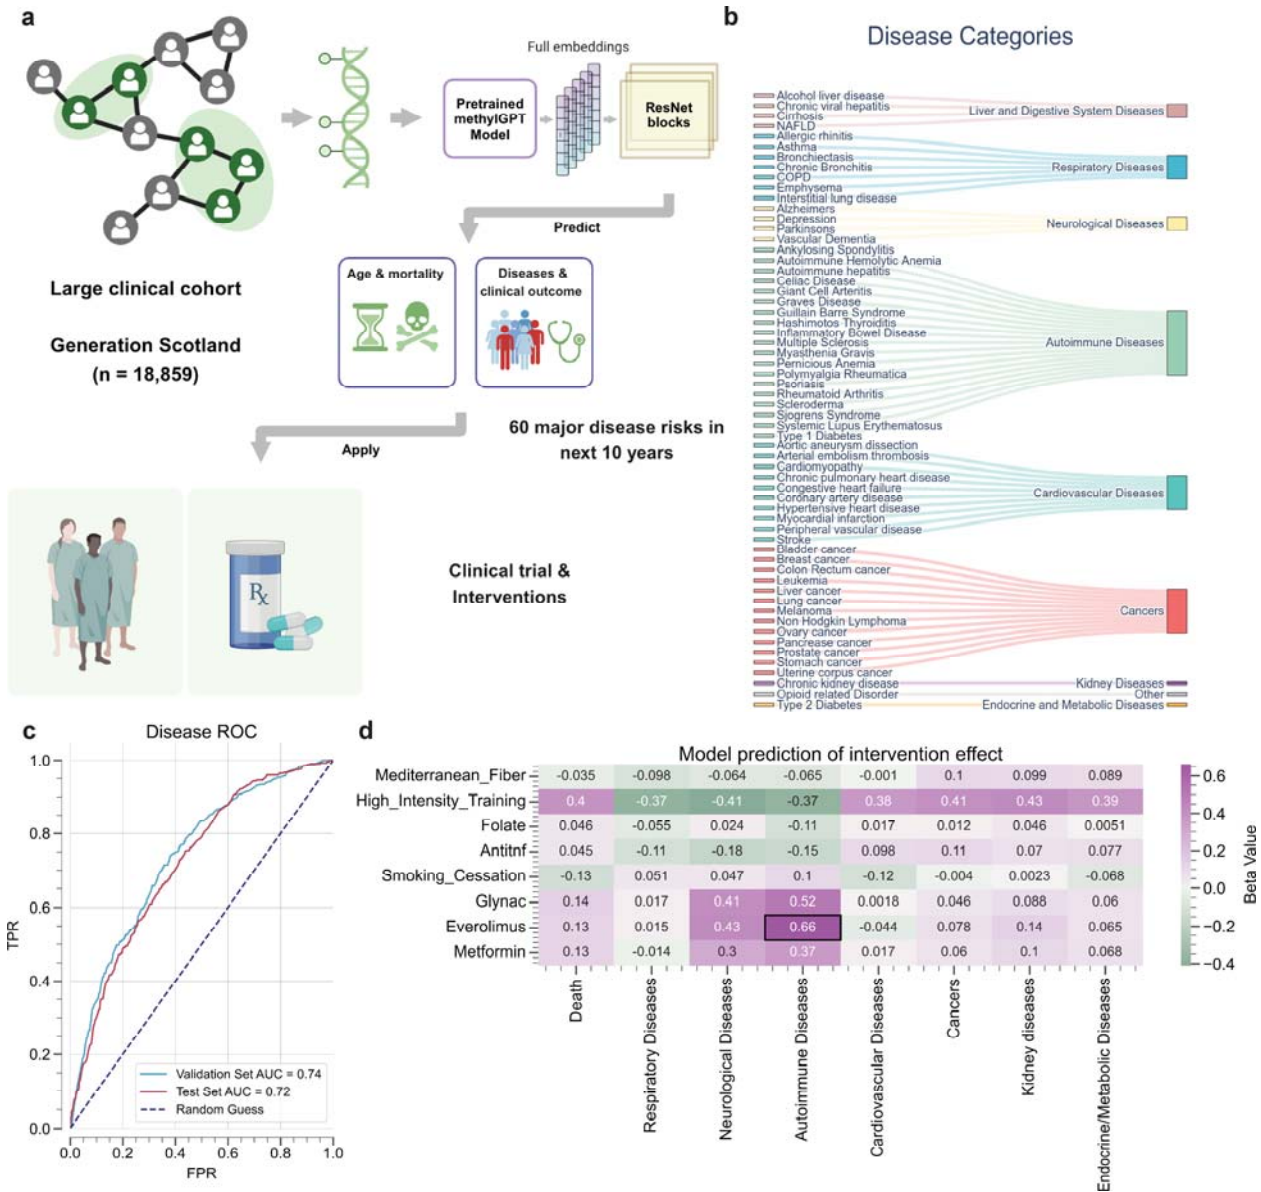
471

472 **Figure 5. Age-specific attention mechanism analysis. a.** Schematic comparison of attention

473 patterns between young and old samples, showing differential CpG site interactions. **b.** Attention

474 score matrices across three age groups (<20, 20-60, >60 years), revealing age-specific patterns.  
475 **c.** Volcano plot of log p-values versus differential mean attention scores identifies a few influen-  
476 tial CpG sites distinguishing the attention pattern of young and old groups. **d.** Heatmap of top  
477 young-important (left) and old-important (right) CpG sites, annotated with associated genes and  
478 EWAS traits, demonstrating age-specific methylation signatures. **e.** Functional enrichment analy-  
479 sis of top young-important (left) and old-important (right) CpG sites, with bars colored according  
480 to -log p-values.

481



482

483 **Figure 6. Disease risk prediction and intervention effects using MethylGPT.** **a.** Schematic  
 484 overview of the disease prediction pipeline using Generation Scotland cohort (n = 18,859). The  
 485 pretrained MethylGPT model processes methylation profiles through ResNet blocks to predict  
 486 age, mortality, and disease risks, which can then be applied to evaluate clinical interventions. **b.**  
 487 Visualization of 60 major diseases organized into disease categories (Liver and Digestive System  
 488 Diseases, Respiratory Diseases, Neurological Diseases, Autoimmune Diseases, Cardiovascular  
 489 Diseases, Cancers, Kidney Diseases, and Endocrine and Metabolic Diseases). **c.** Receiver Oper-  
 490 ating Characteristic (ROC) curves showing the overall performance of MethylGPT disease pre-  
 491 diction model (seven disease classes and overall mortality) on validation (AUC = 0.736) and test

492 (AUC = 0.720) sets. **d.** Heatmap showing predicted effects ( $\beta$  values) of eight different interven-  
493 tions on disease risks across major disease categories (total n=183): Mediterranean fiber (n=36),  
494 high-intensity training (n=5), folate supplementation (n=43), anti-TNF therapy (n=59), smoking  
495 cessation (n=16), glyNAC (n=8), everolimus (n=8), and metformin (n=8). Each intervention in-  
496 cluded an intra-group control as part of the trial design. For phased interventions, only the long-  
497 est duration timepoint was analyzed. Color scale represents effect size, with purple indicating  
498 positive effects (risk reduction) and green indicating negative effects (risk increase). Black box  
499 highlights significant effects. Values represent effect size from the Cohen's d.

500

## 501 **References**

- 502 1. Jones, P. A. Functions of DNA methylation: islands, start sites, gene bodies and beyond. *Nat.*  
503 *Rev. Genet.* **13**, 484–492 (2012).
- 504 2. Weigert, R. *et al.* Dynamic antagonism between key repressive pathways maintains the pla-  
505 cental epigenome. *Nat. Cell Biol.* **25**, 579–591 (2023).
- 506 3. Deniz, Ö., Frost, J. M. & Branco, M. R. Regulation of transposable elements by DNA modi-  
507 fications. *Nat. Rev. Genet.* **20**, 417–431 (2019).
- 508 4. Levenson, V. V. DNA methylation as a universal biomarker. *Expert Rev. Mol. Diagn.* **10**,  
509 481 (2010).
- 510 5. Cappozzo, A. *et al.* A blood DNA methylation biomarker for predicting short-term risk of  
511 cardiovascular events. *Clin. Epigenetics* **14**, 121 (2022).
- 512 6. Horvath, S. DNA methylation age of human tissues and cell types. *Genome Biol.* **14**, R115  
513 (2013).
- 514 7. Belsky, D. W. *et al.* DunedinPACE, a DNA methylation biomarker of the pace of aging.  
515 *eLife* **11**, e73420 (2022).
- 516 8. Lu, A. T. *et al.* DNA methylation GrimAge strongly predicts lifespan and healthspan. *Aging*  
517 **11**, 303–327 (2019).
- 518 9. Ying, K. *et al.* Causality-enriched epigenetic age uncouples damage and adaptation. *Nat. Ag-*  
519 *ing* 1–16 (2024) doi:10.1038/s43587-023-00557-0.
- 520 10. Ying, K. *et al.* High-dimensional Ageome Representations of Biological Aging across Func-  
521 tional Modules. 2024.09.17.613599 Preprint at <https://doi.org/10.1101/2024.09.17.613599>  
522 (2024).
- 523 11. Moqri, M. *et al.* Biomarkers of aging for the identification and evaluation of longevity inter-  
524 ventions. *Cell* **186**, 3758–3775 (2023).

- 525 12. Moqri, M. *et al.* Validation of biomarkers of aging. *Nat. Med.* 1–13 (2024)  
526 doi:10.1038/s41591-023-02784-9.
- 527 13. Kreibich, E., Kleinendorst, R., Barzaghi, G., Kaspar, S. & Krebs, A. R. Single-molecule  
528 footprinting identifies context-dependent regulation of enhancers by DNA methylation. *Mol.*  
529 *Cell* **83**, 787–802.e9 (2023).
- 530 14. Greenberg, M. V. C. & Bourc’his, D. The diverse roles of DNA methylation in mammalian  
531 development and disease. *Nat. Rev. Mol. Cell Biol.* **20**, 590–607 (2019).
- 532 15. Olecka, M. *et al.* Nonlinear DNA methylation trajectories in aging male mice. *Nat. Commun.*  
533 **15**, 3074 (2024).
- 534 16. Ross, J. P. *et al.* Batch-effect detection, correction and characterisation in Illumina  
535 HumanMethylation450 and MethylationEPIC BeadChip array data. *Clin. Epigenetics* **14**, 58  
536 (2022).
- 537 17. Vaswani, A. *et al.* Attention Is All You Need. Preprint at  
538 <https://doi.org/10.48550/arXiv.1706.03762> (2017).
- 539 18. Lin, Z. *et al.* Evolutionary-scale prediction of atomic-level protein structure with a language  
540 model. *Science* **379**, 1123–1130 (2023).
- 541 19. Hayes, T. *et al.* Simulating 500 million years of evolution with a language model.  
542 2024.07.01.600583 Preprint at <https://doi.org/10.1101/2024.07.01.600583> (2024).
- 543 20. Jumper, J. *et al.* Highly accurate protein structure prediction with AlphaFold. *Nature* **596**,  
544 583–589 (2021).
- 545 21. Abramson, J. *et al.* Accurate structure prediction of biomolecular interactions with  
546 AlphaFold 3. *Nature* **630**, 493–500 (2024).



- 547 22. Avsec, Ž. *et al.* Effective gene expression prediction from sequence by integrating long-  
548 range interactions. *Nat. Methods* **18**, 1196–1203 (2021).
- 549 23. Nguyen, E. *et al.* Sequence modeling and design from molecular to genome scale with Evo.  
550 2024.02.27.582234 Preprint at <https://doi.org/10.1101/2024.02.27.582234> (2024).
- 551 24. Theodoris, C. V. *et al.* Transfer learning enables predictions in network biology. *Nature* **618**,  
552 616–624 (2023).
- 553 25. Cui, H. *et al.* scGPT: toward building a foundation model for single-cell multi-omics using  
554 generative AI. *Nat. Methods* **21**, 1470–1480 (2024).
- 555 26. Hao, M. *et al.* Large Scale Foundation Model on Single-cell Transcriptomics.  
556 2023.05.29.542705 Preprint at <https://doi.org/10.1101/2023.05.29.542705> (2023).
- 557 27. Galkin, F. *et al.* Precious3GPT: Multimodal Multi-Species Multi-Omics Multi-Tissue Trans-  
558 former for Aging Research and Drug Discovery. 2024.07.25.605062 Preprint at  
559 <https://doi.org/10.1101/2024.07.25.605062> (2024).
- 560 28. Xiong, Z. *et al.* EWAS Data Hub: a resource of DNA methylation array data and metadata.  
561 *Nucleic Acids Res.* **48**, D890–D895 (2020).
- 562 29. Ying, K. *et al.* ClockBase<sup>2</sup>: a comprehensive platform for biological age profiling in human  
563 and mouse. Preprint at <https://doi.org/10.1101/2023.02.28.530532> (2023).
- 564 30. Battram, T. *et al.* The EWAS Catalog: a database of epigenome-wide association studies.  
565 Preprint at <https://doi.org/10.12688/wellcomeopenres.17598.2> (2022).
- 566 31. de Lima Camillo, L. P., Lapierre, L. R. & Singh, R. A pan-tissue DNA-methylation epigenet-  
567 ic clock based on deep learning. *Npj Aging* **8**, 1–15 (2022).
- 568 32. Zou, H. & Hastie, T. Regularization and Variable Selection Via the Elastic Net. *J. R. Stat.*  
569 *Soc. Ser. B Stat. Methodol.* **67**, 301–320 (2005).

- 570 33. Horvath, S. *et al.* Epigenetic clock for skin and blood cells applied to Hutchinson Gilford  
571 Progeria Syndrome and ex vivo studies. *Aging* **10**, 1758–1775 (2018).
- 572 34. Ohnuki, M. *et al.* Dynamic regulation of human endogenous retroviruses mediates factor-  
573 induced reprogramming and differentiation potential. *Proc. Natl. Acad. Sci.* **111**, 12426–  
574 12431 (2014).
- 575 35. Ospina, N. S. & Papaleontiou, M. Thyroid nodule evaluation and management in older  
576 adults: A review of practical considerations for clinical endocrinologists. *Endocr. Pract. Off.*  
577 *J. Am. Coll. Endocrinol. Am. Assoc. Clin. Endocrinol.* **27**, 261 (2021).
- 578 36. Kao, F.-C. *et al.* Long-Term Use of Immunosuppressive Agents Increased the Risk of Frac-  
579 tures in Patients with Autoimmune Diseases: An 18-Year Population-Based Cohort Study.  
580 *Biomedicines* **11**, 2764 (2023).
- 581 37. Zhang, S. *et al.* DNA methylation profiling to determine the primary sites of metastatic can-  
582 cers using formalin-fixed paraffin-embedded tissues. *Nat. Commun.* **14**, 5686 (2023).
- 583 38. Thompson, M. *et al.* Methylation risk scores are associated with a collection of phenotypes  
584 within electronic health record systems. *Npj Genomic Med.* **7**, 1–11 (2022).
- 585 39. Cheng, Y. *et al.* Development and validation of DNA methylation scores in two European  
586 cohorts augment 10-year risk prediction of type 2 diabetes. *Nat. Aging* **3**, 450–458 (2023).
- 587 40. Mitchell, C., Schneper, L. M. & Notterman, D. A. DNA methylation, early life environment,  
588 and health outcomes. *Pediatr. Res.* **79**, 212–219 (2016).
- 589 41. Hedman, Å. K., Zilmer, M., Sundström, J., Lind, L. & Ingelsson, E. DNA methylation pat-  
590 terns associated with oxidative stress in an ageing population. *BMC Med. Genomics* **9**, 72  
591 (2016).

- 592 42. Lee, K. *et al.* Deduplicating Training Data Makes Language Models Better. Preprint at  
593 <https://doi.org/10.48550/arXiv.2107.06499> (2022).
- 594 43. Devlin, J., Chang, M.-W., Lee, K. & Toutanova, K. BERT: Pre-training of Deep Bidirection-  
595 al Transformers for Language Understanding. Preprint at  
596 <https://doi.org/10.48550/arXiv.1810.04805> (2019).
- 597 44. Dao, T., Fu, D. Y., Ermon, S., Rudra, A. & Ré, C. FlashAttention: Fast and Memory-  
598 Efficient Exact Attention with IO-Awareness. Preprint at  
599 <https://doi.org/10.48550/arXiv.2205.14135> (2022).
- 600 45. Sokolowska, K. E. *et al.* Identified in blood diet-related methylation changes stratify liver  
601 biopsies of NAFLD patients according to fibrosis grade. *Clin. Epigenetics* **14**, 157 (2022).
- 602 46. Pilotto, A. *et al.* Human skeletal muscle possesses an epigenetic memory of high intensity  
603 interval training. Preprint at <https://doi.org/10.1101/2024.05.30.596458> (2024).
- 604 47. Irwin, R. E. *et al.* A randomized controlled trial of folic acid intervention in pregnancy high-  
605 lights a putative methylation-regulated control element at *ZFP57*. *Clin. Epigenetics* **11**, 31  
606 (2019).
- 607 48. Mishra, N. *et al.* Longitudinal multi-omics analysis identifies early blood-based predictors of  
608 anti-TNF therapy response in inflammatory bowel disease. *Genome Med.* **14**, 110 (2022).
- 609 49. Shang, J., Nie, X., Qi, Y., Zhou, J. & Qi, Y. Short-term smoking cessation leads to a univer-  
610 sal decrease in whole blood genomic DNA methylation in patients with a smoking history.  
611 *World J. Surg. Oncol.* **21**, 227 (2023).
- 612 50. Abhimanyu *et al.* TCA metabolism regulates DNA hypermethylation in LPS and *Mycobac-*  
613 *terium tuberculosis* –induced immune tolerance. *Proc. Natl. Acad. Sci.* **121**, e2404841121  
614 (2024).

615 51. Ren, X. & Kuan, P. F. methylGSA: a Bioconductor package and Shiny app for DNA methyl-  
616 ation data length bias adjustment in gene set testing. *Bioinformatics* **35**, 1958–1959 (2019).

617