

# PM1: A FOUNDATION MODEL FUSING GENOTYPE, PHENOTYPE, AND IMAGE FOR PRECISION MEDICINE

Anonymous authors

Paper under double-blind review

## ABSTRACT

Precision medicine aims to personalize disease prevention, prediction, and diagnosis by leveraging genomic patient data. Although patient genomes provide valuable predictive insight, they cannot capture the full complexity of an individual’s health. Integrating genomics with additional patient data modalities, such as clinical phenotypes and medical imaging, enables more accurate and comprehensive disease modeling. We introduce **PM1**, a multimodal foundation model trained on genomic data from 438,668 individuals linked to 3,421 clinical and lifestyle traits and 211,416 retinal fundus photographs drawn from the UK Biobank and EyePACS cohorts. PM1 couples modality-specific encoders with a transformer encoder trained with an information noise-contrastive estimation objective that fuses modalities into a joint latent space, plus generative modality decoders for cross-modal reconstruction and synthesis. A token-level masking schedule lets PM1 use participants with *any* subset of modalities (in UK Biobank only  $\approx 6\%$  have all three), substantially expanding effective training data. Joint modeling of retinal images, clinical traits, and genomic data surpasses single-modality and multimodal baselines. PM1 enables cross-modal genotype inference, raises predictive performance for retinal diseases and systemic conditions, and supports conditioned single nucleotide polymorphism sequence and retinal image generation. As a group-level validation, a GWAS on PM1’s image-conditioned fusion embeddings recovers genome-wide significant HERC2 pigmentation variants.

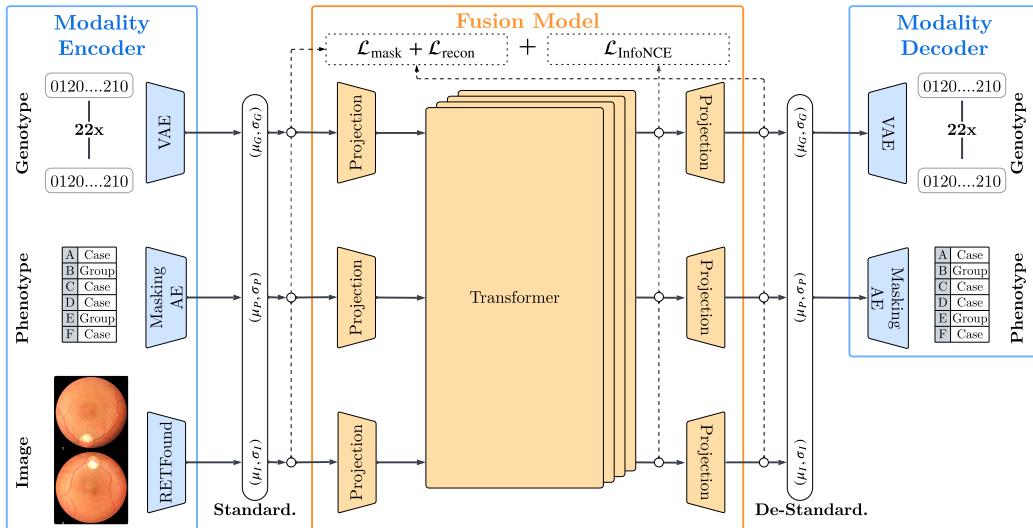


Figure 1: **Overview of PM1 Model Architecture.** Encoders for genotype, phenotypes, and retinal images emit tokens mapped to a shared space, fused by a transformer for alignment and prediction, with modality decoders reconstructing inputs and enabling synthesis when inputs are missing.

054 

## 1 INTRODUCTION

056 Following the first demonstration of generalist multimodal biomedical foundation models in 2023  
 057 (Tu et al., 2024; Li et al., 2023; Zhang et al., 2023; Theodoris et al., 2023), mostly spanning clinical  
 058 text and imaging data, these systems still omit the most heritable axis of human variation—whole-  
 059 genome sequence—hence ignoring information that underlies virtually every complex trait. Inte-  
 060 grating high-dimensional genomic data into multimodal models poses unique challenges: genomic  
 061 sequences are both longer and more structured than typical language or image inputs, exhibiting  
 062 strong correlations through linkage disequilibrium (LD). Meanwhile, biobank-scale resources like  
 063 the UK Biobank (UKB) (Sudlow et al., 2015) provide an invaluable resource for studying the rela-  
 064 tionship between genomes, clinical traits, and medical images on a massive scale, offering millions  
 065 of paired single-nucleotide polymorphisms (SNPs), thousands of clinical records, and retinal fun-  
 066 dus photographs across hundreds of thousands of individuals—providing a unique sandbox for truly  
 067 multimodal precision-medicine models (Szustakowski et al., 2021; DeBoever et al., 2020).

068 We introduce **PM1**<sup>1</sup>, a multimodal foundation model trained on complete biobank-scale data that  
 069 fuses genomic, clinical, and retinal imaging signals. PM1 consists of a three-stream architecture with  
 070 modality-specific encoders and generative decoders, fused by a transformer encoder (Vaswani et al.,  
 071 2017). The model is trained with masked language modeling (Devlin et al., 2019) and an information  
 072 noise-contrastive estimation objective (InfoNCE) (Oord et al., 2018) to encourage fusion of the  
 073 modality embeddings in a shared latent space for inter- and intra-modality prediction. Our training  
 074 scheme allows using *all* available data despite modality-missingness, which is relevant in clinical  
 075 settings, yielding up to **16**× more usable samples than required in complete tri-modal coverage.

076 We show that PM1 consistently outperforms unimodal and multimodal baselines, and adding modal-  
 077 ities generally boosts phenotype prediction for clinical traits, as well as across the 22 autosomes.  
 078 Across phenotypes, PM1 outperforms the medical foundation model MedGemma (Sellergren et al.,  
 079 2025) by an average of 51% relative gain in ROC-AUC, and surpasses ContIG (Taleb et al., 2022)  
 080 by an average of 19% relative gain on the full test set.

081 Finally, we demonstrate the usefulness of PM1 for downstream tasks. Conditioning on PM1’s em-  
 082 beddings, synthetic retinal images can be generated, and a model trained only on generated images  
 083 scores ROC-AUC **0.850** in diabetic retinopathy classification on real held-out data. In a group-level  
 084 genome-wide association study (GWAS) analysis based on PM1, we are able to recover HERC2  
 085 pigmentation variants, and we achieve **>99%** accuracy in an ancestry-inference probe, showcasing  
 086 PM1’s ability to capture biological structure in the fused latent space.

087 

## 2 RELATED WORK

089 

### 2.1 REPRESENTATION LEARNING OF BIOMEDICAL DATA

091 **Genotype** Recent studies in representation learning of genomic data demonstrate that deep models  
 092 can capture structure in large-scale SNP data beyond linear approaches (Zhou et al., 2024a; Vivek  
 093 et al., 2023; Ausmees & Nettelblad, 2022; Geleta et al., 2023), potentially uncovering complex  
 094 associations, e.g., epistasis, missed by linear techniques such as PCA and penalized regression.  
 095 Variational autoencoders (VAEs) (Kingma & Welling, 2013; Rezende et al., 2014), a family of deep  
 096 generative models, are particularly popular for this application. VAEs estimate the mostly intractable  
 097 posterior probability  $p(z|x)$  over latent codes  $z$  given an observed data point  $x$  by maximizing the  
 098 evidence lower bound (ELBO), thereby learning a factorized latent representation of the generative  
 099 factors in the data. In particular, the KL divergence (Kingma & Welling, 2013; Rezende et al.,  
 100 2014) in the ELBO loss enforces continuity and smoothness in the latent space—properties that are  
 101 especially valuable in noisy and fully unsupervised settings such as biomedical data.

102 **Phenotype** Large-scale clinical datasets present unique challenges due to heterogeneous feature  
 103 types (binary, categorical, continuous), systematically missing, sparsity, and measurement noise.  
 104 Recent work tackles these issues with deep generative models and structure-aware embeddings. Au-  
 105 toComplete (An et al., 2023) trains a copy-mask denoising auto-encoder that jointly reconstructs bi-

107 <sup>1</sup>We publish the source code of PM1 at <https://github.com/anonymized-for-submission>  
 (Fig. 1) to ensure reproducibility and enable follow-up work.

108 nary and quantitative traits, outperforming matrix-factorization baselines such as SoftImpute (Hastie  
 109 et al., 2015) as well as earlier deep models HI-VAE (Nazabal et al., 2020) and GAIN (Yoon et al.,  
 110 2018). POPDx (Yang et al., 2023) blends disease-ontology structure with BERT-derived text embed-  
 111 dings, and uses a bilinear network to assign multilabel phenotypes to individuals, retaining accuracy  
 112 for rare or even unseen codes. Earlier generative approaches such as HI-VAE and GAIN remain com-  
 113 petitive on small cohorts but scale poorly to biobank dimensions. PM1 adopts an AutoComplete-  
 114 style encoder that is trained jointly with genomic and imaging modalities.

115 **Image** Deep learning has rapidly become the dominant paradigm for medical-image analysis, en-  
 116 abling fully automated disease detection, segmentation and even biomarker discovery across every  
 117 major radiologic modality. Recent surveys chart this trajectory, noting the transition from early  
 118 CNN classifiers to self-supervised Vision Transformers and masked-autoencoder pre-training that  
 119 now routinely outperform task-specific networks while requiring far fewer labeled images (Bahr  
 120 et al., 2024; Suganyadevi et al., 2021). In ophthalmology, central to our work—systematic reviews  
 121 show that modern fundus and OCT models not only diagnose ocular disorders such as diabetic  
 122 retinopathy or AMD, but can also predict systemic phenotypes (e.g., cardiovascular risk factors)  
 123 from retinal pixels alone, underscoring the retina’s value as a “window” into whole-body health  
 124 (Zhou et al., 2023).

## 126 2.2 FOUNDATION MODELS IN BIOMEDICINE

127 In parallel to advances in biobanking, there is growing interest in generalist biomedical AI systems  
 128 that can interpret multiple data modalities. Recent large-scale models have pushed the state of the  
 129 art in medical NLP and vision-and-language tasks. Med-PaLM (Singhal et al., 2023) and Med-  
 130 PaLM2 (Singhal et al., 2025), for example, are domain-specialized large language model (LLM)  
 131 derived from PaLM LLM (Chowdhery et al., 2022) and fine-tuned for medical QA tasks (Truhn  
 132 et al., 2024). Vision–language models such as LLaVA-Med (Li et al., 2023) and BiomedGPT  
 133 (Zhang et al., 2023) extend this capability to medical imaging using multimodal transformers and  
 134 instruction tuning. Med-PaLM M (Multimodal) (Tu et al., 2024) introduces an early prototype of a  
 135 multimodal model combining text, image, and genomics inputs, though its use of genomic data is  
 136 limited to variant calling. MedGemma (Sellergren et al., 2025) is a Gemma-3 (Team et al., 2025)  
 137 checkpoint that underwent additional pretraining on medical corpora (including retinal and derma-  
 138 tology images, histopathology, and radiology slices), offering open-source evaluation for clinical  
 139 vision-and-language tasks. Other efforts, such as ContIG (Taleb et al., 2022), explore joint image-  
 140 genetic embeddings via contrastive learning.

141 Despite these advances, current generalist biomedical AI models still have important gaps. Vi-  
 142 sion–language models like LLaVA-Med and BiomedGPT excel at image and text understanding,  
 143 but they do not incorporate structured clinical data or genomic sequences. On the other hand, text-  
 144 only LLMs like Med-PaLM 2 have encyclopedic medical knowledge and reasoning ability (Truhn  
 145 et al., 2024), yet remain blind to non-text modalities. Even Med-PaLM M, is a preliminary research  
 146 effort whose handling of genomics addresses only a narrow task. Unlike prior vision–language mod-  
 147 els or contrastive image–genetics models, PM1 directly ingests SNP-level arrays, fuses them with  
 148 phenotypes and retinal images, and supports cross-modal synthesis via modality-specific decoders,  
 149 while training to handle missing modalities at the token level.

## 150 3 DATA

151 Our primary source of data is the UK Biobank (Sudlow et al., 2015), which provides rich genomic,  
 152 phenotypic, and imaging records. It is a study of 500,000 adults aged 40-69 recruited across the UK  
 153 in 2006-2010, providing biospecimens for genome sequencing (Halldorsson et al., 2022; Backman  
 154 et al., 2021), who consented to longitudinal medical record linkage, and underwent extensive base-  
 155 line phenotyping including surveys, physical measurements, and biomarker assays. Uniquely, the  
 156 UKB also collected multimodal imaging on tens of thousands of participants, ranging from retinal  
 157 fundus photographs to brain MRI (Elliott et al., 2018; Gulshan et al., 2016). These resources enable  
 158 learning joint representations across modalities, which is the focus of PM1.

159 Specifically, we use unphased genomes of 438,668 UK participants with 658,720 SNP variants from  
 160 the 22 autosomal chromosomes, and 3,421 tabular phenotypes encompassing clinical and lifestyle

traits, as well as several image modalities. In particular, we use 211,416 Color Fundus Photography (CFP) images from the UKB Eye Imaging Study (Sudlow et al., 2015; Littlejohns et al., 2020; Keane et al., 2016), which provides comprehensive data for 69,600 samples, including CFP images paired with SNP sequences and extensive phenotypic information, including various retinal conditions such as diabetic retinopathy, age-related macular degeneration, glaucoma, and cataract. To balance the three modalities in our dataset, we also incorporate samples from the EyePACS repository (Gulshan et al., 2016), which, while lacking genetic sequences and comprehensive phenotypes, contains 44,351 samples with images of their retinas and annotations for diabetic retinopathy.

## 4 PM1 ARCHITECTURE AND TRAINING SCHEME

PM1 is trained in a two-stage scheme that separates within-modality representation learning from cross-modality fusion. The first stage focuses on modality-specific encoders and decoders, while the second stage leverages these frozen representations to train a fusion model that enables cross-modality prediction via masked token modeling. A schematic overview of the architecture is provided in Fig. 1.

### 4.1 MODALITY ENCODERS AND DECODERS

Multimodal fusion of genomic, clinical, and imaging data requires robust learned representations for each modality. Let each  $j$ -th sample be characterized by a triplet of modalities  $x_j = (x_j^{(G)}, x_j^{(P)}, x_j^{(I)})$ , corresponding to genotype, phenotype, and image data. For each modality input  $x_j^{(m)} \in \mathcal{X}_m$ , we define an encoder-decoder pair  $(\mathcal{E}_m, \mathcal{D}_m)$ , where  $\mathcal{E}_m : \mathcal{X}_m \rightarrow \mathbb{R}^{T_m \times d_m}$  maps raw modality-specific data to a sequence of latent  $T_m$  tokens of dimensionality  $d_m$ , and  $\mathcal{D}_m : \mathbb{R}^{T_m \times d_m} \rightarrow \mathcal{X}_m$  reconstructs the original input in the corresponding modality space. Once trained, the parameters of  $\mathcal{E}_m$  and  $\mathcal{D}_m$  are frozen, and only the encoder forward passes  $\mathcal{E}_m(x_j^{(m)}) = z_j^{(m)}$  are used in fusion.

**Genotype encoder and decoder** ( $\mathcal{E}_G, \mathcal{D}_G$ ) Each sample’s genotype can be represented by 22 autosomal SNP sequences  $x_j^{(G)} = [x_j^{(G_1)}, \dots, x_j^{(G_{22})}]$  where  $x_j^{(G_c)} \in \{0, 1, 2\}^{S_c}$  denotes the SNP sequence of length  $S_c$  for chromosome  $c$ . Given that recombination breaks linkage between chromosomes making genetic variants on different chromosomes largely uncorrelated we decide to train chromosome-specific VAEs. Each encoder  $\mathcal{E}_{G_c}$  maps  $x_j^{(G_c)}$  to a latent vector  $z^{(G_c)} \in \mathbb{R}^{d_G}$ , yielding the full genotype embedding  $z_j^{(G)} = [z_j^{(G_1)}, \dots, z_j^{(G_{22})}] \in \mathbb{R}^{22 \times d_G}$ , after concatenation of all 22 chromosome-level embeddings. Each encoder-decoder pair  $(\mathcal{E}_G, \mathcal{D}_G)$  shares the same architecture. SNPs are embedded via learned token and positional embeddings (added element-wise) into  $\mathbb{R}^{S_c \times d_{\text{model}}}$ . The encoder comprises three stacked residual blocks He et al. (2016) interleaved with downsampling layers; the decoder mirrors this structure with upsampling layers. Inspired by attention-free Transformer alternatives (Yu et al., 2022), each residual block contains a token mixer (depthwise convolution) to capture local LD (Flagel et al., 2019; Sheehan & Song, 2016), followed by channel mixing with feed-forward multilayer perceptrons (MLPs).

To reconstruct SNP logits, the decoder output is linearly projected back to SNP space via a learned un-embedding matrix. We optimize a weighted cross-entropy loss that accounts for allelic imbalance, with an additional KL divergence term scaled by a hyperparameter  $\beta$  following  $\beta$ -VAE (Higgins et al., 2016). Training employs a cyclic learning rate schedule (Smith, 2017) to speed up convergence; gradient clipping to prevent exploding gradients, and KL divergence thresholding for stability (Child, 2020), along with mixed-precision and early stopping. Larger chromosomes (1–10) are trained on a single A100 GPU, while smaller ones are parallelized across up to four NVIDIA A5500 GPUs using data parallelism to increase the effective batch size. A full schematic of the architecture is provided in [Supplementary Fig. 12](#).

**Phenotype encoder and decoder** ( $\mathcal{E}_P, \mathcal{D}_P$ ) Phenotypic measurements are collected as a single dense vector  $x_j^{(P)} \in \mathbb{R}^{3,421}$ , encoding a fixed list of 3,421 phenotypic binary and continuous traits. Continuous traits are standardized to zero mean and unit variance, binary traits are encoded as 0 (controls) and 1 (cases), and missing entries are imputed to zero and their location is saved for use in

216 the training process. The encoder  $\mathcal{E}_P$  maps  $x_j^{(P)}$  to a latent token  $z_j^{(P)} \in \mathbb{R}^{1 \times d_P}$ , and the decoder  $\mathcal{D}_P$   
 217 reconstructs  $p$  from this embedding. The architecture of the phenotype encoder-decoder is inspired  
 218 by AutoComplete (An et al., 2023), using feedforward layers with LeakyReLU activations. The  
 219 model imputes missing values and reconstructs non-missing entries. During the training process,  
 220 copy masking is used to propagate realistic patterns of missingness observed in the clinical training  
 221 data, masking 30% of the observed data. We minimize a mixed reconstruction loss that combines  
 222 mean-squared error for continuous traits and binary cross-entropy for binary traits, evaluated only  
 223 on entries that are observed or copy-masked:

$$\mathcal{L} = \frac{1}{N} \sum_{j=1}^N \sum_{k=1}^P \tilde{M}_{jk} \begin{cases} (\hat{x}_{jk} - x_{jk})^2, & k \in \mathcal{C}, \\ -x_{jk} \log \hat{x}_{jk} - (1 - x_{jk}) \log(1 - \hat{x}_{jk}), & k \in \mathcal{B}, \end{cases} \quad (1)$$

224 where  $\hat{x}_j = \mathcal{D}_P(\mathcal{E}_P(x_j^{(P)}))$ ,  $\mathcal{C}$  and  $\mathcal{B}$  index continuous and binary traits,  $M$  is the true-observation  
 225 mask,  $M^{\text{copy}} \sim \text{CopyMask}(p = 0.3)$  is the copy-mask, and  $\tilde{M} = M \vee M^{\text{copy}}$ . Gradients are  
 226 back-propagated through  $\mathcal{E}_P$  and  $\mathcal{D}_P$  to update all parameters.

227 **Image encoder  $\mathcal{E}_I$**  For retinal image encoding, we employ RETFound (Zhou et al., 2023). RET-  
 228 Found is a state-of-the-art foundation model based on the masked autoencoder (MAE) architec-  
 229 ture (He et al., 2022), pre-trained on a dataset of  $\sim 1.6$ M unlabeled retinal images. To better align  
 230 features with our data distribution, we perform a lightweight self-supervised adaptation on UKB  
 231 images, which were not included in the pretraining, using the original MAE pixel-reconstruction  
 232 objective. We attach low-rank adapters (LoRA) (Hu et al., 2022) to the attention projections and  
 233 output projection in ViT blocks, freezing the base weights and updating only LoRA parameters. We  
 234 use the adapted encoder  $\mathcal{E}_I$  to map left/right fundus images  $x_k^{(I)} = (x_k^{(I_L)}, x_k^{(I_R)}) \in \mathbb{R}^{2 \times H \times W \times 3}$   
 235 into fixed-length tokens  $Z_j^{(I)} = [z_j^{(I_L)}, z_j^{(I_R)}] \in \mathbb{R}^{2 \times d_I}$ ; missing eyes are replaced by mask tokens.

236 The RETFound MAE decoder, due to the asymmetric nature of MAE where the decoder is often less  
 237 powerful than the encoder, is primarily designed to support the encoder’s pretraining objective (i.e.,  
 238 learning good embeddings) and is not optimized for high-fidelity image synthesis (Zhou et al., 2023;  
 239 He et al., 2022). Instead, in the experiments section, we explore diffusion models (Sohl-Dickstein  
 240 et al., 2015; Ho et al., 2020) as an alternative for this decoder.

## 241 4.2 FUSION MODEL $F$

242 After pretraining modality encoders, each  $j$ -th sample is represented by a triplet of latent token  
 243 sequences:

$$244 z_j = [z_j^{(G)} \in \mathbb{R}^{22 \times d_G}, z_j^{(P)} \in \mathbb{R}^{1 \times d_P}, z_j^{(I)} \in \mathbb{R}^{2 \times d_I}]$$

245 We apply token-wise standardization, project each token to a shared space  $\mathbb{R}^{T \times h}$  via learned pro-  
 246 jections, and concatenate them to obtain a unified token stream  $\tilde{z}_j \in \mathbb{R}^{T \times h}$ , where  $T$  is the total  
 247 number of tokens. A learnable [MASK] token replaces masked entries based on a stochastic mask-  
 248 ing scheme (He et al., 2022), and optional Gaussian noise  $\varepsilon \sim \mathcal{N}(0, \sigma^2 I)$  is residually added to  
 249 regularize the representation, resulting in  $\tilde{z}'_j = \tilde{z}_j + \varepsilon$ . Finally, dropout is applied on individual  
 250 dimensions of each token.

251 The fusion model  $F : \mathbb{R}^{T \times h} \rightarrow \mathbb{R}^{T \times h}$  is a transformer encoder configured with multi-head self-  
 252 attention, ReLU activations, and learned positional embeddings. In our design, we opt for a 32-layer  
 253 architecture, with 8 attention heads, and a hidden dimension  $h = 2048$ . Layer normalization is  
 254 applied prior to attention and feedforward operations. A [CLS] token is prepended to the input  
 255 to enable global summarization and ensure that there is always at least one token that is attended.  
 256 The model processes the token stream and outputs contextualized representations  $\hat{z}_j$ , which are  
 257 then segmented back into modality-specific embeddings and passed through frozen decoders  $\mathcal{D}_m$  to  
 258 reconstruct the masked inputs  $\hat{x}_j = (\hat{x}_j^{(G)}, \hat{x}_j^{(P)}, \hat{x}_j^{(I)})$ .

259 The training objective involves a combination of masked token modeling  $\mathcal{L}_{\text{mask}}$ , denoising recon-  
 260 struction  $\mathcal{L}_{\text{recon}}$ , and contrastive InfoNCE  $\mathcal{L}_{\text{InfoNCE}}$  loss terms, jointly optimized to balance fine-  
 261 grained token fidelity and modality-agnostic semantic alignment (Equation 3). The masked to-  
 262 ken modeling component operates similarly to the masked language modeling (MLM) objective in

BERT (Devlin et al., 2019), but generalized to arbitrary data modalities. For each sample, we randomly select a subset of tokens to be replaced with a learned [MASK] vector. The fusion model  $F$  is then tasked with reconstructing the original masked inputs from the surrounding context, including both intra- and inter-modality information.  $\mathcal{L}_{\text{mask}}$  encourages the model to learn conditional dependencies that span across heterogeneous sources, thereby enabling strong cross-modal imputation capabilities. In parallel, the reconstruction loss  $\mathcal{L}_{\text{recon}}$  is computed on unmasked tokens and serves to stabilize training by enforcing consistency of the latent representations, even under mild input Gaussian noise perturbations or dropout applied on individual dimensions of each token. To encourage cross-modal coherence, the InfoNCE loss  $\mathcal{L}_{\text{InfoNCE}}$  aligns pooled modality-specific tokens for the same sample while repelling embeddings from other samples. Specifically, the loss encourages high similarity between anchor average-pooled representations  $\bar{z}_j^{(m)} = \text{POOL}(z_j^{(m)}) \in \mathbb{R}^{1 \times h}$  and  $\bar{z}_j^{(m')}$  for different modalities  $m \neq m'$  for the same  $j$ -th sample, while treating all cross-sample pairs  $\bar{z}_k^{(m')}$  with  $k \neq j$  as negatives. This formulation implicitly maximizes a lower bound on the mutual information  $I(z_j^{(m)}, z_j^{(m')})$  between modality pairs under a shared fusion representation, as formalized by the InfoNCE criterion (Oord et al., 2018) (Equation 2):

$$\mathcal{L}_{\text{InfoNCE}}(\bar{z}_j^{(m)}) = -\log \frac{\sum_{\substack{m' \neq m \\ j \neq k}} e^{\tau^{-1} \bar{z}_j^{(m)\top} \bar{z}_j^{(m')}}}{\sum_{\substack{m' \neq m \\ j \neq k}} e^{\tau^{-1} \bar{z}_j^{(m)\top} \bar{z}_k^{(m')}} + \sum_{m' \neq m} e^{\tau^{-1} \bar{z}_j^{(m)\top} \bar{z}_j^{(m')}}} \quad (2)$$

with a contrastive temperature  $\tau > 0$ . To summarize, for a batch of  $N$  samples, the total fusion loss is formulated as (Equation 3):

$$\mathcal{L}_F = \sum_{j=1}^N \sum_m (1 - \lambda) \left( \mathcal{L}_{\text{mask}}(z_j^{(m)}, \hat{z}_j^{(m)}) + \mathcal{L}_{\text{recon}}(z_j^{(m)}, \hat{z}_j^{(m)}) \right) + \lambda \mathcal{L}_{\text{InfoNCE}}(\bar{z}_j^{(m)}) \quad (3)$$

where  $\mathcal{L}_{\text{mask}}$  and  $\mathcal{L}_{\text{recon}}$  is the mean-squared error and  $\lambda \in [0, 1]$  is the weighting hyperparameter which controls the relative importance of the contrastive loss.

$F$  is trained for 25 epochs ( $\approx 1$  hour/epoch) with batch size 128 on a NVIDIA A100 GPU. We optimize with the quasi-hyperbolic variant of Adam (Ma & Yarats, 2018), with hyperparameters  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\nu_1 = 0.7$ ,  $\nu_2 = 1$  and a weight decay of  $10^{-4}$ , and employ a cosine-annealing learning-rate schedule with a 10-epoch linear warm-up (learning rate from  $0.1 \times$  base to base) and a minimum learning rate floor of  $10^{-6}$ . Gradients are clipped to a max-norm of 10. We use  $\lambda = 0.5$  and  $\tau = 0.03$ , which have shown better performance across our experiments. The total number of trainable parameters is  $> 1,095$  million with  $> 3.5$  TFLOPs per forward pass.

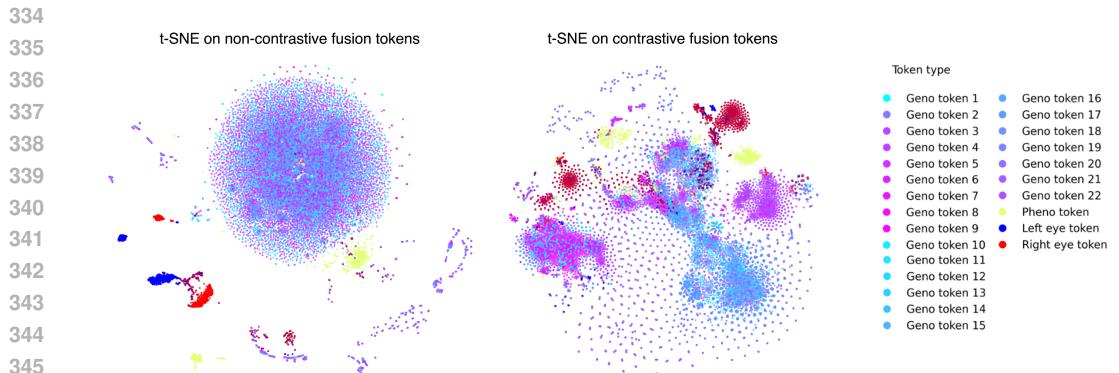
## 5 EXPERIMENTS AND EVALUATIONS

We comprehensively evaluate PM1 on several fronts—latent space exploration, genotype reconstruction and synthesis, retinal image synthesis, phenotype prediction and cross-modal inference—using the test split of UKB to validate both reconstructive fidelity and downstream predictive utility.

**Latent space exploration** We analyze both the modality-specific latent embeddings generated by  $\mathcal{E}_m$  and the pooled fusion embeddings produced by the multimodal transformer  $F$  using principal component analysis (PCA) and t-distributed stochastic neighbor embedding (t-SNE) (Van der Maaten & Hinton, 2008) for projection into two dimensions. Our results reveal that each modality contributes distinct and semantically meaningful structure to the joint representation space. Genotype embeddings, for instance, naturally cluster by ancestry without any supervision, reflecting population stratification effects encoded in autosomal variants. Similarly, phenotype embeddings display unsupervised separation by biological sex and body mass index (BMI) (see **Supplementary Figs. 5 and 6**).

More compelling insights arise from exploring the fused embeddings produced by the transformer  $F$ , which integrates signals from all modalities into a unified latent space. As shown in **Supplemen-**

324 **Supplementary Figs. 7 and 8**, PCA on the pooled fusion tokens reveals cohesive clusters that blend ancestry,  
 325 phenotypic traits, and imaging characteristics, confirming that the fusion model performs mean-  
 326 ingful cross-modal alignment. Crucially, we observe that the inclusion of the contrastive InfoNCE  
 327 loss dramatically improves this alignment: Fig. 2 compares latent fusion spaces trained with and  
 328 without the contrastive objective, under identical architectures. Without contrastive supervision, the  
 329 embeddings segregate cleanly by modality. In contrast, the contrastively trained model aligns cor-  
 330 responding modality embeddings for the same sample: for instance, left and right eye tokens from  
 331 the same individual cluster tightly, demonstrating the model’s success in learning modality-invariant  
 332 representations. This property is particularly advantageous for missing data imputation and cross-  
 333 modal generation, where shared semantics across input views is critical.



344 **Figure 2: t-SNE projections of token-level embeddings from the fusion model across 1,000 UK**  
 345 **Biobank participants.** (Left) Embeddings from a model trained without the contrastive InfoNCE  
 346 loss. (Right) Embeddings obtained with contrastive training. Each point corresponds to a single  
 347 token—either one of 22 chromosome genotype tokens, a phenotype token, or a left and right retinal  
 348 image token—colored by token type. The 2D projections illustrate clear modality-specific clustering  
 349 and demonstrate the model’s improved cross-modal integration when contrastive loss is applied,  
 350 encouraging similar tokens across modalities to align more closely in the embedding space.

351 **Genotype reconstruction and synthesis** We evaluate the capacity of our genotype encoder-  
 352 decoder pair ( $\mathcal{E}_G, \mathcal{D}_G$ ) to reconstruct whole-genome input data and to synthesize realistic genotype  
 353 samples. For reconstruction, we benchmark performance using a weighted accuracy metric that  
 354 accounts for allele frequency imbalance across loci, comparing our genotype encoder architecture  
 355 against two baselines: a linear projection model and a shallow VAE (Supplementary Fig. 9). Our  
 356 genotype encoder consistently outperforms these alternatives across all autosomes.

357 Beyond reconstruction, we assess the realism of our synthetic genotypes by analyzing their LD  
 358 structure and compare to real genotypes (Geleta et al., 2023). Specifically, we compute the  
 359 folded allele frequency spectrum, showing the proportion of SNPs at each minor allele frequency  
 360 (Supplementary Fig. 9) and compute a correlation profile based on variant distance for real and  
 361 simulated sequences. Both measures suggest that our synthetic genotypes reflect patterns found in  
 362 real sequences.

363 **Retinal image synthesis** As an exploratory task to evaluate the richness of the learned multimodal  
 364 representations, we trained a conditional diffusion model for retinal image synthesis using Denoising  
 365 Diffusion Implicit Models (DDIM) (Song et al., 2022) for the backward denoising process. The core  
 366 is a U-Net architecture (Ronneberger et al., 2015) inspired by successful image generation models  
 367 (Rombach et al., 2022; Podell et al., 2024), composed of six downsampling and six symmetric  
 368 upsampling blocks with skip connections, and output channels of 128, 128, 256, 256, 512, and 512.  
 369 Input and output are 3-channels RGB images at configurable resolution. All downsampling blocks,  
 370 except the second, use a stride-2 convolution (LeCun et al., 1989; Krizhevsky et al., 2012) followed  
 371 by two ResNet blocks (He et al., 2016). Crucially, this U-Net incorporates cross-attention layers,  
 372 which attend to keys and values derived from PM1’s fusion embeddings, conditioning the image  
 373 generation process on integrated multimodal representations.

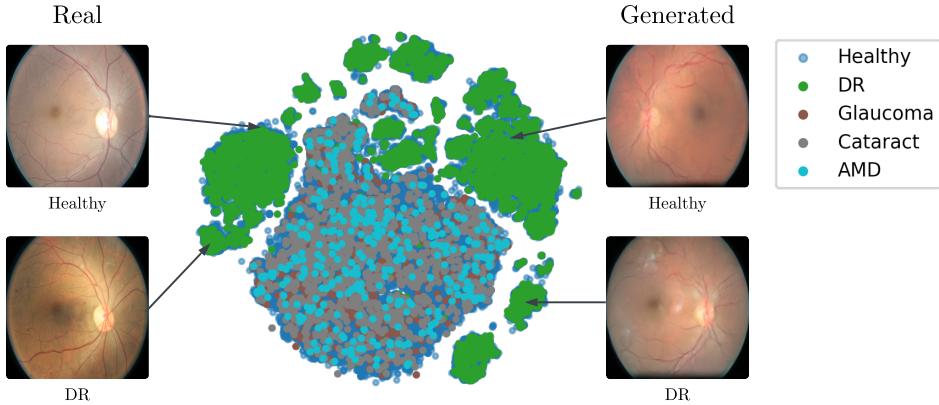


Figure 3: **Retinal image synthesis using PM1-guided conditional diffusion.** We visualize the PM1 image embeddings with t-SNE, color-coded by disease label (Healthy, Diabetic Retinopathy (DR), Glaucoma, Cataract, and Age-related Macular Degeneration (AMD)). The diffusion model is conditioned on PM1’s multimodal fusion tokens via cross-attention and trained to reconstruct realistic retinal images.

Using this setup, we synthesize 50,000 retinal images (25,000 healthy; 25,000 diabetic retinopathy). A ResNet-50 trained exclusively on these synthetic images achieves ROC-AUC **0.8502** on a held-out test set of real UKB and EyePACS images. We further report Fréchet Inception Distance (FID) **50.86** and Kernel Inception Distance (KID) **0.0532**. This experiment demonstrates that the fusion model learns meaningful representations capable of guiding a complex generative task. While the scale is still exploratory and the resulting model is relatively small for producing publication-quality clinical images, the results highlight the potential of this approach if expanded with more data and computational resources.

**Cross-modal inference** We perform cross-modal genotype inference across all 22 autosomes under different modality combinations. We confirm that incorporating more modalities consistently improves genotype prediction (for quantitative results, see [Supplementary Fig. 10](#)). Specifically, we test: (i)  $G+P+I \rightarrow G$ , where genotype, phenotype, and image modalities are jointly provided as context to reconstruct genotype tokens; (ii)  $P+I \rightarrow G$ , where genotype is inferred from phenotypic traits and retinal images; (iii)  $P \rightarrow G$  and  $I \rightarrow G$ , where predictive accuracy drops but still indicates that single modalities capture signals correlated with genetic variation.

For downstream trait inference, we extend evaluation to clinical phenotypes. PM1’s phenotype decoder is trained to reconstruct phenotypes and predict missing ones, and can be enriched with genotype and image inputs. We compare its fusion performance in Table 1. For each target phenotype, the observed phenotypes are masked at the phenotype encoder input of PM1, as well as all other phenotype entries directly related with that trait (more details in [Supplementary Section G.1](#)). Phenotype prediction generally benefits from fusing the other modalities (genotype and image), boosting ROC-AUC. As additional baselines, we evaluate MedGemma (Sellergren et al., 2025) on phenotype prediction using a representative held-out subset of 1,000 UKB participants with complete P+I coverage across nine clinically relevant traits, and ContIG (Taleb et al., 2022) as an image–genetics contrastive baseline on the full test set. Table 1 shows that PM1 consistently outperforms MedGemma across all phenotypes and task settings, achieving average ROC-AUC improvements of +0.32 in the phenotype-to-phenotype ( $P \rightarrow P$ ) setting and +0.28 in the phenotype-plus-image-to-phenotype ( $P+I \rightarrow P$ ) setting. On the full test set, PM1 surpasses ContIG with an average gain of +0.12 ROC-AUC corresponding to **18.6%** relative improvement across phenotypes. More details—including the rationale for the 1,000-sample subset in MedGemma and the benchmarking setup for ContIG—are provided in [Supplementary Sections G.2 and G.3](#).

## 5.1 GWAS AND GROUP-LEVEL SIGNALS IN PM1 EMBEDDINGS

We probed whether PM1’s multimodal embeddings encode biologically meaningful structure. A GWAS on PCA-reduced fusion embeddings (all modalities active) using PLINK 2 (Chang et al.,

432 Table 1: ROC-AUC for phenotype prediction on clinical traits. Left: UKB 1,000 participants with  
 433 P+I coverage comparing PM1 and MedGemma. Right: full test set with PM1 variants and ContIG.  
 434

435 Phenotype	436 1,000 test samples with P+I coverage				437 Full test set				
	438 PM1 <sup>P</sup>	439 MedGemma <sup>P</sup>	440 PM1 <sup>P+I</sup>	441 MedGemma <sup>P+I</sup>	442 PM1 <sup>P</sup>	443 PM1 <sup>P+I</sup>	444 PM1 <sup>P+G</sup>	445 PM1 <sup>P+I+G</sup>	446 ContIG <sup>I+G</sup>
447 Macular degeneration	0.953	0.542	<b>0.954</b>	0.577	0.792	<b>0.802</b>	0.795	0.799	0.644
448 Diabetic eye disease	<b>0.964</b>	0.552	0.962	0.566	0.805	<b>0.820</b>	0.805	0.815	0.731
449 Glaucoma	0.883	0.473	<b>0.921</b>	0.604	0.745	0.750	0.762	<b>0.766</b>	0.639
450 Cataract	0.735	0.578	<b>0.747</b>	0.588	0.686	0.693	0.702	<b>0.708</b>	0.683
451 Heart failure	0.811	0.511	<b>0.823</b>	0.543	<b>0.764</b>	<b>0.764</b>	0.761	0.762	0.673
452 Ischaemic stroke	0.792	0.391	<b>0.802</b>	0.604	0.794	0.795	0.804	<b>0.805</b>	0.591
453 Heart attack	<b>0.799</b>	0.524	0.791	0.521	<b>0.682</b>	<b>0.682</b>	0.664	0.664	0.633
454 Dementia	0.863	0.583	<b>0.874</b>	0.590	0.820	0.819	<b>0.836</b>	0.834	0.667
455 Alzheimer’s disease	<b>0.845</b>	0.612	0.844	0.603	0.818	0.807	<b>0.828</b>	<b>0.828</b>	0.668

447 2015) identified genome-wide hits in the HERC2 gene, on Chromosome 15. Variant *rs1129038* is  
 448 a well-established regulatory SNP involved in human eye pigmentation. It modulates expression  
 449 of the nearby OCA2 gene and is part of a founder haplotype strongly predictive of blue eye color  
 450 in European populations ( $p = 6.2 \times 10^{-46}$ ). This variant has also been linked to other ocular  
 451 traits including glaucoma, macular degeneration, and central corneal thickness. A second variant,  
 452 *rs1667394*, is an intronic SNP that likely influences pigmentation through OCA2 regulation and has  
 453 been previously associated with adiponectin levels.

454 As a complementary probe, using self-reported labels, we construct a 3-class task (European,  
 455 African, Asian) with a 9,152/2,288 train/test split (class-balanced by subsampling Europeans).  
 456 Training simple classifiers (logistic regression, random forest, naive Bayes, k-NN) on PM1’s pooled  
 457 fusion embeddings achieved **>99%** test accuracy and macro-F1, confirming that PM1 captures  
 458 population-level structure without ancestry supervision. See further details in **Supplementary Sec-  
 459 tion C**.

## 460 6 CONCLUSION

461 We present **PM1**, the first multimodal foundation model designed for large-scale biobank data,  
 462 which fuses genomic, phenotypic, and retinal imaging modalities to support cross-modal infer-  
 463 ence, generation, and representation learning. Through a range of experimental validations we  
 464 demonstrate the expressiveness and versatility of the learned multimodal representations. Despite  
 465 its strengths, our work bears several limitations. First, the UK Biobank cohort, while large and  
 466 deeply phenotyped, is known to underrepresent global ancestry diversity, which may limit the gen-  
 467 eralizability of our model to non-European populations. Second, we highlight the inherent imbal-  
 468 ance across modalities: images, phenotypes, and genotypes differ significantly in dimensionality,  
 469 signal-to-noise ratio, and sparsity. These differences raise open questions about optimal weighting  
 470 or modality-specific attention in fusion architectures. Third, the computational demands of train-  
 471 ing large transformer-based models on biobank-scale multimodal data are substantial. While we  
 472 document our architecture and provide code for reproducibility, resource constraints may limit ac-  
 473 cessibility.

474 Nonetheless, the potential scientific and translational benefits of models like PM1 are significant.  
 475 By enabling interpretable cross-modal inference and synthesis, our framework offers new avenues  
 476 for modality imputation, and genotype-phenotype association discovery. In particular, this work  
 477 aligns with broader goals in biomedical research, where the ability to connect rare functional genetic  
 478 variation to deep phenotypic outcomes—including imaging, biomarkers, and clinical traits—can  
 479 help and understand causal mechanisms (Szustakowski et al., 2021).

480 Finally, we recognize that models operating on high-dimensional multimodal health data carry eth-  
 481 ical risks. Concerns around data privacy, re-identification, and disparate performance across de-  
 482 mographic groups are especially salient in genomics and biomedical AI. As larger and more di-  
 483 verse datasets become available, we hope PM1 can serve as a step toward robust, equitable, and  
 484 multimodally-informed precision medicine.

485 **Disclosure** We used large language models to aid in grammar, wording, and style improvements  
 486 during the writing of this paper. The models were not used for generating ideas or analyses.

486 REFERENCES  
487

488 110th United States Congress. Public law 110-233: Genetic information nondiscrimination act of  
489 2008, 5 2008. URL <https://www.govinfo.gov/app/details/PLAW-110publ233>.

490 Julián N Acosta, Guido J Falcone, Pranav Rajpurkar, and Eric J Topol. Multimodal biomedical ai.  
491 *Nature medicine*, 28(9):1773–1784, 2022.

492 Ulzee An, Ali Pazokitoroudi, Marcus Alvarez, Liyan Huang, Silviu Bacanu, Andrew J Schork,  
493 Kenneth Kendler, Päivi Pajukanta, Jonathan Flint, Noah Zaitlen, et al. Deep learning-based phe-  
494 notype imputation on population-scale biobank data increases genetic discoveries. *Nature Genet-*  
495 *ics*, 55(12):2269–2276, 2023.

496 Jason Ansel, Edward Yang, Horace He, Natalia Gimelshein, Animesh Jain, Michael Voznesensky,  
497 Bin Bao, Peter Bell, David Berard, Evgeni Burovski, Geeta Chauhan, Anjali Chourdia, Will  
498 Constable, Alban Desmaison, Zachary DeVito, Elias Ellison, Will Feng, Jiong Gong, Michael  
499 Gschwind, Brian Hirsh, Sherlock Huang, Kshiteej Kalambarkar, Laurent Kirsch, Michael Lazos,  
500 Mario Lezcano, Yanbo Liang, Jason Liang, Yinghai Lu, CK Luk, Bert Maher, Yunjie Pan, Chris-  
501 tian Pührsch, Matthias Reso, Mark Saroufim, Marcos Yukio Siraichi, Helen Suk, Michael Suo,  
502 Phil Tillet, Eikan Wang, Xiaodong Wang, William Wen, Shunting Zhang, Xu Zhao, Keren Zhou,  
503 Richard Zou, Ajit Mathews, Gregory Chanan, Peng Wu, and Soumith Chintala. PyTorch 2: Faster  
504 Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation.  
505 In *29th ACM International Conference on Architectural Support for Programming Languages and*  
506 *Operating Systems, Volume 2 (ASPLOS '24)*. ACM, April 2024. doi: 10.1145/3620665.3640366.  
507 URL <https://pytorch.org/assets/pytorch2-2.pdf>.

508 Kristiina Ausmees and Carl Nettelblad. A deep learning framework for characterization of genotype  
509 data. *G3 (Bethesda)*, 12(3), March 2022.

510 Joshua D Backman, Alexander H Li, Anthony Marcketta, Dylan Sun, Joelle Mbatchou, Michael D  
511 Kessler, Christian Benner, Daren Liu, Adam E Locke, Suganthi Balasubramanian, et al. Exome  
512 sequencing and analysis of 454,787 uk biobank participants. *Nature*, 599(7886):628–634, 2021.

513 Tyler Bahr, Truong A Vu, Jared J Tuttle, and Raymond Iezzi. Deep learning and machine learning  
514 algorithms for retinal image analysis in neurodegenerative disease: systematic review of datasets  
515 and models. *Translational Vision Science & Technology*, 13(2):16–16, 2024.

516 Shaojie Bai, J Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional  
517 and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*, 2018.

518 David Bonet, May Levin, Daniel Mas Montserrat, and Alexander G Ioannidis. Machine learning  
519 strategies for improved phenotype prediction in underrepresented populations. In *Pacific Sympo-*  
520 *sium on Biocomputing. Pacific Symposium on Biocomputing*, volume 29, pp. 404, 2024.

521 Gary Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 11 2000. URL <https://www.drdobbs.com/open-source/the-opencv-library/184404319>.

522 Clare Bycroft, Colin Freeman, Desislava Petkova, Gavin Band, Lloyd T. Elliott, Kevin Sharp, Allan  
523 Motyer, Damjan Vukcevic, Olivier Delaneau, Jared O'Connell, Adrian Cortes, Samantha Welsh,  
524 Alan Young, Mark Effingham, Gil McVean, Stephen Leslie, Naomi Allen, Peter Donnelly, and  
525 Jonathan Marchini. The uk biobank resource with deep phenotyping and genomic data. *Nature*,  
526 562(7726):203–209, October 2018. ISSN 1476-4687. URL <https://doi.org/10.1038/s41586-018-0579-z>.

527 Francisco Carrillo-Perez, Marija Pizurica, Kathleen Marchal, and Olivier Gevaert. Synthetic multi-  
528 modal data modelling for data imputation. *Nature Biomedical Engineering*, pp. 1–5, 2024.

529 Christopher C Chang, Carson C Chow, Laurent CAM Tellier, Shashaank Vattikuti, Shaun M Purcell,  
530 and James J Lee. Second-generation plink: rising to the challenge of larger and richer datasets.  
531 *GigaScience*, 4(1), February 2015. ISSN 2047-217X. doi: 10.1186/s13742-015-0047-8. URL  
532 [www.cog-genomics.org/plink/2.0/](http://www.cog-genomics.org/plink/2.0/).

533 Rewon Child. Very deep VAEs generalize autoregressive models and can outperform them on im-  
534 ages. *arXiv preprint arXiv:2011.10650*, 2020.

540 Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam  
 541 Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh,  
 542 Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam  
 543 Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James  
 544 Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Lev-  
 545 skaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin  
 546 Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret  
 547 Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick,  
 548 Andrew M. Dai, Thanumalayan Sankaranarayana Pillai, Marie Pellat, Aitor Lewkowycz, Erica  
 549 Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Bren-  
 550 nan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas  
 551 Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. Palm: Scaling language modeling with pathways,  
 552 2022. URL <https://arxiv.org/abs/2204.02311>.

553 Marçal Comajoan Cara, Daniel Mas Montserrat, and Alexander Ioannidis. Popgenadapt: Semi-  
 554 supervised domain adaptation for genotype-to-phenotype prediction in underrepresented pop-  
 555ulations. In *Pacific Symposium on Biocomputing 2024*, pp. 327–340. World Scientific Pub-  
 556lishing Company, 1 2024. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10906137>.

557 Francis Crick. Central dogma of molecular biology. *Nature*, 227(5258):561–563, August 1970.

558 Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. In  
 559 *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=mZn2Xyh9Ec>.

560 Christopher DeBoever, Yosuke Tanigawa, Matthew Aguirre, Greg McInnes, Adam Lavertu, and  
 561 Manuel A Rivas. Assessing digital phenotyping to enhance genetic studies of human diseases.  
 562 *The American Journal of Human Genetics*, 106(5):611–622, 2020.

563 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A Large-Scale  
 564 Hierarchical Image Database. *CVPR*, 2009.

565 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of  
 566 deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and  
 567 Thamar Solorio (eds.), *Proceedings of the 2019 Conference of the North American Chapter of  
 568 the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long  
 569 and Short Papers)*, pp. 4171–4186. Association for Computational Linguistics, June 2019. URL  
<https://aclanthology.org/N19-1423>.

570 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas  
 571 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An  
 572 image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint  
 573 arXiv:2010.11929*, 2020.

574 Emma Dugas, Jared, Jorge, and Will Cukierski. Diabetic retinopathy detection, 2015. URL <https://kaggle.com/competitions/diabetic-retinopathy-detection>.

575 Lloyd T Elliott, Kevin Sharp, Fidel Alfaro-Almagro, Sinan Shi, Karla L Miller, Gwenaëlle Douaud,  
 576 Jonathan Marchini, and Stephen M Smith. Genome-wide association studies of brain imaging  
 577 phenotypes in uk biobank. *Nature*, 562(7726):210–216, 2018.

578 Lex Flagel, Yaniv Brandvain, and Daniel R Schrider. The unreasonable effectiveness of convolu-  
 579 tional neural networks in population genetic inference. *Mol. Biol. Evol.*, 36(2):220–238, February  
 580 2019.

581 Matthew Fredrikson, Eric Lantz, Somesh Jha, Simon Lin, David Page, and Thomas Ristenpart.  
 582 Privacy in pharmacogenetics: An {End-to-End} case study of personalized warfarin dosing. In  
 583 *23rd USENIX security symposium (USENIX Security 14)*, pp. 17–32, 2014.

584 Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. Convolutional  
 585 sequence to sequence learning. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th*

594 *International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning*  
 595 *Research*, pp. 1243–1252. PMLR, 06–11 Aug 2017. URL <https://proceedings.mlr.press/v70/gehring17a.html>.

596

597 Margarita Geleta, Daniel Mas Montserrat, Xavier Giro-i Nieto, and Alexander G Ioannidis. Deep  
 598 variational autoencoders for population genetics. *biorxiv*, pp. 2023–09, 2023.

599

600 Varun Gulshan, Lily Peng, Marc Coram, Martin C. Stumpe, Derek Wu, Arunachalam  
 601 Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Tom Madams, Jorge Cuadros, Ra-  
 602 masamy Kim, Rajiv Raman, Philip C. Nelson, Jessica L. Mega, and Dale R. Webster. Devel-  
 603 opment and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in  
 604 Retinal Fundus Photographs. *JAMA*, 316(22):2402–2410, 12 2016. ISSN 0098-7484. URL  
 605 <https://doi.org/10.1001/jama.2016.17216>.

606

607 Bjarni V Halldorsson, Hannes P Eggertsson, Kristjan HS Moore, Hannes Hauswedell, Ogmundur  
 608 Eiriksson, Magnus O Ulfarsson, Gunnar Palsson, Marteinn T Hardarson, Asmundur Oddsson,  
 609 Brynjar O Jansson, et al. The sequences of 150,119 genomes in the uk biobank. *Nature*, 607  
 610 (7920):732–740, 2022.

611

612 Trevor Hastie, Rahul Mazumder, Jason D Lee, and Reza Zadeh. Matrix completion and low-rank  
 613 SVD via fast alternating least squares. *The Journal of Machine Learning Research*, 16(1):3367–  
 614 3402, 2015.

615

616 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-  
 617 nition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.  
 618 770–778, 2016.

619

620 Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollar, and Ross Girshick. Masked  
 621 autoencoders are scalable vision learners. In *2022 IEEE/CVF Conference on Computer Vision*  
 622 and *Pattern Recognition (CVPR)*, pp. 15979–15988. IEEE, June 2022. URL <https://doi.org/10.1109/CVPR52688.2022.01553>.

623

624 Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (GELUs). *arXiv*, 2016. URL  
 625 <https://arxiv.org/abs/1606.08415>.

626

627 Irina Higgins, Loïc Matthey, Arka Pal, Christopher P. Burgess, Xavier Glorot, Matthew M.  
 628 Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts  
 629 with a constrained variational framework. In *International Conference on Learning Representa-  
 630 tions*, 2016. URL <https://api.semanticscholar.org/CorpusID:46798026>.

631

632 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on*  
 633 *Deep Generative Models and Downstream Applications*, 2021. URL <https://openreview.net/forum?id=qw8AKxfYbI>.

634

635 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In  
 636 H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neu-  
 637 ral Information Processing Systems*, volume 33, pp. 6840–6851. Curran Associates, Inc.,  
 638 2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf).

639

640 Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and  
 641 Weizhu Chen. Lora: Low-rank adaptation of large language models. In *International Confer-  
 642 ence on Learning Representations*, 2022.

643

644 Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by  
 645 reducing internal covariate shift. In *International conference on machine learning*, pp. 448–456.  
 646 pmlr, 2015.

647 Konrad J Karczewski and Michael P Snyder. Integrative omics for health and disease. *Nature*  
 648 *Reviews Genetics*, 19(5):299–310, 2018.

648 Pearse A. Keane, Carlota M. Grossi, Paul J. Foster, Qi Yang, Charles A. Reisman, Kinpui Chan,  
 649 Tunde Peto, Dhanes Thomas, and Praveen J. Patel. Optical coherence tomography in the uk  
 650 biobank study - rapid automated analysis of retinal thickness for large population-based studies.  
 651 *PLOS ONE*, 11(10), October 2016. ISSN 1932-6203. URL <https://doi.org/10.1371/journal.pone.0164095>.

653 Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint*  
 654 *arXiv:1312.6114*, 2013.

656 Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep  
 657 convolutional neural networks. In F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger  
 658 (eds.), *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc.,  
 659 2012. URL [https://proceedings.neurips.cc/paper\\_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf).

661 Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hub-  
 662 bard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition.  
 663 *Neural computation*, 1(4):541–551, 1989.

665 Benjamin Lefauze, Francisco Massa, Diana Liskovich, Wenhan Xiong, Vittorio Caggiano, Sean  
 666 Naren, Min Xu, Jieru Hu, Marta Tintore, Susan Zhang, Patrick Labatut, Daniel Haziza, Luca  
 667 Wehrstedt, Jeremy Reizenstein, and Grigory Sizov. xformers: A modular and hackable trans-  
 668 former modelling library, 2022. URL <https://github.com/facebookresearch/xformers>.

670 Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Nau-  
 671 mann, Hoifung Poon, and Jianfeng Gao. Llava-med: training a large language-and-vision assis-  
 672 tant for biomedicine in one day. In *Proceedings of the 37th International Conference on Neural*  
 673 *Information Processing Systems*, NIPS '23, Red Hook, NY, USA, 2023. Curran Associates Inc.

674 Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband:  
 675 A novel bandit-based approach to hyperparameter optimization. *Journal of Machine Learning*  
 676 *Research*, 18(185):1–52, 2018.

678 Shanchuan Lin, Bingchen Liu, Jiashi Li, and Xiao Yang. Common diffusion noise schedules and  
 679 sample steps are flawed. In *Proceedings of the IEEE/CVF Winter Conference on Applications*  
 680 *of Computer Vision (WACV)*, pp. 5404–5411, 1 2024. URL [https://openaccess.thecvf.com/content/WACV2024/papers/Lin\\_Common\\_Diffusion\\_Noise\\_Schedules\\_and\\_Sample\\_Steps\\_Are\\_Flawed\\_WACV\\_2024\\_paper.pdf](https://openaccess.thecvf.com/content/WACV2024/papers/Lin_Common_Diffusion_Noise_Schedules_and_Sample_Steps_Are_Flawed_WACV_2024_paper.pdf).

683 Thomas J. Littlejohns, Jo Holliday, Lorna M. Gibson, Steve Garratt, Niels Oesingmann, Fidel  
 684 Alfaro-Almagro, Jimmy D. Bell, Chris Boultwood, Rory Collins, Megan C. Conroy, Nicola Crab-  
 685 tree, Nicola Doherty, Alejandro F. Frangi, Nicholas C. Harvey, Paul Leeson, Karla L. Miller, Ste-  
 686 fan Neubauer, Steffen E. Petersen, Jonathan Sellors, Simon Sheard, Stephen M. Smith, Cathie  
 687 L. M. Sudlow, Paul M. Matthews, and Naomi E. Allen. The uk biobank imaging enhancement  
 688 of 100,000 participants: rationale, data collection, management and future directions. *Nature*  
 689 *Communications*, 11(1), May 2020. ISSN 2041-1723. URL <https://doi.org/10.1038/s41467-020-15948-9>.

691 Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *Interna-  
 692 tional Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Skq89Scxx>.

694 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Confer-  
 695 ence on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.

698 Jerry Ma and Denis Yarats. Quasi-hyperbolic momentum and adam for deep learning. *arXiv preprint*  
 699 *arXiv:1810.06801*, 2018.

700 Siddharth Mangrulkar, Quentin Lhoest, et al. Peft: Parameter-efficient fine-tuning of transformers.  
 701 <https://github.com/huggingface/peft>, 2023.

702 Alicia R. Martin, Masahiro Kanai, Yoichiro Kamatani, Yukinori Okada, Benjamin M. Neale, and  
 703 Mark J. Daly. Clinical use of current polygenic risk scores may exacerbate health disparities.  
 704 *Nature Genetics*, 51(4):584–591, March 2019. ISSN 1546-1718. URL <https://doi.org/10.1038/s41588-019-0379-x>.

705

706 Neo Wei Ming, Zhehui Wang, Cheng Liu, Rick Siow Mong Goh, and Tao Luo. Ma-bert: To-  
 707 wards matrix arithmetic-only bert inference by eliminating complex non-linear functions. In *The*  
 708 *Eleventh International Conference on Learning Representations*, 2022.

709

710 Yves Moreau. Crack down on genomic surveillance. *Nature*, 576(7785):36–38, December 2019.  
 711 ISSN 1476-4687. URL <https://doi.org/10.1038/d41586-019-03687-x>.

712

713 Muhammad Muneeb, Samuel Feng, and Andreas Henschel. Transfer learning for geno-  
 714 type–phenotype prediction using deep learning models. *BMC Bioinformatics*, 23(1), November  
 715 2022. ISSN 1471-2105. doi: 10.1186/s12859-022-05036-8. URL <http://dx.doi.org/10.1186/s12859-022-05036-8>.

716

717 Alfredo Nazabal, Pablo M Olmos, Zoubin Ghahramani, and Isabel Valera. Handling incomplete  
 718 heterogeneous data using vaes. *Pattern Recognition*, 107:107501, 2020.

719

720 Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predic-  
 721 tive coding. *arXiv preprint arXiv:1807.03748*, 2018.

722

723 Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe  
 724 Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image  
 725 synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. URL  
<https://openreview.net/forum?id=di52zR8xgf>.

726

727 Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging.  
*SIAM Journal on Control and Optimization*, 30(4):838–855, July 1992. ISSN 1095-7138. URL  
<https://doi.org/10.1137/0330046>.

728

729 Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and ap-  
 730 proximate inference in deep generative models. In *International conference on machine learning*,  
 731 pp. 1278–1286. PMLR, 2014.

732

733 Maria Rigaki and Sebastian Garcia. A survey of privacy attacks in machine learning. *ACM Com-  
 734 puting Surveys*, 56(4):1–34, 2023.

735

736 Robin Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis  
 737 with latent diffusion models. In *2022 IEEE/CVF Conference on Computer Vision and Pattern*  
 738 *Recognition (CVPR)*, pp. 10674–10685. IEEE Computer Society, 6 2022. URL <https://doi.ieee.org/10.1109/CVPR52688.2022.01042>.

739

740 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for  
 741 biomedical image segmentation. In *International Conference on Medical Image Comput-  
 742 ing and Computer-Assisted Intervention - MICCAI 2015*, pp. 234–241. Springer Interna-  
 743 tional Publishing, 2015. ISBN 9783319245744. URL [https://doi.org/10.1007/978-3-319-24574-4\\_28](https://doi.org/10.1007/978-3-319-24574-4_28).

744

745 David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-  
 746 propagating errors. *Nature*, 323(6088):533–536, October 1986.

747

748 David Ruppert. Efficient estimations from a slowly convergent robbins-monro process. Technical  
 749 report, Cornell University Operations Research and Industrial Engineering, 1988. URL <https://hdl.handle.net/1813/8664>.

750

751 Andrew Sellergren, Sahar Kazemzadeh, Tiam Jaroensri, Atilla Kiraly, Madeleine Traverse, Timo  
 752 Kohlberger, Shawn Xu, Fayaz Jamil, Cían Hughes, Charles Lau, Justin Chen, Fereshteh Mahvar,  
 753 Liron Yatziv, Tiffany Chen, Bram Sterling, Stefanie Anna Baby, Susanna Maria Baby, Jeremy  
 754 Lai, Samuel Schmidgall, Lu Yang, Kejia Chen, Per Bjornsson, Shashir Reddy, Ryan Brush, Ken-  
 755 neth Philbrick, Mercy Asiedu, Ines Mezerreg, Howard Hu, Howard Yang, Richa Tiwari, Sunny  
 Jansen, Preeti Singh, Yun Liu, Shekoofeh Azizi, Aishwarya Kamath, Johan Ferret, Shreya Pathak,

756 Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane  
 757 Riviere, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos,  
 758 Edouard Yvinec, Michelle Casbon, Elena Buchatskaya, Jean-Baptiste Alayrac, Dmitry Lepikhin,  
 759 Vlad Feinberg, Sebastian Borgeaud, Alek Andreev, Cassidy Hardin, Robert Dadashi, Léonard  
 760 Hussenot, Armand Joulin, Olivier Bachem, Yossi Matias, Katherine Chou, Avinatan Hassidim,  
 761 Kavi Goel, Clement Farabet, Joelle Barral, Tris Warkentin, Jonathon Shlens, David Fleet, Vic-  
 762 tor Cotruta, Omar Sanseviero, Gus Martins, Phoebe Kirk, Anand Rao, Shravya Shetty, David F.  
 763 Steiner, Can Kirmizibayrak, Rory Pilgrim, Daniel Golden, and Lin Yang. Medgemma technical  
 764 report, 2025. URL <https://arxiv.org/abs/2507.05201>.

765 Sara Sheehan and Yun S Song. Deep learning for population genetic inference. *PLoS Comput. Biol.*,  
 766 12(3):e1004845, March 2016.

767 Xinghua Shi and Xintao Wu. An overview of human genetic privacy. *Annals of the New York  
 768 Academy of Sciences*, 1387(1):61–72, 2017. URL <https://nyaspubs.onlinelibrary.wiley.com/doi/10.1111/nyas.13211>.

771 Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan  
 772 Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfahl, et al. Large language models encode  
 773 clinical knowledge. *Nature*, 620(7972):172–180, 2023.

774 Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Mohamed Amin, Le Hou,  
 775 Kevin Clark, Stephen R Pfahl, Heather Cole-Lewis, et al. Toward expert-level medical question  
 776 answering with large language models. *Nature Medicine*, pp. 1–8, 2025.

777 Leslie N Smith. Cyclical learning rates for training neural networks. In *2017 IEEE winter conference  
 778 on applications of computer vision (WACV)*, pp. 464–472. IEEE, 2017.

780 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsuper-  
 781 vised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei (eds.),  
 782 *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Pro-  
 783 ceedings of Machine Learning Research*, pp. 2256–2265. PMLR, 7 2015. URL <https://proceedings.mlr.press/v37/sohl-dickstein15.html>.

785 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Inter-  
 786 national Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=St1giarCHLP>.

788 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models, 2022. URL  
 789 <https://arxiv.org/abs/2010.02502>.

791 Cathie Sudlow, John Gallacher, Naomi Allen, Valerie Beral, Paul Burton, John Danesh, Paul  
 792 Downey, Paul Elliott, Jane Green, Martin Landray, Bette Liu, Paul Matthews, Giok Ong, Jill  
 793 Pell, Alan Silman, Alan Young, Tim Sprosen, Tim Peakman, and Rory Collins. UK Biobank:  
 794 An Open Access Resource for Identifying the Causes of a Wide Range of Complex Diseases of  
 795 Middle and Old Age. *PLOS Medicine*, 12(3):1–10, 03 2015. doi: 10.1371/journal.pmed.1001779.  
 796 URL <https://doi.org/10.1371/journal.pmed.1001779>.

797 S. Suganyadevi, V. Seetalakshmi, and K. Balasamy. A review on deep learning in medical image  
 798 analysis. *International Journal of Multimedia Information Retrieval*, 11(1):19–38, 2021. doi:  
 799 10.1007/s13735-021-00218-1.

800 Joseph D Szustakowski, Suganthi Balasubramanian, Erika Kvikstad, Shareef Khalid, Paola G Bron-  
 801 son, Ariella Sasson, Emily Wong, Daren Liu, J Wade Davis, Carolina Haefliger, et al. Advancing  
 802 human genetics research and drug discovery through exome sequencing of the uk biobank. *Nature  
 803 genetics*, 53(7):942–948, 2021.

804 Aiham Taleb, Matthias Kirchler, Remo Monti, and Christoph Lippert. Contig: Self-supervised mul-  
 805 timodal contrastive learning for medical imaging with genetics. In *Proceedings of the IEEE/CVF  
 806 conference on computer vision and pattern recognition*, pp. 20908–20921, 2022.

808 Yosuke Tanigawa, Junyang Qian, Guhan Venkataraman, Johanne Marie Justesen, Ruilin Li, Robert  
 809 Tibshirani, Trevor Hastie, and Manuel A Rivas. Significant sparse polygenic risk scores across  
 813 traits in uk biobank. *PLoS Genetics*, 18(3):e1010105, 2022.

810 Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej,  
 811 Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical  
 812 report. *arXiv preprint arXiv:2503.19786*, 2025.

813

814 Christina V. Theodoris, Ling Xiao, Anant Chopra, Mark D. Chaffin, Zeina R. Al Sayed, Matthew C.  
 815 Hill, Helene Mantineo, Elizabeth M. Brydon, Zexian Zeng, X. Shirley Liu, and Patrick T.  
 816 Ellinor. Transfer learning enables predictions in network biology. *Nature*, 618(7965):616–  
 817 624, June 2023. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-023-06139-9. URL  
 818 <https://www.nature.com/articles/s41586-023-06139-9>.

819

820 Daniel Truhn, Jan-Niklas Eckardt, Dyke Ferber, and Jakob Nikolas Kather. Large language models  
 821 and multimodal foundation models for precision oncology. *NPJ Precision Oncology*, 8(1):72,  
 822 2024.

823

824 Tao Tu, Shekoofeh Azizi, Danny Driess, Mike Schaeckermann, Mohamed Amin, Pi-Chuan Chang,  
 825 Andrew Carroll, Chuck Lau, Ryutaro Tanno, Ira Ktena, Anil Palepu, Basil Mustafa, Aakanksha  
 826 Chowdhery, Yun Liu, Simon Kornblith, David Fleet, Philip Mansfield, Sushant Prakash, Renee  
 827 Wong, Sunny Virmani, Christopher Semturs, Sara Mahdavi, Bradley Green, Ewa Dominowska,  
 828 Blaise Aguera-Arcas, Joelle Barral, Dale Webster, Greg Corrado, Yossi Matias, Karan Singhal,  
 829 Pete Florence, Alan Karthikesalingam, and Vivek Natarajan. Towards Generalist Biomedical AI.  
 830 *NEJM AI*, 2024. URL <https://ai.nejm.org/doi/full/10.1056/AIoa2300138>.

831

832 Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine  
 833 learning research*, 9(11), 2008.

834

835 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,  
 836 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Informa-  
 837 tion Processing Systems*, 30, 2017.

838

839 Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and  
 840 composing robust features with denoising autoencoders. In *Proceedings of the 25th Interna-  
 841 tional Conference on Machine Learning*, ICML '08, pp. 1096–1103. Association for Computing  
 842 Machinery, 2008. ISBN 9781605582054. URL [https://doi.org/10.1145/1390156.  
 843 1390294](https://doi.org/10.1145/1390156.1390294).

844

845 Sithara Vivek, Jessica Faul, Bharat Thyagarajan, and Weihua Guan. Explainable variational autoen-  
 846 coder (E-VAE) model using genome-wide SNPs to predict dementia. *J. Biomed. Inform.*, 148  
 847 (104536):104536, December 2023.

848

849 Sean Whalen, Rebecca M Truty, and Katherine S Pollard. Enhancer–promoter interactions are  
 850 encoded by complex genomic signatures on looping chromatin. *Nat. Genet.*, 48(5):488–496, May  
 851 2016.

852

853 Lu Yang, Sheng Wang, and Russ B Altman. Popdx: an automated framework for patient pheno-  
 854 typing across 392 246 individuals in the uk biobank study. *Journal of the American Medical  
 855 Informatics Association*, 30(2):245–255, 2023.

856

857 Jinsung Yoon, James Jordon, and Mihaela Schaar. Gain: Missing data imputation using generative  
 858 adversarial nets. In *International conference on machine learning*, pp. 5689–5698. PMLR, 2018.

859

860 Weihao Yu, Mi Luo, Pan Zhou, Chenyang Si, Yichen Zhou, Xinchao Wang, Jiashi Feng, and  
 861 Shuicheng Yan. MetaFormer is actually what you need for vision. In *Proceedings of the  
 862 IEEE/CVF conference on computer vision and pattern recognition*, pp. 10819–10829, 2022.

863

864 Shuangfei Zhai, Walter Talbott, Nitish Srivastava, Chen Huang, Hanlin Goh, Ruixiang Zhang,  
 865 and Josh Susskind. An attention free transformer, 2021. URL <https://arxiv.org/pdf/2105.14103.pdf>.

866

867 Kai Zhang, Jun Yu, Eashan Adhikarla, Rong Zhou, Zhiling Yan, Yixin Liu, Zhengliang Liu, Lifang  
 868 He, Brian Davison, Xiang Li, et al. BiomedGPT: A unified and generalist biomedical generative  
 869 pre-trained transformer for vision, language, and multimodal tasks. *arXiv e-prints*, pp. arXiv-  
 870 2305, 2023.

Wei Zhou, Zhengxiao Yan, and Liting Zhang. A comparative study of 11 non-linear regression models highlighting autoencoder, DBN, and SVR, enhanced by SHAP importance analysis in soybean branching prediction. *Sci. Rep.*, 14(1):5905, March 2024a.

Yukun Zhou, Mark A. Chia, Siegfried K. Wagner, Murat S. Ayhan, Dominic J. Williamson, Robert R. Struyven, Timing Liu, Moucheng Xu, Mateo G. Lozano, Peter Woodward-Court, Yuka Kihara, Naomi Allen, John E. J. Gallacher, Thomas Littlejohns, Tariq Aslam, Paul Bishop, Graeme Black, Panagiotis Sergouniotis, Denize Atan, Andrew D. Dick, Cathy Williams, Sarah Barman, Jenny H. Barrett, Sarah Mackie, Tasanee Braithwaite, Roxana O. Carare, Sarah Ennis, Jane Gibson, Andrew J. Lotery, Jay Self, Usha Chakravarthy, Ruth E. Hogg, Euan Paterson, Jayne Woodside, Tunde Peto, Gareth Mckay, Bernadette Mcguinness, Paul J. Foster, Konstantinos Balaskas, Anthony P. Khawaja, Nikolas Pontikos, Jugnoo S. Rahi, Gerassimos Lascaratos, Praveen J. Patel, Michelle Chan, Sharon Y. L. Chua, Alexander Day, Parul Desai, Cathy Egan, Marcus Fruttiger, David F. Garway-Heath, Alison Hardcastle, Sir Peng T. Khaw, Tony Moore, Sobha Sivaprasad, Nicholas Strouthidis, Dhanes Thomas, Adnan Tufail, Ananth C. Viswanathan, Bal Dhillon, Tom Macgillivray, Cathie Sudlow, Veronique Vitart, Alexander Doney, Emanuele Trucco, Jeremy A. Guggenheim, James E. Morgan, Chris J. Hammond, Katie Williams, Pirro Hysi, Simon P. Harding, Yalin Zheng, Robert Luben, Phil Luthert, Zihan Sun, Martin McKibbin, Eoin O'Sullivan, Richard Oram, Mike Weedon, Chris G. Owen, Alicja R. Rudnicka, Naveed Sattar, David Steel, Irene Stratton, Robyn Tapp, Max M. Yates, Axel Petzold, Savita Madhusudhan, Andre Altmann, Aaron Y. Lee, Eric J. Topol, Alastair K. Denniston, Daniel C. Alexander, and Pearse A. Keane. A foundation model for generalizable disease detection from retinal images. *Nature*, 622(7981):156–163, September 2023. ISSN 1476-4687. URL <https://doi.org/10.1038/s41586-023-06555-x>.

Zhanke Zhou, Jianing Zhu, Fengfei Yu, Xuan Li, Xiong Peng, Tongliang Liu, and Bo Han. Model inversion attacks: A survey of approaches and countermeasures. *arXiv preprint arXiv:2411.10023*, 2024b.

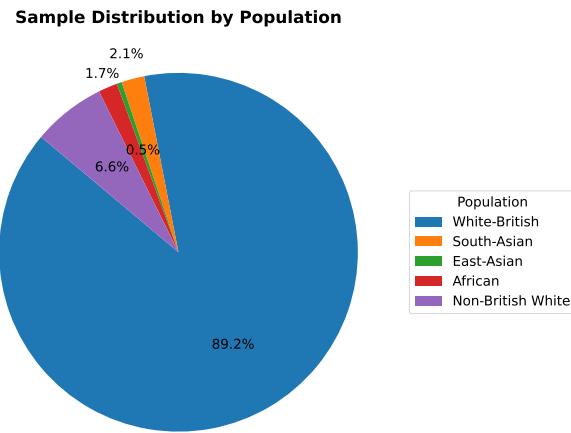
864  
865  
866  
867  
868  
869  
870  
871  
872  
873  
874  
875  
876  
877  
878  
879  
880  
881  
882  
883  
884  
885  
886  
887  
888  
889  
890  
891  
892  
893  
894  
895  
896  
897  
898  
899  
900  
901  
902  
903  
904  
905  
906  
907  
908  
909  
910  
911  
912  
913  
914  
915  
916  
917

## 918 A LIMITATIONS

920 While PM1 demonstrates strong performance across a range of multimodal inference tasks, several  
 921 limitations highlight avenues for enhancement.

923 **Data limitations** First, our dataset is primarily derived from the UK Biobank (Bycroft et al., 2018)  
 924 and EyePACS (Gulshan et al., 2016) cohorts, both of which significantly overrepresent individuals  
 925 of European ancestry. Supplementary Fig. 4 illustrates the population distribution in the UKB, high-  
 926 lighting the predominance of individuals of White British ancestry, who constitute 89.2% of the  
 927 cohort. As such, the generalizability of PM1 to non-European populations is limited, and perfor-  
 928 mance disparities may arise when applied to ancestrally diverse cohorts. Addressing this requires  
 929 future evaluation and fine-tuning on more representative datasets. Possible mitigations could also  
 930 include the use of domain adaptation techniques (Muneeb et al., 2022; Comajoan Cara et al., 2024)  
 931 or population-conditional resampling (Bonet et al., 2024).

932 Second, the input modalities—genotypes, phenotypes, and retinal images—differ markedly in terms  
 933 of data density, noise characteristics, and missingness patterns. These modality-specific imbalances  
 934 may lead to biased feature representations, especially when one modality dominates the signal.  
 935 Moreover, the imaging modality used in PM1 is limited to retinal fundus photographs. While this  
 936 choice is motivated by data availability and the known predictive power of retinal features for both  
 937 ocular and systemic conditions, it restricts the usefulness of the image modality to other types of  
 938 images. Importantly, our framework is modular and could be extended to other types such as X-rays  
 939 or MRI, pending appropriate encoder substitution and retraining.

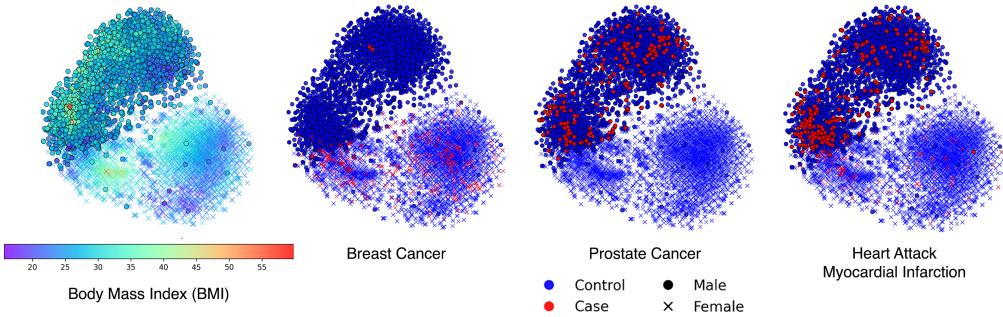


956 **Supplementary Figure 4: Population distribution of the UKB.** The cohort is heavily skewed toward  
 957 individuals of European ancestry, with 89.2% of participants self-identifying as White-British.

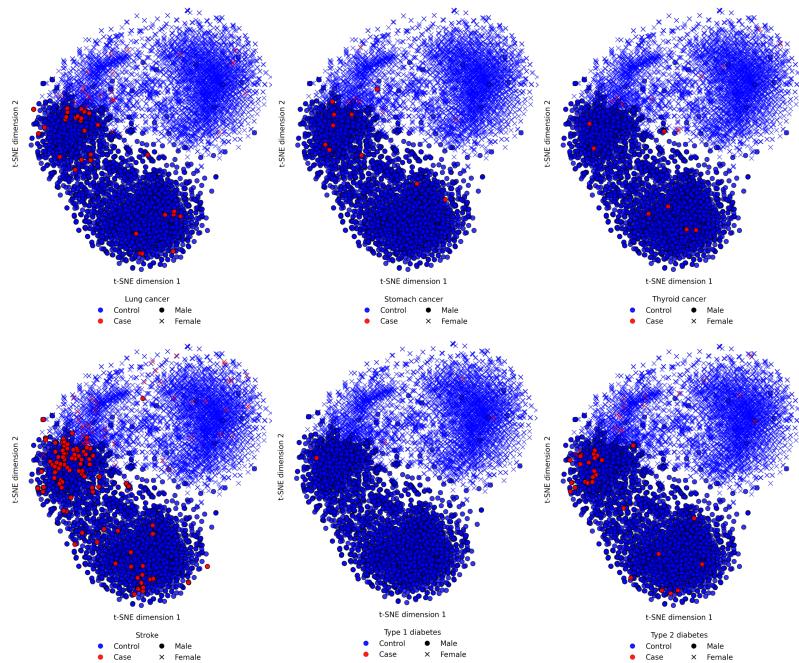
959 **Architectural limitations** Given that genetic recombination breaks linkage disequilibrium (LD)  
 960 between chromosomes—rendering variants on different chromosomes largely uncorrelated—we  
 961 adopt a chromosome-specific modeling strategy. Each of the 22 autosomes is encoded independently  
 962 using a dedicated variational autoencoder (VAE), following the approach introduced in (Geleta et al.,  
 963 2023). While this approach aligns with biological priors, it imposes limitations on modeling capac-  
 964 ity. Compressing entire chromosomes overlooks the finer functional organization of the genome,  
 965 such as gene-level or regulatory element groupings. This coarser partitioning may hinder the model’s  
 966 ability to uncover pathogenic associations that arise within or between biologically meaningful loci,  
 967 both intra- and inter-chromosomally (Whalen et al., 2016). Moreover, due to limited compute re-  
 968 sources, we choose to train PM1 in two sequential stages: first, modality-specific encoders and  
 969 decoders are trained independently to learn within-modality representations; then, their weights are  
 970 frozen before training the fusion transformer. While this design reduces computational cost, it pre-  
 971 cludes end-to-end fine-tuning across modalities and may introduce alignment mismatches. Finally,  
 972 for retinal image reconstruction, we deliberately omit the RETFound decoder. This decision is moti-

vated by two key limitations: (i) the RETFound decoder is not designed for high-fidelity generation and performed poorly in our preliminary experiments, and (ii) retraining an alternative full image autoencoder is infeasible due to the proprietary nature of RETFound’s pretraining data. Instead, we opt for a diffusion model conditioned on PM1’s fused embeddings to synthesize images, but the available data only allows us to do an exploratory analysis of this component.

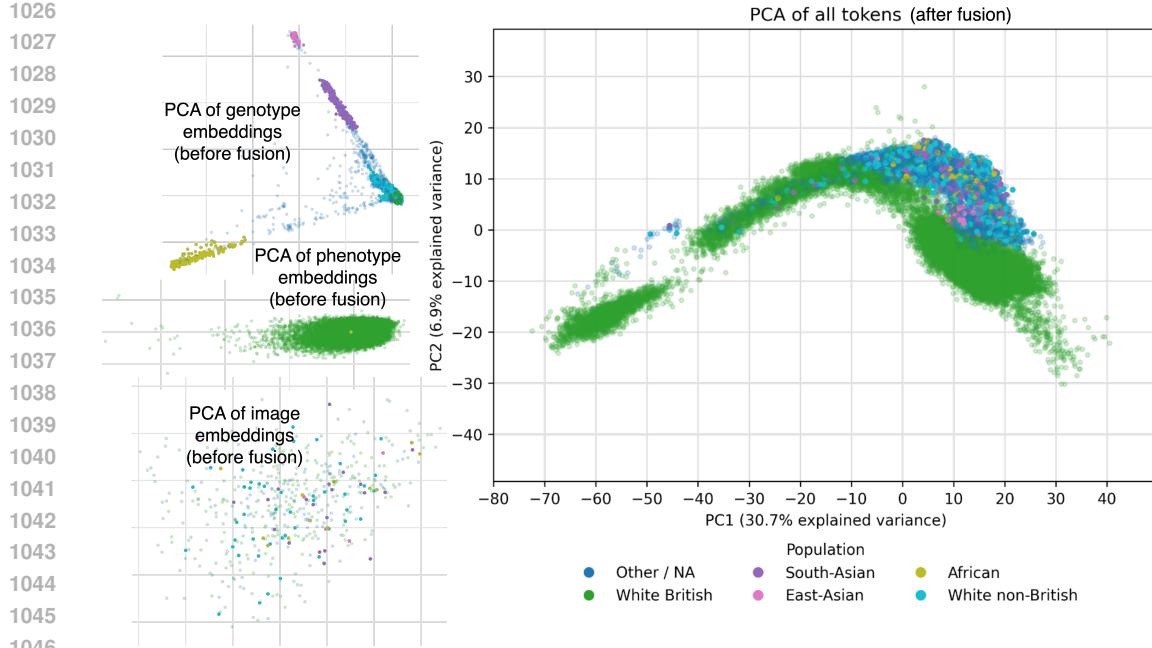
## B SUPPLEMENTARY FIGURES



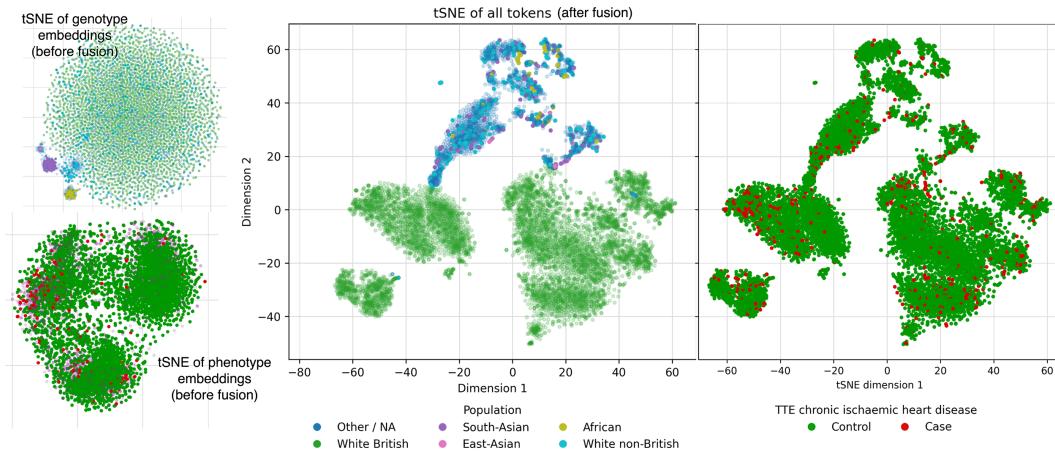
Supplementary Figure 5: **t-SNE projections of phenotype embeddings generated by  $\mathcal{E}_P$  across 1,000 UKB samples.** Observe that diseases like breast and prostate cancer project onto clearly delineated axes aligned with the self-reported biological sex covariate—while breast and prostate cancers are highly correlated with respective sexes, myocardial infarction appears across sexes but with a concentration in higher BMI regions, suggesting latent structure correlating with cardiometabolic risk.



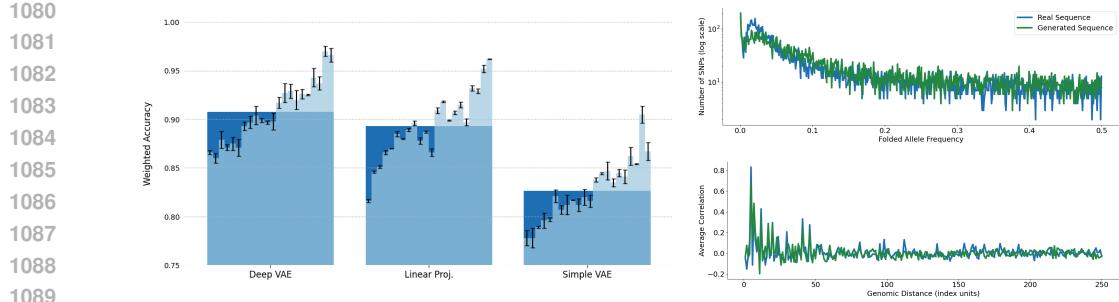
Supplementary Figure 6: **t-SNE projections of phenotype embeddings generated by  $\mathcal{E}_P$  across 1,000 UKB samples, colored by positive cases of various cancers, stroke, and diabetes.** Although covariates such as sex and body mass index (BMI) are not explicitly provided as input, the learned embeddings reflect their influence through emergent structure. Distinct clusters and phenotype-specific separations indicate that the model captures latent demographic and physiological factors associated with disease risk.



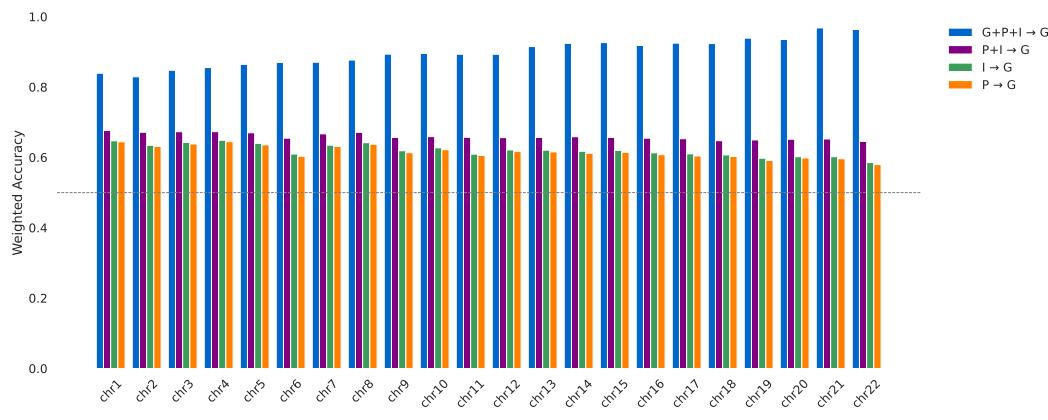
**Supplementary Figure 7: Visualization of modality-specific and fused latent representations in PM1 across >30,000 UKB samples.** The left panel shows PCA projections of pre-fusion latent embeddings, color-coded by ancestral origin. The right panel plots the post-fusion embeddings, illustrating how the transformer fusion network aligns heterogeneous modalities into a unified representation. This demonstrates PM1’s capacity to integrate semantically rich but structurally diverse inputs into a coherent latent space.



**Supplementary Figure 8: t-SNE projections of PM1’s latent space before and after modality fusion across >30,000 UKB samples.** The left panel shows t-SNE projections of genotype and phenotype embeddings independently, where genotype embeddings are color-coded by ancestry, while the phenotype embeddings are color-coded by an indicator on chronic ischaemic heart disease condition. In the middle and right panels we have the t-SNE-embedded genotype/phenotype/image tokens after fusion, color-coded by ancestry and heart disease, respectively.



Supplementary Figure 9: **Genotype reconstruction and synthesis evaluation metrics.** (Left) Per-chromosome ordered from left to right by size (light blue) and mean reconstruction accuracy (dark blue), comparing our custom genotype encoder to a linear projection model and a shallow VAE. (Right) Assessment of the realism of synthetic genotypes generated by our genotype Encoder. We compare the folded allele frequency spectrum and the LD decay—measured via correlation with neighboring variants—between real (blue) and synthetic (green) genotypes.



Supplementary Figure 10: **Cross-modal genotype inference across all 22 autosomes under different modality combinations.** Each group of bars corresponds to a specific chromosome (chromosome 1 through chromosome 22), and each bar within a group represents a distinct cross-modal input configuration used to reconstruct the genotype sequence for that chromosome.

### C ANCESTRY INFERENCE AS A PROBING TASK

To demonstrate the usefulness of the learned representations, we performed an additional downstream evaluation focused on continental-ancestry inference. Starting with UKB participants' self-reported ancestry, we defined a target variable consisting of three ancestral classes, specifically: European (comprising *White British* samples), African, and Asian (including both *East Asian* and *South Asian*). To mitigate class imbalance we randomly subsampled 10,000 (from a total of 35,181) European individuals while retaining every African (587) and Asian (853) sample, yielding a cohort we split into 9,152 train and 2,288 test samples. For each individual we extracted the pooled Fusion embeddings produced by PM1 and concatenated them into a single latent representation (a 6,144-dimensional vector that integrates genotype, phenotype, and retinal information) and trained four off-the-shelf standard machine learning classifiers—logistic regression, random forest, naive Bayes, and  $k$ NN (with  $k \in \{3, 10, 25\}$ ).

All models surpassed  $>99\%$  test and balanced accuracy, macro-F1  $>99\%$ , and weighted-F1  $>99\%$ , confirming that ancestral groups are cleanly separable in the learned embedding space. These results demonstrate that PM1 captures population-level structure without any ancestry supervision and its representations transfer effectively to an independent classification task.

1134 **D ETHICAL IMPLICATIONS AND SOCIETAL IMPACTS**  
11351136 The development and deployment of multimodal biomedical foundation models like PM1 raises im-  
1137 portant ethical considerations that warrant careful examination. As these models integrate highly  
1138 sensitive genetic, clinical, and imaging data at unprecedented scale, we must address both the trans-  
1139 formative potential for advancing precision medicine and the risks of perpetuating or amplifying  
1140 existing healthcare disparities.1141  
1142 **D.1 POTENTIAL RISKS**  
11431144 **Privacy and Re-identification Risks** PM1 encodes highly sensitive personal health  
1145 data—including genome sequences, phenotypic traits, and medical images—into dense multi-  
1146 modal representations. Although UKB data is de-identified at the source (Bycroft et al., 2018),  
1147 concerns remain regarding the risk of re-identification from model outputs or learned repres-  
1148 entations. In particular, model inversion attacks, which aim to reconstruct original training data  
1149 by exploiting access to a trained model, could potentially leak sensitive information (Rigaki &  
1150 Garcia, 2023; Zhou et al., 2024b). For instance, (Fredrikson et al., 2014) demonstrated that phar-  
1151 macogenetic models could leak sensitive genetic markers. However, successfully reconstructing  
1152 entire high-dimensional training samples, such as full genomes, from model outputs or embeddings  
1153 presents substantial challenges. Recent surveys on model inversion attacks highlight that successful  
1154 model inversion attacks on complex, high-dimensional data typically rely on exploiting specific  
1155 model vulnerabilities or having access to substantial prior knowledge, rather than a simple decoding  
1156 of learned representations (Rigaki & Garcia, 2023; Zhou et al., 2024b).1157 The model’s cross-modal inference capabilities also raise particular concerns about genetic discrim-  
1158 ination (Moreau, 2019). Although legislation like the Genetic Information Nondiscrimination Act  
1159 (GINA) (110th United States Congress, 2008) offers some protection, PM1’s ability to infer genetic  
1160 information from non-genetic modalities (e.g., predicting SNPs from retinal images or phenotypes)  
1161 could potentially circumvent these safeguards (Shi & Wu, 2017). Furthermore, as foundation models  
1162 like PM1 compress data from numerous individuals, ensuring data privacy and compliance with reg-  
1163 ulations becomes paramount. This includes the challenge of implementing methods for selectively  
1164 “forgetting” patient data upon request (Carrillo-Perez et al., 2024).1165 **Representational Bias and Health Disparities** PM1 is trained predominantly on UKB data (By-  
1166 croft et al., 2018), which overrepresents individuals of European ancestry. This creates a concrete  
1167 risk of performance disparities: the model’s predictions are likely to be less accurate for individuals  
1168 from underrepresented populations (Martin et al., 2019). Since disease prevalence, presentation,  
1169 and genetic architecture vary across ancestries, a model trained primarily on European data may  
1170 produce biased clinical predictions for non-European individuals, potentially exacerbating existing  
1171 healthcare inequities. While (Carrillo-Perez et al., 2024) suggest that synthetic data generated by  
1172 such models *after* training might offer a way to balance datasets, the initial training on biased data  
1173 remains a core challenge that could lead to the perpetuation of existing biases.1174 Furthermore, the quality and nature of the input data significantly impact model performance and  
1175 the reliability of its outputs. As highlighted by (Carrillo-Perez et al., 2024) a key concern arises  
1176 when data *missing not at random* (MNAR)—that is, the likelihood of data being missing is related  
1177 to its actual value or other unobserved factors—models trained on such data may perpetuate biases  
1178 or generate inaccurate imputations. For instance, if certain tests are more frequently performed on  
1179 sicker patients, the model might learn a skewed representation of the general population, impacting  
1180 the fidelity of generated or imputed data for less represented or healthier subgroups.1181 **D.2 POSITIVE APPLICATIONS**  
11821183 Despite these concerns, multimodal foundation models like PM1 offer significant potential benefits  
1184 for advancing healthcare and biomedical research:1185 **Democratizing Access and Reducing Burden** PM1’s ability to infer information across modal-  
1186 ities can enhance accessibility to advanced medical insights. By inferring genetic risk from more  
1187 readily available and less expensive data, such as retinal images or clinical records, it could reduce

1188 reliance on costly or invasive procedures. Indeed, it has been proposed that such models could  
 1189 obviate the need for invasive procedures by imputing desired information from already collected  
 1190 non-invasive data. Moreover, they offer a way to lessen the economic strain associated with acquir-  
 1191 ing certain medical modalities. This could democratize aspects of precision medicine, especially  
 1192 in resource-limited settings, and make complex diagnostics more widely available (Carrillo-Perez  
 1193 et al., 2024).

1194

1195 **Enhancing Early Diagnosis and Personalized Medicine** Multimodal integration, as demon-  
 1196 strated by PM1, enables a more comprehensive and holistic assessment of disease risk. Combining  
 1197 genetic predispositions with subtle changes in imaging or clinical data could lead to earlier identi-  
 1198 fication of at-risk individuals, often before clinical symptoms manifest, aligning with the goals of  
 1199 precision medicine (Acosta et al., 2022).

1200

1201 **Accelerating Research and Discovery** The integrated representations learned by PM1 can signif-  
 1202 icantly accelerate biomedical research. The model’s capacity for cross-modal inference might help  
 1203 reveal novel genotype-phenotype associations and guide hypothesis generation. PM1 has the  
 1204 potential to be used for “*in silico* hypothesis testing”, (Carrillo-Perez et al., 2024), where researchers  
 1205 could study the effects of altering specific features in one modality on others within a simulated  
 1206 environment. This, coupled with improved data imputation and the generation of high-quality syn-  
 1207 thetic samples, can enhance the diversity and availability of data for research, particularly in scarce-  
 1208 data settings (Carrillo-Perez et al., 2024). The integration of these data is expected to substantially  
 1209 improve our comprehension of human health and enable the development of more precise, personal-  
 1210 ized approaches to prevention, diagnosis, and treatment (Karczewski & Snyder, 2018; Acosta et al.,  
 1211 2022).

1212

1213

1214 PM1, as introduced in this paper, exemplifies the increasing capability of multimodal foundation  
 1215 models to contribute to precision medicine. The integration of diverse biomedical data sources  
 1216 offers clear avenues for enhanced diagnostic insight and accelerated research. However, such ad-  
 1217 vancements are intrinsically linked with significant ethical considerations, including data privacy,  
 1218 security, and the imperative to avoid perpetuating health inequities through biased representations.  
 1219 It is crucial that ongoing and future research in this field embeds rigorous ethical frameworks and  
 1220 fosters interdisciplinary collaboration to navigate these complexities and guide the responsible de-  
 1221 velopment of these impactful technologies.

1222

1223

## E DATA PROCESSING

1224

1225

### E.1 DATA SPLIT

1226

1227

We adopt a fixed 80/10/10 split of our 521,269 samples for training, validation, and testing, using a  
 1228 random seed of 42. This split is applied to the outer join of all available samples across modalities,  
 1229 ensuring consistent partitioning for the fusion model. For training the modality-specific encoders,  
 1230 we use the same global split and select the corresponding 80%, 10%, and 10% subsets of each  
 1231 individual modality that fall into the respective buckets. This strategy ensures alignment between  
 1232 modality-specific and fusion training while preventing data leakage across evaluation stages. Sup-  
 1233 plementary Tables 2, 3, and 4 show summary statistics of genotype (G), image (I), and phenotype (P)  
 1234 data availability across training, validation, and test splits, including sample counts and missingness  
 1235 for each disease phenotype.

1236

1237

### E.2 GENOTYPE DATA

1238

1239

1240

1241

For the genotype data, we use UKB SNP genotype arrays spanning 487,409 individuals, each rep-  
 1242 resented with 658,720 variants. The SNP arrays in the UK Biobank are significantly shorter than  
 1243 the total number of known variants (Halldorsson et al., 2022), as not all positions were initially  
 1244 sequenced. For this initial version of the foundation model, we decide to rely on the shorter SNP  
 1245 arrays, which are more widely used and readily accessible.

1242

1243

Supplementary Table 2: Modality missingness per data split.

1244

1245

Split	$n_{\text{samples}}$ (% of total)	$n_{\text{geno}}$ (missing)	$n_{\text{img}}^*$ (missing left, missing right)	$n_{\text{pheno}}$ (missing**)
Train	417,015 (80%)	350,753 (66262)	91,212 (left: 333,124, right: 331,585)	269,617 (147,398)
Val	52,127 (10%)	44,055 (8072)	11,354 (left: 41,709, right: 41,524)	33,771 (18,356)
Test	52,127 (10%)	43,860 (8267)	11,385 (left: 41,676, right: 41,504)	33,741 (18,386)
SUM	521,269 (100%)			

1246

1247

\* Number of samples for which at least one retina scan is available.

1248

1249

\*\* Number of samples for which no phenotype data is available. Phenotype-specific missingness can be found in the phenotype rows.

1250

1251

Supplementary Table 3: Phenotype (eye diseases) counts per split shown as case (positive), total (case and control), and missing.

1252

1253

1254

1255

1256

1257

Split	Macular degeneration			Diabetic eye			Glaucoma			Cataract		
	Case	Total	Miss	Case	Total	Miss	Case	Total	Miss	Case	Total	Miss
train	433	380,857	7	1,021	380,269	7	6,256	126,401	248,640	16,640	125,092	239,565
val	50	47,753	1	141	47,662	1	782	15,780	31,242	1,996	15,624	30,184
test	49	47,669	0	132	47,586	0	769	15,754	31,195	2,053	15,594	30,071

1258

1259

1260

1261

1262

1263

1264

Supplementary Table 4: Phenotype (cardiovascular diseases and neurodegenerative diseases) counts per split shown as case (positive), total (case and control), and missing.

1265

1266

1267

1268

1269

Split	Heart failure			Ischaemic stroke			Heart attack			Dementia			Alzheimer’s disease		
	Case	Total	Miss	Case	Total	Miss	Case	Total	Miss	Case	Total	Miss	Case	Total	Miss
train	7,082	37,4208	7	4,557	376,734	6	15,802	365,488	7	436	380,855	6	745	380,546	6
val	898	46,905	1	613	47,190	1	1,966	45,837	1	57	47,746	1	101	47,702	1
test	903	46,815	0	584	47,134	0	2,017	45,701	0	61	47,657	0	89	47,629	0

1270

1271

1272

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295

1296

1297

1298

1299

1300

1301

1302

1303

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

1350

1351

1352

1353

1354

1355

1356

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403

1404

1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

1458

1459

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

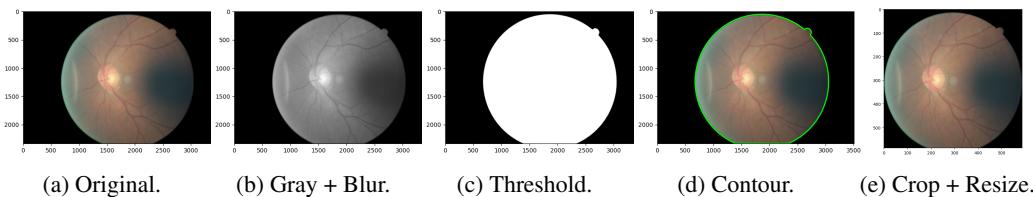
1481

1296 contour encompassing the largest area is heuristically considered to be the retina. Following the  
 1297 identification of the retinal contour, its bounding box is determined, and from this, the center of the  
 1298 retina is computed. The dimensions for cropping are then established by taking the maximum of the  
 1299 bounding box's width and height, and incorporating an additional 5-pixel margin on each side.

1300 The original color image is then cropped according to these calculated dimensions, carefully ensuring  
 1301 that the cropped area is contained within the original image boundaries and that the retina is  
 1302 positioned centrally. As a final step in this phase, the cropped image is resized to a target resolution  
 1303 of  $587 \times 587$  pixels and stored in PNG format. It is important to emphasize that the initial grayscaling,  
 1304 blurring, and thresholding operations are performed on a temporary version of the image solely  
 1305 for the purpose of accurate retina localization; the color information of the final CFP image remains  
 1306 unaltered throughout this process.

1307 An incidental outcome of this preprocessing methodology was the automatic exclusion of certain  
 1308 images. Specifically, if no contour could be detected, the image was omitted from further processing.  
 1309 Manual examination of these discarded images revealed them to be entirely black, rendering their  
 1310 exclusion advantageous for the subsequent analysis. Given the extensive size of the dataset and the  
 1311 author's limited domain expertise in retinal pathology, no additional manual curation or filtering of  
 1312 the images was undertaken.

1313 Subsequent to these initial preprocessing steps, and prior to their input into the machine learning  
 1314 models, the images undergo further transformations facilitated by the torchvision library (Ansel  
 1315 et al., 2024). These transformations consist of resizing the images to a more compact resolution,  
 1316 converting them into PyTorch tensor format, and normalizing their pixel values. This normalization,  
 1317 which adjusts pixel values to achieve a zero mean and unit variance, is a common practice that  
 1318 typically enhances the stability and convergence rate of model training.



1328 **Supplementary Figure 11: Demonstration of the image preprocessing pipeline.** The pipeline is  
 1329 applied to the CFP image of the left retina of sample 498 in the EyePACS dataset (Dugas et al.,  
 1330 2015; Gulshan et al., 2016).

## E.5 FUSION OF MODALITIES

1335 The multimodal dataset consists of  $N_G = 487,409$  genotype samples,  $N_P = 476,569$  phenotype  
 1336 samples, and  $N_I = 211,416$  retina image samples, of which 97,465 include both eyes and 113,951  
 1337 include only one. To maximize data utilization, we perform an outer join across all modalities in-  
 1338 stead of an inner join, which would restrict the dataset to only 47,432 samples, to align the data.  
 1339 Naturally, this implies that many samples will have missing modalities. To allow the fusion model  
 1340 to handle such missing data, we pre-compute a binary data mask  $M_{\text{data}} \in \{0, 1\}^T$  for each sample,  
 1341 indicating the presence (1) or absence (0) of each sequence position. Missing tokens are replaced  
 1342 with a learned mask embedding. In addition, we define a loss mask  $M_{\text{loss}} \in \{0, 1\}^T$  for each sam-  
 1343 ple, which identifies the token positions used in the masked modeling objective. For each sample,  
 1344 the loss mask  $M_{\text{loss}} \in \{0, 1\}^{25}$  is sampled from a Bernoulli distribution with success probability  
 1345  $p_{\text{mask,mod}_i, \text{seq}_j}$ , independently of all sequence positions.

1346 Let  $n_{\text{mod}} = 3$  denote the number of modalities: genotype ( $G$ ), phenotype ( $P$ ), and image ( $I$ ). Each  
 1347 modality  $\text{mod}_i \in \{G, P, I\}$  is masked with equal probability (Supplementary Equation 4):

$$p_{\text{mask,mod}_i} = \frac{1}{n_{\text{mod}}} = \frac{1}{3} \quad (4)$$

1350 Each modality contains a different number of tokens:  $n_{\text{seq},G} = 22$ ,  $n_{\text{seq},P} = 1$ , and  $n_{\text{seq},I} = 2$ .  
 1351 To evenly distribute the modality masking probability across its tokens, we define the per-token  
 1352 masking probability as (Supplementary Equation 5):  
 1353

$$p_{\text{mask,mod}_i,\text{seq}_j} = p_{\text{mask,mod}_i} \cdot \frac{1}{n_{\text{seq,mod}_i}} \quad (5)$$

1357 This results in per-token masking probabilities of approximately 0.015 for each genotype sequence  
 1358 position, 0.333 for the phenotype sequence position, and 0.167 for each image sequence position.  
 1359 The loss mask is resampled independently for each sample in every epoch.

1360 Before feeding the modality-specific embeddings into the fusion model—or reconstructing them  
 1361 using the respective decoders—we standardize each modality to have zero mean and unit variance to  
 1362 mitigate the impact of scale mismatches between modalities, which can lead to imbalanced gradient  
 1363 contributions during loss optimization. This normalization is performed using the mean and standard  
 1364 deviation computed on the training set only, ensuring no information leakage into the validation or  
 1365 test sets.  
 1366

## F MODEL SPECIFICATIONS

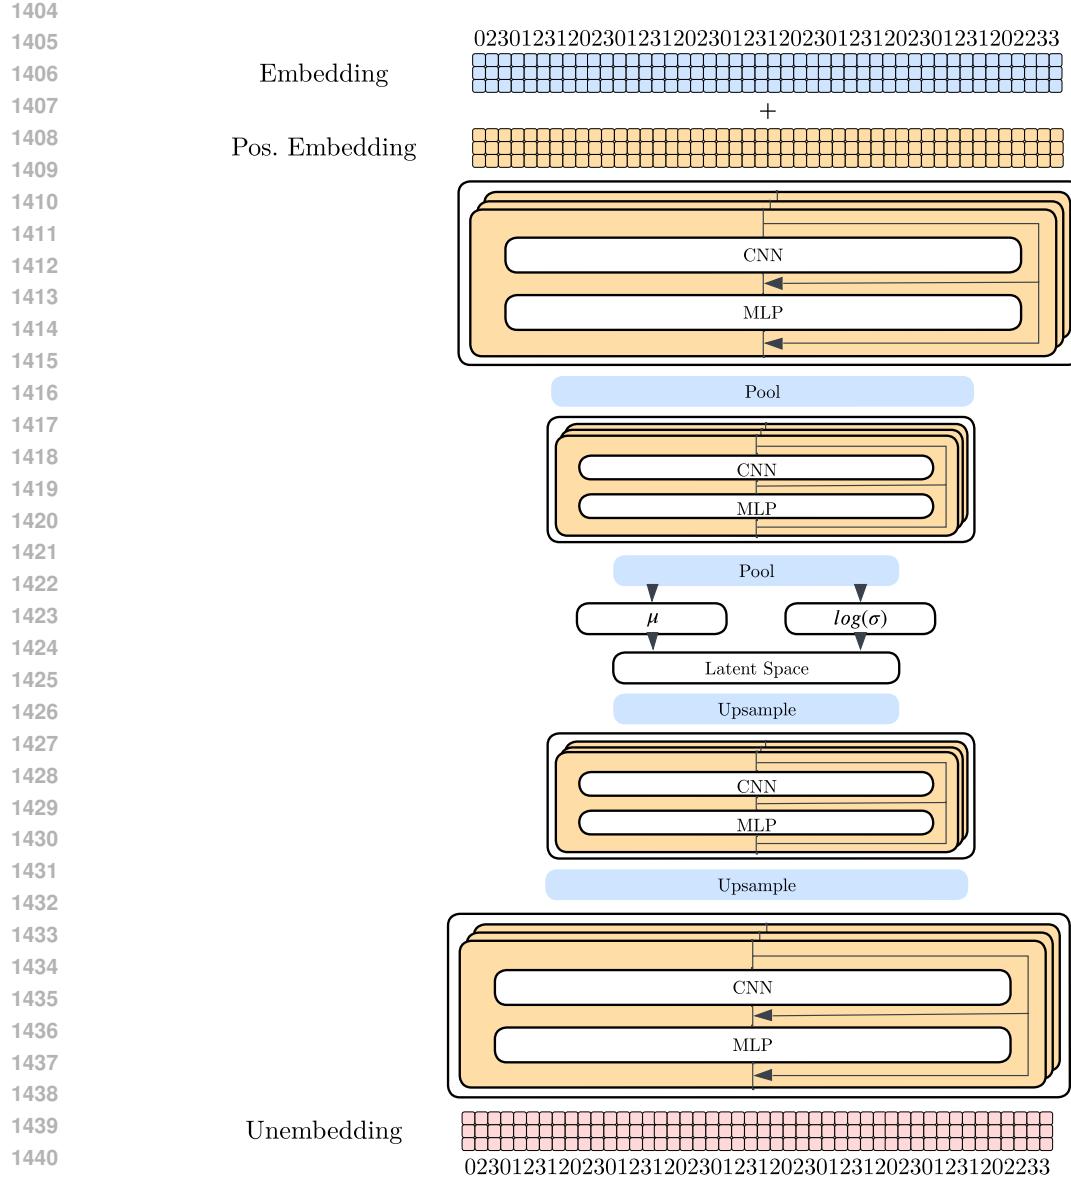
1367 This section provides a detailed overview of the model architecture, training protocol, and hyper-  
 1368 parameter choices for PM1’s modality-specific encoders and decoders, and the transformer-based  
 1369 fusion network, as used throughout our experiments. We emphasize transparency and reproducibil-  
 1370 ity by documenting our design rationale and training infrastructure.  
 1371

### F.1 GENOTYPE ENCODER

1372 We design a modality-specific encoder-decoder architecture for genotype data, implemented as a  
 1373 VAE. A schematic overview of the model architecture is provided in Supplementary Fig. 12. The  
 1374 architecture leverages a compact and expressive latent space of dimensionality  $d_{\text{latent}} = 2048$ ,  
 1375 optimized to capture high- and low-level genomic features while minimizing information loss. Each  
 1376 SNP position is embedded into a  $d_{\text{model}} = 64$ -dimensional vector using the sum of learned token  
 1377 and positional embeddings (Vaswani et al., 2017). For a chromosome with  $S_c$  SNPs, the embedded  
 1378 input lies in  $\mathbb{R}^{S_c \times 64}$ .  
 1379

1380 The encoder  $\mathcal{E}_G$  and decoder  $\mathcal{D}_G$  consist of  $\text{num\_blocks} = 3$  blocks, each composed of  
 1381  $\text{block\_depth} = 3$  residual sub-layers. These blocks are separated by spatial resolution-changing  
 1382 layers (downsampling in the encoder and upsampling in the decoder) to construct a hierarchical  
 1383 representation, reflecting the hierarchical organization of genomic data (Crick, 1970). The number  
 1384 of channels changes by a factor of  $\text{seq\_change\_rate} = 2$  at each resolution change. To ensure  
 1385 compatibility with the hierarchical structure of the model, the input sequence of each chromosome  
 1386 is padded to the next larger number divisible by  $\text{seq\_change\_rate}^{\text{num\_blocks}} = 2^3 = 8$ . The  
 1387 reconstructed output for the padded sequence positions is discarded for evaluation.  
 1388

1389 Inspired by recent attention-free Transformer alternatives (Yu et al., 2022; Zhai et al., 2021), each  
 1390 residual sub-layer is composed of two modules: a token-mixer and a channel-mixer. The token-  
 1391 mixer is implemented as a sequence of four depthwise convolutional layers. The first and last layers  
 1392 preserve the input channel dimension, while the two intermediate layers expand the channel size by  
 1393 a factor of  $\text{conv\_mult} = 1.5$  before reducing it back, enhancing local receptive fields as seen in  
 1394 RNNs (Bai et al., 2018). This structure is particularly well-suited for capturing local dependencies in  
 1395 genomic sequences such as those introduced by linkage disequilibrium (LD). The channel-mixer is  
 1396 a three-layer feed-forward MLP (Rumelhart et al., 1986), where the hidden layer width is increased  
 1397 by a factor of  $\text{mlp\_mult} = 2.0$  relative to the input and output dimensions. Residual connections  
 1398 (He et al., 2016) are included in both mixers to stabilize training, while normalization layers such  
 1399 as batch norm (Ioffe & Szegedy, 2015) are intentionally excluded due to observed instabilities.  
 1400 The output of the final encoder block is flattened and passed through a linear projection layer that  
 1401 maps it to the parameters of a Gaussian distribution: a mean vector  $\mu$  and a log-variance vector  
 1402  $\log \sigma$ , each of dimension  $d_{\text{latent}} = 2048$ . The latent representation  $\mathbf{z}$  is then sampled using the  
 1403 reparameterization trick (Kingma & Welling, 2013; Rezende et al., 2014), enabling backpropagation



Supplementary Figure 12: Schematic overview of genotype VAE architecture.

through the stochastic sampling step. The reconstructed output, after running through the decoder  $\mathcal{D}_G$ , is projected back into the original SNP input space using a learned un-embedding matrix, producing logits for each SNP.

The training loss used to train the genotype encoder entails a weighted cross-entropy term to correct for allelic imbalance and a KL divergence regularizer scaled by  $\text{beta\_kl} = 0.0913$ , in accordance with the  $\beta$ -VAE formulation (Higgins et al., 2016). KL divergence is annealed over the first  $\text{kl\_anneal\_steps} = 1000$  steps to prevent divergence of the KL divergence loss within the first few training steps.

The training process uses a cyclic learning rate schedule (CyclicLR) with a triangular mode. The learning rate cycles between  $\text{base\_lr} = 1.21 \times 10^{-6}$  and  $\text{max\_lr} = 3 \times 10^{-6}$  with a step size of  $\text{step\_size} = 2000$  updates. As the optimizer we use Adam (Ma & Yarats, 2018) with the default hyperparameters  $\text{betas} = (0.5, 0.999)$ . Mixed-precision training is enabled using the `float16` datatype to allow for larger batch sizes ( $\text{batch\_size} = 32$ ) on available hardware.

1458 To avoid instability during training, gradient clipping is applied with a maximum norm of  
 1459 `max_norm` = 1.0 and steps in which the KL divergence loss includes NaN values are skipped.  
 1460 Training runs for up to `epochs` = 20 epochs, with early stopping applied if no validation improve-  
 1461 ment is observed for `patience` = 5 epochs.

1462 Training is conducted on a single NVIDIA A100 GPU for large chromosomes (1–10), and paral-  
 1463 lelized across up to four NVIDIA A5500 GPUs using data parallelism for smaller chromosomes.  
 1464 Training times range from 13.22 hours (chromosome 1, 13 epochs) to 6.84 hours (chromosome 22,  
 1465 20 epochs), depending on sequence length and convergence behavior.

## 1467 F.2 PHENOTYPE ENCODER

1468 We instantiate  $(\mathcal{E}_P, \mathcal{D}_P)$  as a denoising auto-encoder with 5 fully-connected layers on the encoder  
 1469 side and a symmetric decoder. Layer widths decrease linearly from 3,421 inputs to a 2,048-  
 1470 dimensional bottleneck (`depth=5`, `latent_dim=2048`). Each linear layer is followed by a  
 1471 LeakyReLU activation ( $\alpha = 0.01$ ). Model parameters are optimized with stochastic gradient de-  
 1472 scent (SGD); gradient norms are clipped to 10. Training runs on a single NVIDIA A5500 (24  
 1473 GB) GPU with early stopping after 50 epochs without improvement in validation loss, which cor-  
 1474 responds to roughly 24-hour wall-clock time. During training, every input row is corrupted with  
 1475 copy-masking with probability 0.3, where the observed/missing pattern of a randomly chosen donor  
 1476 row from the training set is applied and the newly masked entries are set to zero. During evaluation  
 1477 for a given target phenotype, the column corresponding to that phenotype and all other phenotypes  
 1478 that are directly related are fully masked.

1479 Hyperparameters were selected via Hyperband (Li et al., 2018) with up to 150 trials minimizing  
 1480 validation loss. The search space comprised: learning rate  $\in [10^{-3}, 10^{-1}]$  (log-uniform), batch  
 1481 size  $\in \{512, 1024, 2048\}$ , depth  $\in \{4, \dots, 9\}$ , and momentum  $\in \{0.8, 0.9, 0.95\}$ . The best con-  
 1482 figuration (learning rate =  $4.78 \times 10^{-2}$ , batch size = 512, depth = 5, and momentum = 0.8, was  
 1483 adopted for all results reported in the paper.

## 1485 F.3 IMAGE ENCODER

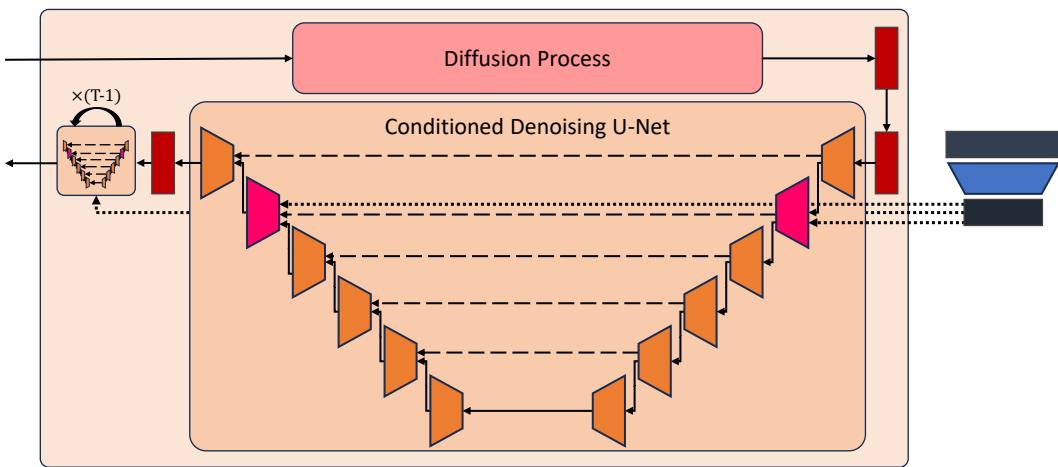
1486 PM1’s retinal image encoder is adapted from the pre-trained RETFound model (Zhou et al., 2023).  
 1487 This encoder is a Vision Transformer (ViT) (Dosovitskiy et al., 2020), specifically configured to  
 1488 process inputs at a resolution of 224x224 pixels, utilizing a patch size of 16x16 pixels. Archi-  
 1489 tecturally, it comprises 24 Transformer blocks, each with an embedding dimension of 1024. The  
 1490 pre-training regimen for this encoder involved an initial phase on natural images sourced from the  
 1491 ImageNet-1k dataset (Deng et al., 2009). This was subsequently followed by continued training on  
 1492 an extensive collection of 904,170 CFP retinal images via masked image modeling (He et al., 2022).  
 1493 The majority of these retinal images were drawn from the proprietary Moorfields Diabetic Image  
 1494 Dataset (MEH-MIDAS), as detailed in the RETFound publication (Zhou et al., 2023). The specific  
 1495 pre-trained checkpoint leveraged in PM1 was acquired from the HuggingFace Hub, corresponding  
 1496 to the repository ID YukunZhou/RETFound\_mae\_natureCFP.

1497 We adapt the RETFound image encoder (checkpoint `RETFound_mae_natureCFP.pth`) to UKB  
 1498 fundus images using by inserting LoRA adapters (Hu et al., 2022) (rank  $r=8$ ,  $\alpha=16$ , dropout 0.1)  
 1499 into the self-attention projections (`qkv`) and the output projection (`proj`) in all ViT blocks with  
 1500 the PEFT (Mangrulkar et al., 2023) library. Base weights remain frozen; only LoRA parameters  
 1501 are trainable. We train for 10 epochs on the UKB training split (inputs  $224 \times 224$ , ImageNet nor-  
 1502 malization), using AdamW (learning rate  $1 \times 10^{-4}$ ) and batch size 256 with the original masked-  
 1503 autoencoding reconstruction loss with mask ratio 0.75 and without any labels. After training, we  
 1504 merge the adapters into the encoder weights. The adapted encoder produces left/right tokens  $z^{(I_L)}$ ,  
 1505  $z^{(I_R)}$  consumed by the fusion transformer. The MAE decoder is not used within PM1; diffusion  
 1506 models handle image generation.

## 1508 F.4 DIFFUSION MODEL

1509 The diffusion model employed for retinal image synthesis, as explored in Section 5, implements  
 1510 the Denoising Diffusion Implicit Models (DDIM) scheduler (Song et al., 2022) for the diffusion  
 1511 process, incorporating the specific modifications detailed in (Lin et al., 2024), and the denoising

1512 process is based on a conditional U-Net architecture (Ronneberger et al., 2015), drawing inspiration  
 1513 of previous works (Rombach et al., 2022). We utilize 1000 diffusion steps during training and 50  
 1514 denoising steps during inference.  
 1515



1516  
 1517  
 1518  
 1519  
 1520  
 1521  
 1522  
 1523  
 1524  
 1525  
 1526  
 1527  
 1528  
 1529  
 1530  
 1531 **Supplementary Figure 13: Overview diagram of the diffusion component architecture.** The  
 1532 model is composed of a forward diffusion process (DDIM (Song et al., 2021)), and a conditioned  
 1533 denoising U-Net (Ronneberger et al., 2015) which is applied for  $T$  timesteps to reverse the diffusion  
 1534 process. The trapezoids shrinking from right to left are downsampling blocks, and the ones growing  
 1535 are upsampling blocks. The pink ones, additionally, incorporate cross-attention (Vaswani et al.,  
 1536 2017) between the hidden states of the model and the embedding of the condition. The dashed lines  
 1537 represent skip connections, and the dotted lines the conditioning with the embedding. A linear layer  
 1538 (in blue on the right) is used to obtain the embedding used for conditioning.  
 1539

1540 Our implementation of the U-Net (Ronneberger et al., 2015) consists of six downsampling blocks  
 1541 and six symmetrically structured upsampling blocks. Skip connections are utilized to concatenate  
 1542 feature channels from the contracting path to the corresponding layers in the expanding path. The  
 1543 sequence of feature channels produced at the output of these blocks is 128, 128, 256, 256, 512, and  
 1544 512, respectively (these are visually distinguished as orange or magenta in Supplementary Fig. 13).  
 1545 The model is designed to process 3-channel RGB 224x224 images. Except for the second downsam-  
 1546 pling block, each such block comprises a convolutional layer with a stride of 2 (to achieve spatial  
 1547 dimension reduction), which is then followed by two ResNet blocks (He et al., 2016) incorporating  
 1548 SiLU activation functions (Hendrycks & Gimpel, 2016).

1549 To enable conditional image generation, the U-Net architecture integrates cross-attention layers  
 1550 (highlighted in magenta in Supplementary Fig. 13). These layers attend to key and value projec-  
 1551 tions derived from the embeddings generated by PM1's fusion model. This mechanism effectively  
 1552 conditions the image synthesis process on the integrated multimodal representations. The fusion  
 1553 embeddings are projected to a dimensionality of 768, matching that of the attention keys and values,  
 1554 through a linear layer (indicated in blue on the right of Supplementary Fig. 13) which also utilizes  
 1555 a SiLU activation function (Hendrycks & Gimpel, 2016). To guide the diffusion process via these  
 1556 embeddings, we employ classifier-free guidance (Ho & Salimans, 2021).

1557 The model parameters were optimized using the AdamW optimizer (Loshchilov & Hutter, 2019)  
 1558 for a total of 100 epochs. The learning rate was managed by a cosine annealing schedule coupled  
 1559 with a linear warmup phase (Loshchilov & Hutter, 2017). To ensure numerical stability and prevent  
 1560 issues like exploding gradients, gradient clipping with a maximum norm of 1 was applied, and  
 1561 gradient accumulation was performed over 10 steps. Furthermore, following practices from models  
 1562 such as Stable Diffusion XL (Podell et al., 2024), an Exponential Moving Average (EMA) of the  
 1563 model weights is maintained (Ruppert, 1988; Polyak & Juditsky, 1992). To optimize computational  
 1564 performance, the diffusion model's attention mechanisms leverage the FlashAttention-2 algorithm  
 1565 (Dao, 2024), implemented via the xFormers library (Lefauze et al., 2022). To further mitigate  
 1566 memory demands during the training of the diffusion model, mixed-precision training with `bf16` is  
 1567 used.

1566 The diffusion model was trained using data parallelism distributed across 4 NVIDIA A100 GPUs,  
 1567 which provided a cumulative total of 320 GB of GPU memory. The decision to not pursue a larger  
 1568 or more complex model architecture was predicated on the available dataset size and computational  
 1569 resource limitations. Moreover, the primary goal of this phase was an exploratory investigation into  
 1570 the feasibility and potential of conditional retinal image synthesis using the described multimodal  
 1571 inputs, rather than an exhaustive optimization of model scale or performance.

1572

1573

1574

## 1575 F.5 FUSION MODEL

1576

1577

1578 The PM1 fusion model is instantiated as a transformer encoder architecture that operates over tok-  
 1579 enized embeddings derived from genotype, phenotype, and retinal image modalities. Specifically,  
 1580 we use a 32-layer transformer with a hidden size of  $d_{\text{model}} = 2048$ , 8 attention heads per layer, and  
 1581 ReLU non-linearities. While GELU activations are common in BERT-style transformers (Devlin  
 1582 et al., 2019), we opt for ReLU due to its more favorable compute-to-throughput ratio (Ming et al.,  
 1583 2022). Each token stream is linearly projected into this shared hidden space, with input dimen-  
 1584 sionalities being 2048 for both genotypes and phenotypes, and 1024 for fundus image embeddings.

1585

1586

1587

1588 Input sequences are composed of 22 chromosome-level genotype embeddings, a single phenotype  
 1589 token, and two image tokens (left and right eye), with respective sequence lengths of 22, 1, and 2. A  
 1590 learnable positional embedding (Gehring et al., 2017) with values initialized from a truncated normal  
 1591 distribution  $\mathcal{N}(0, 0.02^2)$ , where values are drawn from a normal distribution centered at zero with a  
 1592 standard deviation of 0.02 and clipped to remain within two standard deviations. This initialization  
 1593 ensures stable early training dynamics while enabling the model to capture relative positional de-  
 1594 pendencies, particularly important for modalities with inherent structure such as SNP ordering along  
 1595 chromosomes or bilateral symmetry in fundus images. We prepend a learnable [CLS] token and  
 1596 optionally mask selected tokens as part of the training objective. Input corruption is implemented  
 1597 through Gaussian noise (standard deviation  $\sigma = 0.08$ ) and token dropout ( $p = 0.4$ ), reflecting a de-  
 1598 noising autoencoder framework and aligning with the strategies proposed in (Vincent et al., 2008).  
 1599 Token corruption is sampled from a categorical distribution over four modes: no corruption (20%  
 1600 probability), additive Gaussian noise (25%), token dropout (25%), and combined noise + dropout  
 1601 (30%). This heterogeneously noised input encourages the model to learn robust representations un-  
 1602 der varying perturbation levels. In parallel, a token-level masking strategy is applied independently  
 1603 to each modality with a uniform masking probability of 0.5. Masked tokens are replaced with a  
 1604 shared learnable [MASK] vector, as in (He et al., 2022), and the model is trained to reconstruct their  
 1605 clean embeddings through  $\mathcal{L}_{\text{mask}}$ .

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

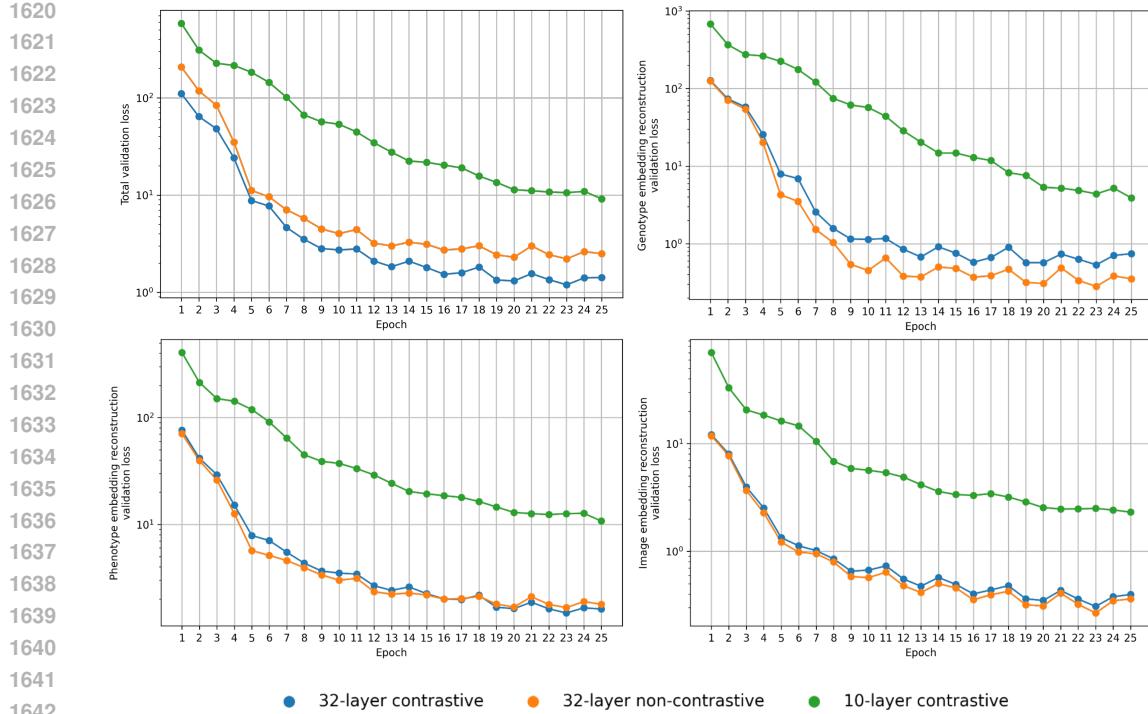
1618

1619

The total training objective is a weighted sum of three terms: masked token reconstruction loss  
 $\mathcal{L}_{\text{mask}}$ , denoising loss  $\mathcal{L}_{\text{recon}}$ , and a cross-modal contrastive loss  $\mathcal{L}_{\text{InfoNCE}}$  as detailed in the main text.  
 Each loss term contributes equally to the optimization signal. The contrastive loss is scaled by a  
 factor  $\lambda = 0.5$  and uses a softmax temperature  $\tau = 0.03$ .

Training proceeds for 25 epochs with a batch size of 128 using the quasi-hyperbolic variant of the  
 Adam optimizer (Ma & Yarats, 2018). We use a base learning rate of  $5 \cdot 10^{-5}$  and a weight decay of  
 $10^{-4}$ . The optimizer's quasi-hyperbolic momentum terms are set to  $\nu_1 = 0.7$  and  $\nu_2 = 1.0$ . A cosine  
 learning rate schedule is used with 10 epochs of linear warm-up starting from 10% of the base rate.  
 A learning rate floor of  $10^{-6}$  is imposed to avoid collapse. While we considered dynamic learning  
 rate adjustment via plateau detection, the cosine scheduler offered smoother convergence for our  
 setup. To prevent exploding gradients, we enforce max-norm gradient clipping with a threshold of  
 10. Mixed precision training was explored but ultimately disabled due to loss gradient instability  
 when scaling with batch size. However, we anticipate that future training runs could benefit from  
 mixed precision.

All experiments were monitored continuously using logging tools. Checkpoints were saved at every  
 epoch, including full model weights, optimizer states, and metrics (refer to Supplementary Fig. 14  
 for validation curves using different fusion model configurations). Training was conducted on a  
 compute cluster equipped with NVIDIA A100 and A5500 GPUs and 150 GB RAM for data loading  
 and preprocessing. For the full fusion model, a single A100 was used per training run. Validation  
 was performed on a held-out stratified 10% split of the data.



Supplementary Figure 14: **Validation loss curves across model variants.** Total and modality-specific reconstruction validation losses are shown over training epochs for three fusion model configurations: a 32-layer transformer with contrastive loss (blue), a 32-layer transformer without contrastive loss (orange), and a smaller 10-layer contrastive model (green). The top-left panel shows overall validation loss, while the remaining plots show per-modality losses for genotype, phenotype, and image embeddings, respectively. The 32-layer contrastive model consistently achieves the lowest total loss, suggesting better overall integration across modalities.

## G BENCHMARK SETUP

### G.1 PHENOTYPE MASKING

Supplementary Table 5 shows a taxonomy of target phenotypes including their corresponding related phenotypes that were masked during prediction to prevent label leakage. Each row shows the target phenotype, a brief description, and the set of semantically or clinically related traits that were excluded from the input.

### G.2 MEDGEMMA

MedGemma (Sellergren et al., 2025), is a checkpoint of the multimodal Gemma 3 (Team et al., 2025) foundation model fine-tuned for medical applications. Relative to the base Gemini model, MedGemma benefits from continued pretraining on a broad internal medical corpus, including retinal fundus images, dermatology images across 210 skin conditions, histopathology patch-text pairs, and 2D radiology slices (Team et al., 2025). All following benchmarks are conducted with the largest available 27 billion parameter multimodal MedGemma checkpoint to ensure a fair comparison.

We benchmark MedGemma on phenotype prediction (ROC-AUC) on a representative subset of 1,000 samples from our test set across 9 diseases and systemic conditions. We run inference using the HuggingFace checkpoint (Sellergren et al., 2025) on 2 A100 GPUs (80 GB VRAM). With an average runtime of 64.4 s per sample across nine phenotypes, inference on the full test set of 52,127 participants would require  $\approx 8,392$  hours ( $\approx 1$  year). Given the strong class imbalance (on average only 739.7 cases per phenotype in the test set, see Supplementary Tables 3 and 4), we instead evaluate on a representative 1,000-participant subset per phenotype, containing all (or up to 500)

1674  
1675  
1676  
1677  
1678  
1679  
1680  
1681  
1682  
1683  
1684  
1685  
1686  
1687  
1688  
1689  
1690  
1691  
1692  
1693  
1694  
1695  
1696  
1697  
1698  
1699  
1700  
1701  
1702  
1703  
1704  
1705  
1706  
1707  
1708  
1709  
1710  
1711  
1712  
1713  
1714  
1715  
1716  
1717  
1718  
1719  
1720  
1721  
1722  
1723  
1724  
1725  
1726  
1727

Supplementary Table 5: Taxonomy of target phenotypes including their corresponding related phenotypes that were masked during prediction to prevent label leakage.

Target Phenotype	Description	Masked Phenotypes
Macular degeneration	A progressive eye disease that causes vision loss in the center of the visual field.	Macular degeneration; Which eye(s) affected by macular degeneration (Right eye); Which eye(s) affected by macular degeneration (Left eye); Eye problems/disorders Macular degeneration; Age macular degeneration diagnosed
Diabetic eye disease	Eye damage resulting from diabetes, including retinopathy and macular edema.	Diabetic eye disease; Which eye(s) affected by diabetes-related eye disease (Right eye); Which eye(s) affected by diabetes-related eye disease (Left eye); Eye problems/disorders Diabetic eye disease; Age when diabetes-related eye disease diagnosed
Eye problems/disorders Glaucoma	A group of eye conditions that damage the optic nerve, often due to high intraocular pressure.	Glaucoma; Eye problems/disorders Glaucoma; Ever had surgery for glaucoma or high eye pressure; Ever had laser treatment for glaucoma or high eye pressure; Which eye(s) affected by glaucoma (Right eye); Which eye(s) affected by glaucoma (Left eye); TTE glaucoma; Age glaucoma diagnosed; Retinal problem; Retinal detachment; Retinal artery/vein occlusion; TTE chorioretinal inflammation; TTE retinal detachments and breaks; TTE retinal vascular occlusions; TTE other retinal disorders; TTE retinal disorders in diseases classified elsewhere
Eye problems/disorders Cataract	Clouding of the lens in the eye leading to a decrease in vision.	Cataract; Eye problems/disorders Cataract; Ever had cataract surgery; Which eye(s) are affected by cataract (Right eye); Which eye(s) are affected by cataract (Left eye); TTE senile cataract; TTE other cataract; Age cataract diagnosed
Heart failure	A chronic condition where the heart doesn't pump blood as well as it should.	Heart failure/pulmonary oedema; TTE heart failure
AD ischaemic stroke	A type of stroke caused by an obstruction within a blood vessel supplying blood to the brain.	AD ischaemic stroke; Transient ischaemic attack (TIA); TTE transient cerebral ischaemic attacks and related syndromes; TTE cerebral infarction; TTE occlusion and stenosis of precerbral arteries, not resulting in cerebral infarction; TTE stroke, not specified as haemorrhage or infarction; Age stroke diagnosed; TTE other acute ischaemic heart diseases; TTE chronic ischaemic heart disease
Heart attack	Also known as myocardial infarction, occurs when blood flow to the heart is blocked.	Heart attack/myocardial infarction; AD myocardial infarction; Vascular/heart problems diagnosed by doctor Heart attack; TTE acute myocardial infarction; TTE subsequent myocardial infarction; Age heart attack diagnosed; TTE other acute ischaemic heart diseases; TTE chronic ischaemic heart disease
TTE dementia in Alzheimer's disease	A dementia subtype caused by Alzheimer's disease pathology.	AD Alzheimer's disease; Alzheimer's disease/dementia; TTE dementia in Alzheimer's disease; TTE Alzheimer's disease
Alzheimer's disease	A progressive neurodegenerative disorder and the most common cause of dementia.	AD Alzheimer's disease; Alzheimer's disease/dementia; TTE dementia in Alzheimer's disease; TTE Alzheimer's disease; AD all cause parkinsonism; AD Parkinson's disease

cases and the remainder controls. Since metrics such as ROC-AUC depend on ranking rather than prevalence, this provides a reasonable yet representative simplification that reduces compute time to  $\approx 161$  hours.

For phenotype-to-phenotype prediction ( $P \rightarrow P$ ), we curate a natural language prompt from the input phenotypes per sample, masking any traits related to the target according to the masking protocol detailed in [Supplementary Section G.1](#), and instruct the model to predict the probability of the sample to be be diagnosed with the target phenotype. We subsequently extract the prediction from the answer string and compute an ROC-AUC score. For the multimodal setting ( $P+I \rightarrow P$ ), we also included fundus images of both eyes, resized and patched together into a single  $896 \times 896$  image and normalized to the  $[-1, 1]$  range, as required by MedGemma, in the prompt. Again, the model was then prompted to produce a probability reflecting the likelihood that the individual is diagnosed with the target phenotype.

### G.3 CONTIG

We pretrain ContIG following Taleb et al. (2022) using the UK Biobank raw SNP data and retinal fundus images. During pretraining, the model maximizes agreement between image and genetic representations using a contrastive loss. For downstream evaluation, we fine-tune the pretrained encoder on each phenotype prediction task.