Computer Methods and Programs in Biomedicine 277 (2026) 109250

Contents lists available at ScienceDirect

Computer Methods and Programs in Biomedicine

journal homepage: www.sciencedirect.com/journal/computer-methods-and-programs-in-biomedicine

Shape matters: Predicting Huntington's disease using progression modelling

Mohsen Ghofrani-Jahromi [a], Susmita Saha [a], Adeel Razi [a], Pubu M. Abeyasinghe [a], Govinda R. Poudel [b], Jane S. Paulsen [c], Sarah J. Tabrizi [d], Nellie Georgiou-Karistianis [a,*]

[a] Turner Institute for Brain and Mental Health, Monash University, Clayton, VIC 3800, Australia
[b] Mary MacKillop Institute for Health Research, Australian Catholic University, Melbourne, VIC 3000, Australia
[c] Department of Neurology, University of Wisconsin-Madison, 1685 Highland Avenue, Madison, Wisconsin, USA
[d] UCL Huntington's Disease Centre, UCL Queen Square Institute of Neurology, UK Dementia Research Institute, Department of Neurodegenerative Diseases, University College London, London, UK

ARTICLE INFO

Keywords:
Huntington's Disease
Biomarkers
Neuroimaging
Clinical Trials
Subcortical Shape
Deep Learning

ABSTRACT

Background: Despite evidence of group-level differences in striatal morphometry among persons with Huntington's Disease (PwHD), current models of HD progression used for participant selection and assessment of treatment outcomes in clinical trials do not leverage shape information.
Methods: We first validated the capability of a discriminative deep neural network to derive descriptors of shape from all subcortical structures affected by HD, utilizing 2,932 brain scans in 615 PwHD across three longitudinal datasets (TRACK-HD, PREDICT-HD, and IMAGE-HD). We then trained a conditional generative model that used shape descriptors, alongside conventional volumetric, genetic, as well as composite cognitive, motor, and functional features at baseline to predict biomarkers of disease progression at subsequent time points.
Results: We observed that the anatomical shapes of subcortical structures, including putamen, lateral ventricle, pallidum, caudate, thalamus, and accumbens, exhibited strong associations with HD progression, as measured by a commonly used prognostic score. Furthermore, within-stage heterogeneity, along the continuum of disease progression, was better captured: when shape descriptors were aggregated using principal component analysis, they showed a high correlation with disease stage (Spearman's correlation: $\rho = 0.72$), compared to volumetric measurements in cubic millimetres ($\rho = 0.45$). Finally, incorporating subcortical shape into the generative model improved predictive performance, compared to the same model that relied solely on brain volumes.
Conclusion: This study demonstrates that subcortical brain shape is associated with HD progression, enables capturing fine-grained within-stage variability, and improves the predictability of characteristic biomarkers. The findings could potentially optimize future clinical trials through more targeted participant recruitment and more objective post-intervention assessments of treatment efficacy.

1. Introduction

Efforts to discover disease-modifying treatments for Huntington's disease (HD) critically rely on stratifying trial participants into cohorts with sufficiently homogeneous severity of disease progression [1]. Specifically, overcoming heterogeneity of participants in clinical trials is essential for an objective assessment of how successful a designed intervention is in achieving its endpoints [2,3]. While the conventional approach to evaluating drug efficacy involves comparing treatment and placebo arms with respect to the clinical endpoints, machine learning (ML) offers a potentially more objective evaluation by providing subject-specific predictions of how the outcome measure would manifest without intervention [4,5]. This approach could be particularly useful for persons with (PwHD), where large natural history datasets now offer

Abbreviation: CAG, cytosine-adenine-guanine; CAP, CAG-Age Product; cUHDRS, Composite Unified Huntington's Disease Rating Scale; cVAE, conditional variational autoencoder; DL, Deep Learning; ML, machine learning; non-HGEC, non-huntingtin gene expansion carriers; PIN, normalized prognostic index; PwHD, Persons with HD; SDMT, Symbol Digit Modalities Test; SWR, Stroop Word Reading Test; TMS, Total Motor Score; t-SNE, t-distributed Stochastic Neighbour Embedding.

* Corresponding author at: Turner Institute for Brain and Mental Health, 18 Innovation Walk, Monash University, Clayton, VIC 3800, Australia.
E-mail address: nellie.georgiou-karistianis@monash.edu (N. Georgiou-Karistianis).

https://doi.org/10.1016/j.cmpb.2026.109250
Received 16 April 2025; Received in revised form 14 December 2025; Accepted 14 January 2026
Available online 15 January 2026

a rich set of information, enabling such modelling to yield significant insights.

Various techniques are proposed for modelling the trajectories of HD progression, including hidden Markov [6,7], autoregressive [8], event-based [9,10], Gaussian process regression [11], mixed effect [12,13], and disease progression models [14,15]. These computational models have been trained using features such as the genetic burden, motor and cognitive assessment biomarkers, as well as image-derived volumetric measurements obtained from automated segmentation of structural magnetic resonance imaging (MRI) [16,17].

Conversely, numerous studies have explored the associations between HD progression and the morphology of atrophy in subcortical nuclei. For instance, during early HD stages, shape displacements have been observed in the tail of the caudate nucleus, the medial side of the left putamen, the dorsal part of the pallidum, and the accumbens nucleus [18]. Moreover, Younes et al. observed substantial early-stage atrophy in the caudate, with relative preservation of the dorsolateral and dorsomedial regions of its mid-body [19]. They also identified significant atrophy in the anterior dorsomedial section of the left putamen and in the globus pallidus in later-stage individuals, with atrophic changes more pronounced in elevated CAP scores (CAG age product). From a functional perspective, Tang et al. found that subregions of caudate and putamen connected to limbic and executive cortices exhibited atrophy, especially in higher CAP groups [20]. Furthermore, extensive shape analyses of the striatum, hippocampus, and thalamus have been conducted using manually segmented MRI from IMAGE-HD [21–24]. However, such findings about progressive deformation in subcortical shapes have not yet been utilized in predictive models of HD.

Incorporating features of anatomical shape as inputs to ML models trained with large sample sizes is non-trivial. Deep learning (DL) [25] allows for automated feature extraction, by enabling the discovery of underlying representations of anatomical deformations throughout the disease course. For example, PointNet [26] is a deep neural network capable of directly processing a set of points on the surface of a segmented brain structure. It can derive optimal representations of anatomical shapes informed by the severity of disease progression. These representations can then be utilized for stratification, diagnosis or prognosis, as was done in other neurodegenerative conditions [27–30].

With the emergence of various DL-based generative models, researchers have utilized them for synthesizing different modalities of brain imaging not only for data augmentation but also for understanding aging and disease-related mechanisms, and as tools for generating digital twins for clinical trials [31–33]. Moreover, conditional generative models have been proposed to predict follow-up characteristic biomarkers using baseline measurements [34]. In such models, the observed data is assumed to be governed by an underlying latent (hidden) space, where model training involves mapping the inputs onto latent representations disentangled into subject-specific and age-related components [35]. For instance, Choi et al. followed this approach to predict age-related brain metabolic degeneration by training a model that generates a follow-up PET image from baseline [36]. Chai et al. predicted future cortical thickness maps by synthetizing the age-conditioned trajectory of cortical morphology [37].

In this study, we employed DL, for the first time in HD, to develop effective descriptors of shape for subcortical brain structures that exhibited strong anatomical associations with disease progression across the continuum of PwHD. Subsequently, we integrated such shape descriptors along with conventional HD biomarkers, summarized in the composite Unified Huntington's Disease Rating Scale (cUHDRS), to develop and evaluate a conditional generative model for predicting disease progression at subsequent longitudinal points. We considered the practical necessity for a predictive model potentially employed for clinical trials to rely on a single brain scan at baseline.

The present study should be viewed primarily as a methodological proof of concept with clear translational potential. It presents an end-to-end framework with two core components: one for shape-informed feature learning and one for individualized forecasting of HD progression. Each component of the pipeline could be further refined, providing a flexible foundation for integrating anatomical shape into prognostic models of HD. We hypothesized that incorporating shape descriptors into a prognostic model would improve the accuracy of predicting disease progression compared to using volumetric measures alone. The significance of the proposed model lies in its capacity to deliver individualized predictions for biomarkers of HD progression, thereby aiding in the detection of treatment effects on a subject-specific basis.

2. Materials and methods

2.1. Data preparation

We incorporated data from cohorts described in previous studies: TRACK-HD [14,38], TrackON-HD [39], PREDICT-HD [40,41], and IMAGE-HD [42,43]. Each dataset featured repeated visits, during which T1-weighted MRI was acquired from PwHD, alongside genetic and clinical assessments, detailed in the Supplementary materials.

We included 3Tesla MRI form the pooled datasets comprising a total of 2,932 brain scans form 615 PwHD. Fig. 1(A) shows the histogram for the number of MRI scans per individual, and Fig. 1(B) shows the distribution of interval between scans. While on average approximately 3 scans were acquired (SD = 1.94), there is substantial variability, with many participants contributing between 1 and 6 timepoints, within a mean interval of 13 months (SD = 6) between scans. FreeSurfer [44] (version 6.0) was used for MRI pre-processing and segmentation of eight structures including the lateral ventricle, thalamus, caudate, putamen, pallidum, hippocampus, amygdala, and accumbens. The quality control procedure was detailed in our previous work [12].

We used the normalized prognostic index (PIN), a measure of HD progression status [45], as the target variable for training PointNets. Our goal was to determine whether a PointNet could derive shape descriptors from a segmented structure that correlated with the concurrent PIN scores.

Subsequently, we trained a conditional generative model to predict changes from baseline to follow-up in outcome measures commonly used in clinical trials: the volumes of the caudate, putamen, and lateral ventricle, as well as the cUHDRS [46]. We considered the cUHDRS in our prognostic modelling, as it combines multiple domains into a single score and has recently been used as an intermediate endpoint in HD clinical trials due to its sensitivity [47]. Moreover, the volume of the caudate nucleus is among the earliest structural biomarkers, showing consistent atrophy years before clinical onset [48]. We also included lateral ventricle volume as it is a robust and reliable imaging biomarker with lower susceptibility to segmentation noise. It is commonly used as a surrogate endpoint in neurodegeneration studies due to its strong correlation with global brain atrophy, particularly in later disease stages [49].

We used the CAP_{100} score ($CAP_{100} = Age \times \frac{CAG-30}{6.49}$), a measure of genetic burden [50], along with sex and age as the condition vector. Table 1 summarizes the characteristics of the PwHD, separated by the HD integrated staging system (HD-ISS) [51]. The HD-ISS online calculator was used, and k-nearest neighbours (kNN) imputation was applied to handle missing values. It is important to note that the HD-ISS was not incorporated into any of our modelling schemes but was used solely for reporting and visualization purposes.

2.2. Deriving descriptors of shape from segmented structures

PointNet [26] is a discriminative neural network that serves as a simple yet powerful universal approximator. It processes a set of N unordered points on the surface of a segmented structure, schematically shown in Fig. 2 (A) and (B), identically and independently treating each point $p_i = (x_i, y_i, z_i)$ in the 3D space, in two steps:

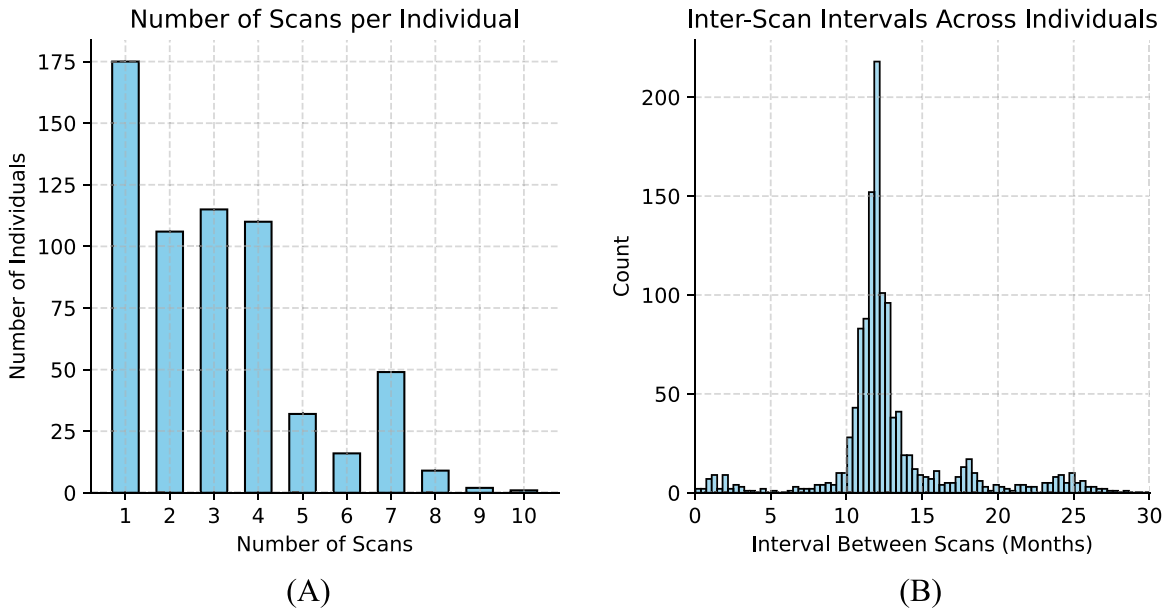

Fig. 1. (A) Histogram for the number of acquired scans from individuals (B) Distribution of inter-scan intervals.

Table 1
Demographics and clinical scores in the pooled dataset separated by HD-ISS. N shows the number of longitudinally acquired images in each stage.

HD-ISS	Stage 0	Stage 1	Stage 2	Stage 3
Age	40.6 ± 11.3	43.3 ± 10.3	46.5 ± 9.6	50.9 ± 9.8
CAP_{100}	73.9 ± 15.4	82.1 ± 12.6	90.3 ± 11.8	104.1 ± 13.4
PIN	-0.41 ± 0.68	-0.01 ± 0.56	0.98 ± 0.86	2.71 ± 1.25
cUHDRS	17.3 ± 1.8	17.0 ± 1.3	15.1 ± 1.9	10.2 ± 3.2
N	516	836	851	729

CAP: CAG age product, cUHDRS: composite Unified Huntington's Disease Rating Scale, PIN: Prognostic Index. Values represent mean ± standard deviation.

1) Mapping to a high-dimensional space: Each input point p_i is mapped onto a high-dimensional space using a function, $h(p_i)$. In practice, $h(.)$ is implemented by N fully connected multilayer perceptrons (MLPs) that share similar weights and take $P_{N\times 3}$ as input and give $H_{N\times F}$ as output, where F is the length of the shape descriptor vector.
2) Applying a symmetric operator: A symmetric operator, denoted as $g(.)$, which is essentially invariant to the permutation of its inputs is applied column-wise on the matrix $H_{N\times F}$ produced from step 1 above. This operator is typically implemented by a *max pooling* layer, which outputs a $1\times F$vector. Hence, from N input points, F shape descriptors that are dominant in the high-dimensional space are selected with each corresponding input point called a critical point.

Mathematically, these two steps can be described as:

$$f(\{p_1,\ldots,p_N\}) = g(h(p_1),\ \ldots, h(p_N)) \quad (1)$$

For the complete PointNet architecture, refer to the original paper [26]. While we employed the entire network in our implementation, for brevity, Fig. 2(C) does not include the first module in the network where points are transformed by a 3×3 matrix to account for rotation to a consistent orientation. We used PyTorch [52] (version 2.3) to train eight independent PointNets, one per each segmented structure, with the concurrent PIN score as the target variable. From the total 2,932 images, 80% were used for the training set, 10% for the validation set to examine overfitting, and 10% in the test set.

Within each epoch of training, 25% of the points situated on the surface of a structure were randomly sampled, then adjusted to relocate the centroid to the origin, and subsequently normalized to a unit sphere by dividing all distances from the centroid by the furthest point. To maintain simplicity in the models, we focused our analyses on the left hemisphere, as existing literature indicates that while lateralization may not be a significant factor in HD, the left hemisphere is more prominently affected by neurodegenerative changes [53,54], thereby allowing for reduced computational complexity. The trained networks were evaluated using a test set of unseen individuals.

We applied principal component analysis (PCA) and t-distributed Stochastic Neighbour Embedding (t-SNE) to both subcortical volumetric measures (in mm^3) and aggregated shape descriptors to compare how these two domains reflect disease progression and to assess whether shape descriptors demonstrate superior sensitivity.

2.3. A conditional generative model of disease progression

This section describes the training of a conditional variational autoencoder (cVAE) [55] designed to predict biomarkers of disease progression at follow-up based on their baseline values [37,56]. The model was trained to specifically estimate changes in these biomarkers ($\boldsymbol{\Delta B}$) over time. To benchmark the shape-informed model (shown in Fig. 3), we also employed: (1) a naïve predictor assuming no change at follow-up ($\Delta\boldsymbol{B}=0$), and (2) a cVAE that used only baseline cUHDRS and volumes, without shape descriptors.

Assuming that the observed data is governed by a latent generative process, the joint conditional probability is given by:

$$\boldsymbol{p}(\boldsymbol{X},\boldsymbol{z}|\boldsymbol{c}) = \boldsymbol{p}(\boldsymbol{X}|\boldsymbol{z},\boldsymbol{c})\boldsymbol{p}(\boldsymbol{z}) \quad (2)$$

where $\boldsymbol{X}$ is a vector comprising shape descriptors ($\boldsymbol{F}$) derived using a stack of PointNets along with biomarkers of interest ($\boldsymbol{B}$), $\boldsymbol{z}$ is the latent space representations, and $\boldsymbol{c}$ is the condition vector comprising confounding variables: [age, sex, and the CAP score].

To estimate the latent variables $\boldsymbol{z}$ given the observed data $\boldsymbol{X}$ and condition vector $\boldsymbol{c}$, we need to approximate the posterior distribution $p(\boldsymbol{z}|\boldsymbol{X},\ \boldsymbol{c})$, which is intractable to compute directly. Approximate variational inference provides a practical solution by using a simpler distribution, called the variational posterior, to approximate the true posterior. This surrogate distribution, denoted as $q_\theta(\boldsymbol{z}|\boldsymbol{X},\boldsymbol{c})$, is often chosen to be a Gaussian distribution (known as a Laplace approximation). In this case, the parameters $\theta=[\boldsymbol{\mu},\ \boldsymbol{\Sigma}]$ represent the mean ($\boldsymbol{\mu}$) and covariance ($\boldsymbol{\Sigma}$) of the Gaussian. The goal is to find the values of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ that make$q_\theta(\boldsymbol{z}|\boldsymbol{X},\boldsymbol{c})$ as close as possible to the true posterior $p(\boldsymbol{z}|\boldsymbol{X},\ \boldsymbol{c})$. For

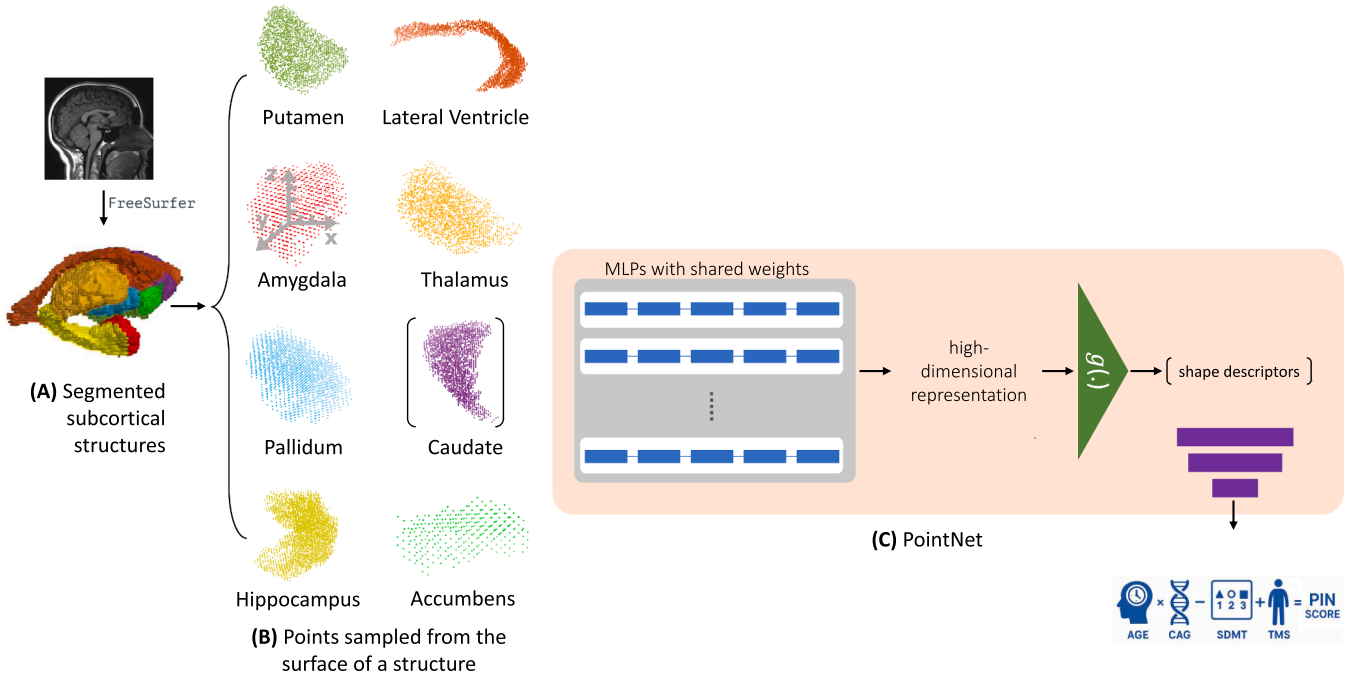

Fig. 2. (A) Binary arrays of segmented volumes were derived using FreeSurfer. (B) For each structure, points on the surface were randomly sampled and fed into the PointNets. Eight PointNets were independently trained. (C) Each PointNet constitutes fully connected multilayer perceptrons (MLPs) that map the input points onto a high-dimensional space, then applies a symmetric operator to derive global descriptors of shape. Finally, another fully connected layer estimates the PIN score from the shape descriptors.

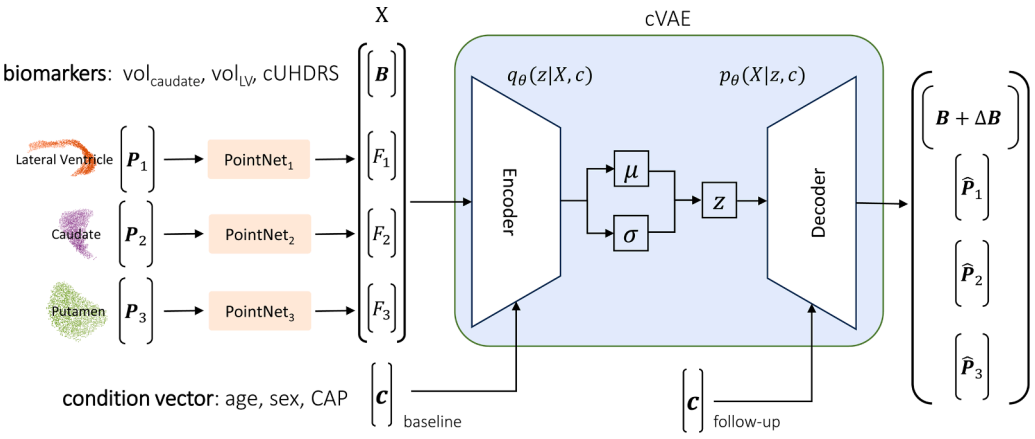

Fig. 3. Training the conditional variational autoencoder (cVAE) involves finding optimal weights for the encoder, the decoder, as well as the two PointNets. The input points, P, from the caudate, putamen, and the lateral ventricle are fed into the corresponding PointNets to obtain two vectors of shape descriptors. The encoder should estimate the parameters of the latent space distribution from the concatenation of shape descriptors and biomarkers B which includes the cUHDRS and volumes of caudate, putamen, and lateral ventricle, and the condition vector c all at baseline. The decoder should reconstruct the desired output, based on condition vector $\boldsymbol{c}$ at follow-up.

a detailed derivation of this approach and the associated KL divergence term, we refer the reader to Kingma and Welling [57].

Consequently, training a cVAE (Fig. 3) involves simultaneous optimization of the weights of stacked PointNets, the encoder which infers the latent code from the input, as well as the decoder which reconstructs the output given the latent representations and the condition vector.

In each epoch during training, we minimized a combination of the point cloud reconstruction loss ($\mathscr{L}_{rec}$), the biomarker prediction loss ($\mathscr{L}_{pred}$), and the latent representation loss ($\mathscr{L}_{latent}$):

$$\mathscr{L} = \boldsymbol{w}_r\mathscr{L}_{rec} + \boldsymbol{w}_p\mathscr{L}_{pred} + \boldsymbol{w}_l\mathscr{L}_{latent} \tag{3}$$

To quantify the point cloud reconstruction loss, we used the Chamfer distance [58] between the target point cloud $\boldsymbol{P}$ and the reconstructed point cloud $\widehat{\boldsymbol{P}}$ generated by the model. It penalizes the model if the shape of $\widehat{\boldsymbol{P}}$ does not resemble the desired shape. This is computed as the sum of Euclidean distances between each point in one set and the nearest point in the other set:

$$\mathscr{L}_{rec} = \boldsymbol{d}_{CH}(\boldsymbol{P}, \widehat{\boldsymbol{P}}) = \sum_{\boldsymbol{p}\in\boldsymbol{P}} \min_{\widehat{\boldsymbol{p}}\in\widehat{\boldsymbol{P}}} (\| \boldsymbol{p} - \widehat{\boldsymbol{p}}\|^2) + \sum_{\widehat{\boldsymbol{p}}\in\widehat{\boldsymbol{P}}} \min_{\boldsymbol{p}\in\boldsymbol{P}} (\| \boldsymbol{p} - \widehat{\boldsymbol{p}}\|^2) \tag{4}$$

$\mathscr{L}_{pred}$ was chosen as the mean squared error (MSE) between the actual and estimated $\boldsymbol{\Delta B}$.

The latent representation loss enforces the variational posterior to match a normal distribution by minimizing the Kullback-Leibler (KL) divergence between them:

$$\mathscr{L}_{latent} = D_{KL}(q_\theta(\boldsymbol{z}|\boldsymbol{X},\boldsymbol{c}),\mathcal{N}(0,\boldsymbol{I})) = -\frac{1}{2}\sum_{j=1}^{J} 1 + \log(\sigma_j)^2 - \mu_j^2 - \sigma_j^2 \quad (5)$$

where J is the dimensionality of the latent space, and μ_j and σ_j denote the mean and standard deviation of the j-th latent variable, respectively. This expression assumes a factorized Gaussian posterior $q_\theta(\boldsymbol{z}|\boldsymbol{X},\boldsymbol{c}) = \mathcal{N}(\mu, \Sigma)$, where Σ is diagonal and σ_j^2 corresponds to its j-th diagonal element. The KL divergence acts as a regulariser, encouraging the learned latent representations to remain close to the prior $\mathcal{N}(0,\boldsymbol{I})$, which improves generalization and enables sampling during inference.

To balance the loss terms in Eq. 3, we equally prioritized the point cloud reconstruction and autoencoding (prediction) losses by setting their weights to 1 ($w_r = w_p = 1$). However, for $\mathscr{L}_{latent}$, we applied KL cost annealing proposed by Bowman and Vilnis [59], initialising w_l at zero and increasing it linearly to 1 over the first 10 epochs. This approach allows the model to initially focus on learning meaningful reconstructions and predictions before gradually applying the regularizing effect of the KL term, improving training stability and helping to prevent posterior collapse [60].

We trained the cVAE within 100 epochs, with each batch including 36 individuals. During training, the total loss was iteratively minimized as the model was provided with pairs of inputs and targets in each epoch. The input–target pairs, along with their corresponding condition vectors, could either span baseline and follow-up time points, requiring the model to learn to make accurate predictions, or represent the same time point, in which case the model is trained to replicate the input (i.e., autoencode it) to encourage a more consistent latent space. We recorded the model weights after each epoch and ultimately picked the weights that resulted in the minimum validation loss.

3. Results

3.1. The capability of shape descriptors to associate with disease progression

We initially explored whether a PointNet could effectively learn to correlate the anatomical shape of a subcortical structure with HD progression, as represented by a wide range of PIN scores in the training set. Fig. 4 illustrates the estimated versus the actual values of PIN scores in an unseen test set for each trained PointNet per structure, with HD-ISS stages color-coded. Supplementary Figure S1 shows the same plot for the training set, with mean absolute error (MAE) reported for different stages.

For each structure, training and validation losses were tracked across epochs. Training loss steadily decreased, while validation loss plateaued, indicating overfitting. The final model was selected at the onset of the validation plateau, balancing low training and validation losses (See Supplementary Figure S2). We observed that most structures showed strong shape–PIN associations, except for the hippocampus and amygdala, which demonstrated weak associations (See Supplementary Figure S3).

Training a PointNet, consisting of F shape descriptors, essentially involves locating a selection of F critical points from the N input points and also an optimal relationship among these critical points in the high-dimensional representation, such that a skeleton of the shape could be defined with respect to the target task, in this case, regression of the PIN score. Fig. 5 displays the critical points, highlighted in red, for an individual with five consecutive MRIs at ages 33 to 39. Since preprocessing involves residing all points within unit spheres, the trained model becomes relatively agnostic to the variations in volumes (mm^3) per se and instead concentrates on the associations among the shape descriptors.

3.2. Shape descriptors effectively captured the disease continuum

We used one-dimensional PCA for dimensionality reduction and to quantitatively compare the effectiveness of the stacked shape descriptors and volumetric measures in reflecting disease progression in the unseen test set. The Spearman's rank correlation between HD-ISS stage and the principal component was higher for shape descriptors ($\rho = 0.72$ vs $\rho = 0.45$) indicating better capture of disease progression in this feature space. Pairwise stage comparisons are given in the supplementary materials.

Fig. 6 presents two-dimensional visualizations of both feature spaces using PCA and t-SNE. These projections provide qualitative insights into the structure of the data and highlight the improved separability of disease stages within the shape descriptor space. We also performed an

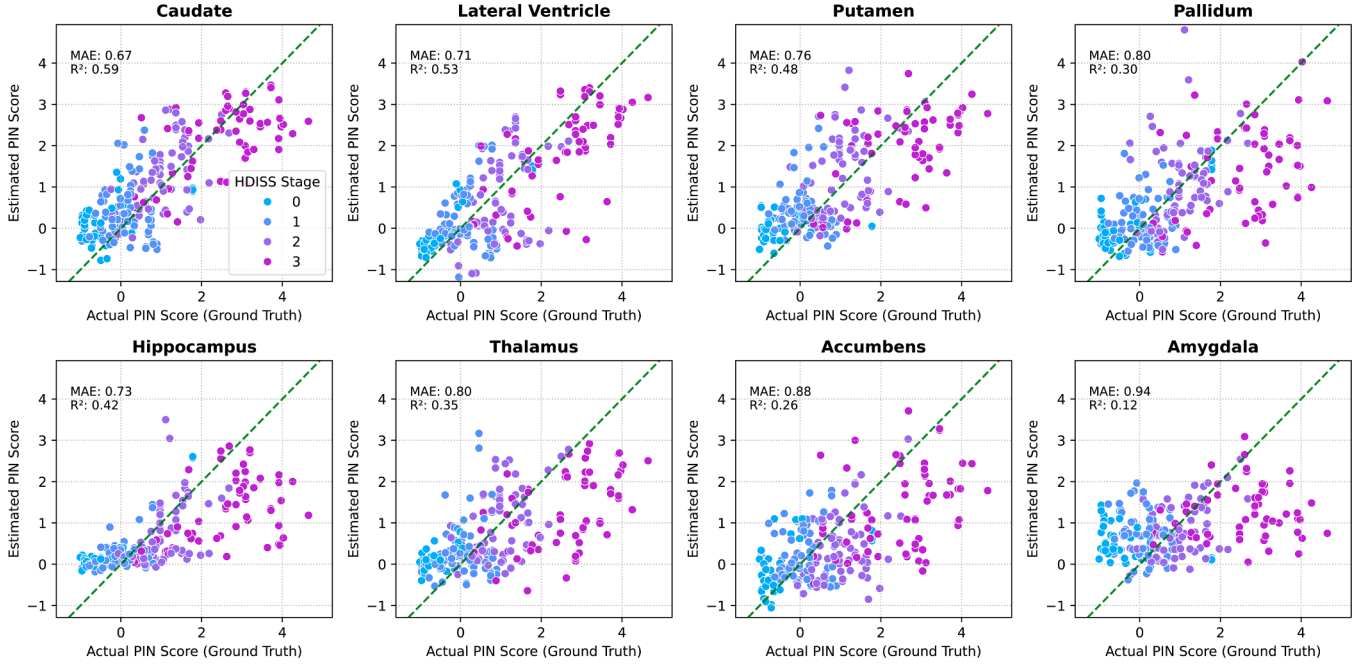

Fig. 4. A comparison between the capability of shape descriptors derived from different subcortical structures in estimating disease progression measured by the concurrent PIN score, in an unseen test dataset. HD-ISS stages are colour coded.

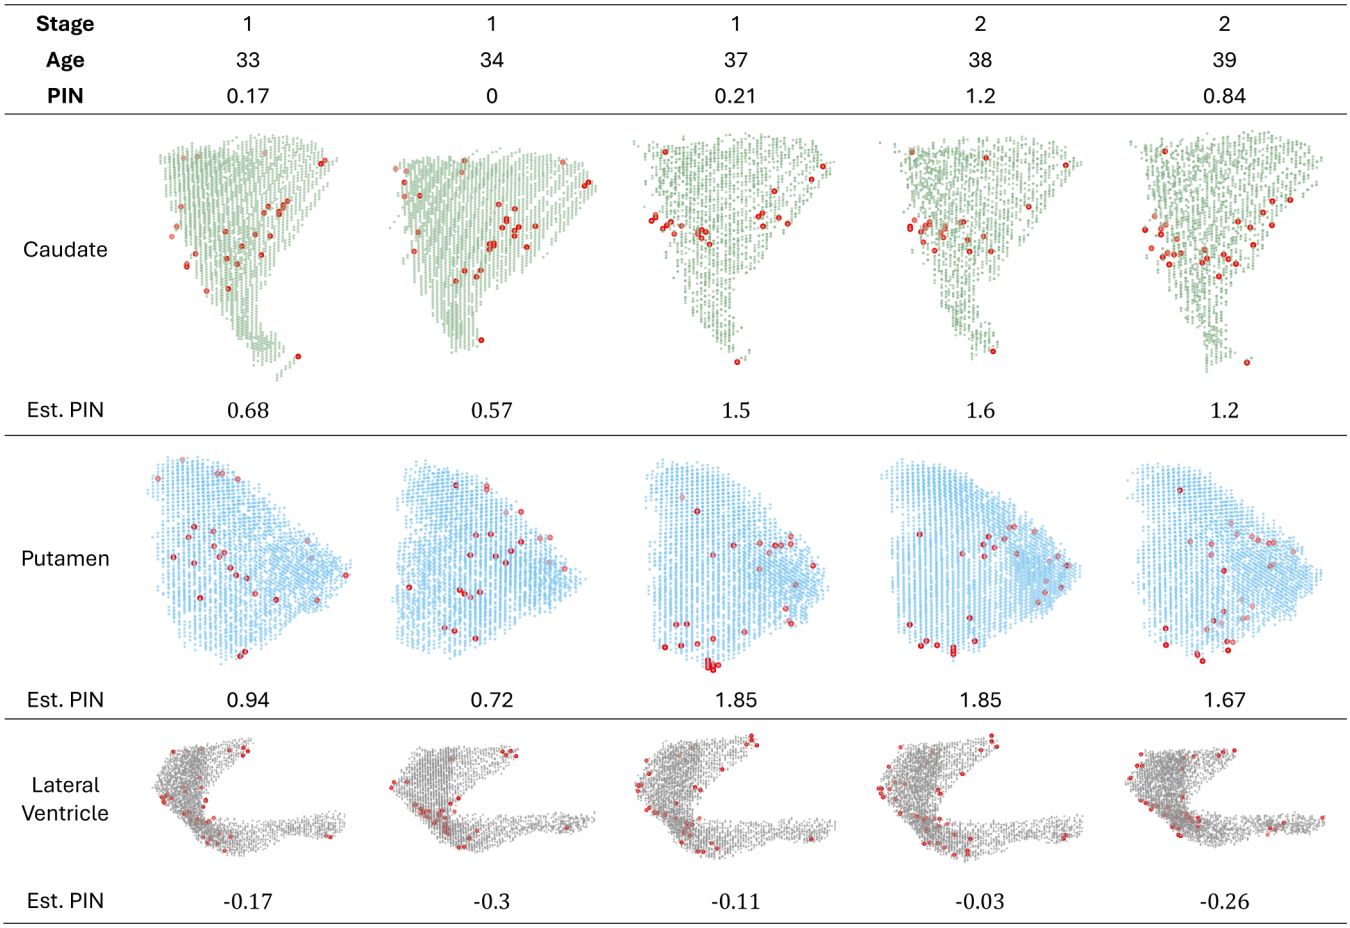

Fig. 5. Critical points for estimating PIN scores from subcortical shapes in PointNets: Each segmented structure corresponds to MRI scans acquired from an HD individual aging from 33 to 39 years. The red points, called critical points in the context of a PointNet model, are points that give rise to the shape descriptors in the max pooling layer.

additional classification experiment in which the two feature spaces were used to assign HD-ISS stages to inputs (see Supplementary Figures S4 and S5).

3.3. Subject-specific prediction of biomarkers of disease progression

We assessed the performance of the trained cVAE in predicting follow-up cUHDRS and volumes of the caudate, putamen, and lateral ventricle. Using an unseen test dataset that included individuals from all stages, $\boldsymbol{X}$ and $\boldsymbol{c}$ at baseline time point were provided to the encoder to obtain the latent code, $\boldsymbol{z}$. Subsequently, this latent code was concatenated with the follow-up conditions $\boldsymbol{c}$ (namely, sex as well as age and the CAP score in the follow-up visit) and was fed into the decoder for output generation under the follow-up conditions.

Improvements in MAE, expressed as percentage reductions relative to the naïve model, are shown in Table 2 for both the shape-lacking and shape-informed cVAE models. For reference, the naïve model's MAEs are given in the first column. Supplementary Figures S5 and S6 show individual trajectories from baseline to follow-up, comparing ground-truth values with predictions in the models.

4. Discussion

For the first time, we aimed to harness the information embedded in the shape of subcortical structures to improve the accuracy of predicting temporal changes in HD biomarkers (both the cUHDRS and volumetric), which are commonly used as primary or secondary outcome measures in clinical trials. We demonstrated the feasibility of training a discriminative DL architecture (PointNet [26] to extract meaningful descriptors from brain subcortical structures that associate shape with HD progression, as measured by the PIN score.

Our findings confirmed that shape descriptors derived from the putamen, caudate, pallidum, and the lateral ventricle exhibited strong predictive associations with HD progression, while thalamus and accumbens had weaker associations. Moreover, consistent with previous research [19], we found that shape features derived from amygdala and hippocampus were less indicative of disease progression, compared to the striatum.

Shape descriptors represented the trajectory of disease progression from HD-ISS stage 0 to 3 more granularly, when compared to volumetric measures (Fig. 6). This comparative advantage over simple volumetrics represents a key strength of the proposed framework, suggesting that shape-derived features may offer a complementary and precise tool for prognosis and for tracking progression across the HD continuum. This improvement could be due to the fact that aggregate metrics like volume fail to fully capture the intricate details of subcortical atrophy [61]. For example, same-day MRI acquisitions available in PREDICT or TRACK [62] exhibited volume inconsistencies, which were attributed to unavoidable scanning discrepancies or the limited robustness of segmentation algorithms [63]. In contrast, the higher dimensionality of shape descriptors potentially offers greater robustness to segmentation errors and allows for the detection of localized structural variations [64].

PointNet, compared to harmonics-based shape analysis methods [65–67], does not rely on predefined templates; instead, it exclusively seeks underlying morphometric interactions that are tailored to the specific predictive task [27], in our case PIN association. In fact, our

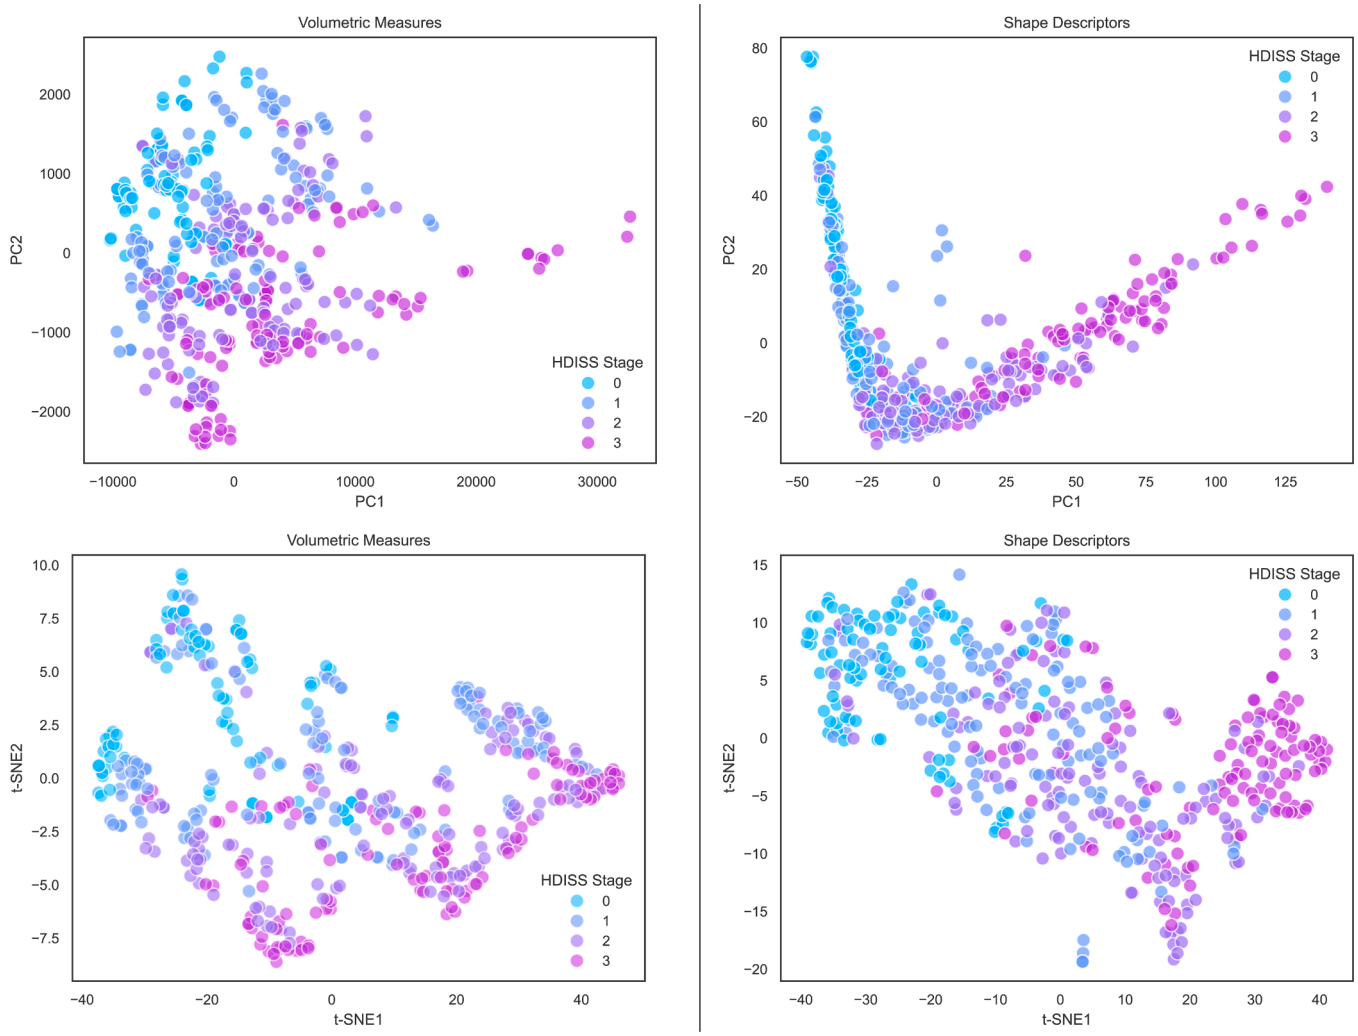

Fig. 6. Two-dimensional visualization of feature spaces using t-SNE and PCA for volumetric measures and shape descriptors from the putamen, pallidum, caudate, lateral ventricle, thalamus, and accumbens. t-SNE values are arbitrary and show relative relationships in the space of shape descriptors reflecting disease progression across HD-ISS stages.

Table 2
Mean absolute error (MAE) for the Naïve model and MAE reduction in percentage for the shape-informed cVAE models. Higher percentages indicate greater improvement. LV: lateral ventricle.

	Naïve	Shape-lacking	Shape-informed (L.V. + Caudate)	Shape-informed (L.V. + Caudate + Putamen)
cUHDRS	0.85	0.85 (0.0%)	0.8 (2.4%)	0.83 (2.4%)
Caudate vol.	97.7	81.3 (16.9%)	78.5 (19.7%)	81.00 (17.2%)
Putamen vol.	172.6	153.7 (10.9%)	—	158.5 (8.8%)
LV vol.	693.7	657.3 (5.3%)	571.5 (17.6%)	551.2 (20.4%)

qualitative analysis of the critical points giving rise to PIN estimations in the corresponding trained PointNets highlighted consistent patterns of decisive points representing shapes (Fig. 5).

Building on the findings from the discriminative PointNets, linking shapes to concurrent PIN scores, in this study we developed for the first time a conditional generative model capable of predicting the cUHDRS and volumes of caudate, putamen, and lateral ventricle at follow-up using their baseline values. We incorporated the prior clinical knowledge that age and its cumulative effect on disease burden should be incorporated into the distribution of the latent space governing the conditional model. Compared to a basic cVAE, that solely relied on volumetric measures for making predictions, the shape-informed counterpart performed more accurate, highlighting the complementary predictive value of subcortical morphometric deformations for HD prognosis. One other important aspect is that the prediction accuracy for the cUHDRS, conventionally used as a primary outcome measure in clinical trials, was lower compared to that of caudate and putamen volumes. This further emphasizes the critical role of incorporating volumetric measures as secondary endpoints, specifically in early stage preventive drug development trials [68]. In clinical trials, due to practicality constraints, it is often not feasible or cost effective to collect longitudinal data for patient stratification. Therefore, we adhered to an architecture that used single-time-point data at baseline to predict the decline in biomarkers over time. However, in scenarios where longitudinal data is available, approaches such as recurrent neural networks are recommended [69].

Recently, there has been growing interest in the shapes of the cortex and subcortex in relation to normal neural function and heritability. Pang et al. found that brain geometry imposes fundamental constraints on brain dynamics, with brain activity being better explained by geometry than by complex interregional connectivity, emphasizing the role of geometry in brain dynamics [70]. Roshchupkin et al. found that the shapes of subcortical brain structures are highly heritable, with

genetic factors contributing significantly to localized shape variability in subcortical structures [61]. The present study brings attention to the value of subcortical shape subtleties in HD beyond volumetric features. While our focus in this study was on predicting conventional clinical and volumetric biomarkers in HD, the developed model could potentially be employed to generate point sets corresponding to follow-up visits. Future HD studies could leverage this model to analyse longitudinal shape differences and investigate patterns of diffeo-morphometry in structural subdivisions, providing insights into disease progression. Beyond its methodological advantages, understanding shape of a structure carries clinical relevance. Clinical trials may benefit from shape analysis as a monitoring strategy, offering the ability to anatomically localize where neuropathological changes are expected and to assess whether an investigational therapeutic has been able to halt progression of these changes [20]. Supplementary Figure S9 illustrates this by showing longitudinal comparisons between predicted and actual follow-up point clouds, highlighting the potential to track subject-specific structural deformation. To be practically implemented in trial settings, however, such applications would require a dedicated diffeomorphic framework in which longitudinal anatomical change is explicitly incorporated into model training [71]. On the other hand, as suggested by the evidence in Fig. 6, trial enrichment strategies may also be improved by incorporating shape information at baseline, thereby capturing within-stage heterogeneity and enabling more precise stratification prior to enrolment.

One limitation of representing segmented volumes with point clouds, compared to established volumetric masks, is the interpretability of results. Sarasua et al. have comprehensively compared alternative representations of volumetric shapes for Alzheimer's disease prognosis [72]. For example, volumetric textures, which include hypo/hyper-intensities, may contain information related to neurodegeneration [73] which is lost in surface representations such as point clouds. Furthermore, the approach used in the present work may cause the loss of fine-grained local geometric details due to surface sampling along with the inherent focus of PointNets on global shape features. Architectures such as PointNet++ [74] could be explored to address this issue. Yet, a key limitation of DL-based shape models such as PointNet is their restricted interpretability relative to classical statistical shape modelling or harmonics-based approaches, where spatially explicit patterns of shape change can be directly visualized and attributed to disease progression. The reliance on FreeSurfer for structural segmentation introduces variability inherent in its parcellation algorithms [75]. Additionally, the generative model is computationally demanding, which prevented the inclusion of data from non-huntingtin gene expansion carriers, which would have allowed the model to learn to disentangle atrophy due to normal aging from neuropathology-driven deformations. Simpler approaches may involve developing novel composite disease prognostication scores [76] based on the underlying trajectory of the shape-informed representations, as seen in Fig. 6. Moreover, the current pipeline does not incorporate cortical features or wet biomarkers such as neurofilament light chain (NfL), which may provide additional predictive value. While the results demonstrate a robust association between shape features and concurrent disease severity (descriptive validity), the added value of shape descriptors in forecasting future biomarker decline (predictive validity) remains modest. The present findings should therefore be interpreted as preliminary evidence of feasibility rather than conclusive proof of clinical utility. Another important limitation of the present study is its exclusive focus on subcortical structures. Cortical thinning [77], gyrification changes [78], and network disconnection [79] are well documented in HD, including in earlier stages, and have been shown to contribute meaningfully to disease progression. By restricting our framework to imaging-derived shape descriptors of subcortical structures, we may overemphasize the role of subcortical morphology alone. Future work should therefore integrate cortical features and measures of structural and functional connectivity to provide a more comprehensive and biologically grounded model of HD progression.

In summary, we proposed a novel end-to-end pipeline for integrating shape information into a predictive model of disease progression in HD. By incorporating shape descriptors, the model demonstrated a modest but consistent improvement in forecasting HD biomarkers at follow-up, particularly for caudate and putamen volumes. Such advanced methods offer a promising approach for refining enrolling homogeneous HD populations in future treatment trials and serve as valuable post-hoc tools for assessing treatment effects in a subject-specific manner.

Ethics statement

IMAGE-HD was approved by the Monash University and Melbourne Health Human Research Ethics Committees as a single-site study in Melbourne, Australia. TRACK-HD was approved by the local ethics committees at each study site in the Netherlands, UK, France, and Canada. PREDICT-HD procedures were approved by institutional review boards at each site (32 sites across the United States, Canada, Australia, and Europe). For all 3 studies, each participant provided written informed consent.

Funding

TRACK-HD and TrackON-HD were supported by the CHDI Foundation, a not-for-profit organisation dedicated to finding treatments for Huntington's disease. IMAGE-HD was supported by CHDI Foundation research agreement A-3433 and the National Health and Medical Research Council (NHMRC) Australia grant 606650 (N.G.-K.). The PREDICT-HD study was funded by the NCATS and the NIH (NIH; R01—NS040068, U01—NS105509, U01—NS103475).

Data availability

Data from the PREDICT-HD, TRACK-HD, and IMAGE-HD studies can be accessed through CHDI foundation upon request (enroll-hd.org). The written scripts used for model development and evaluation in this study are available at: https://github.com/mghofrani/HD_Shape_Descriptors.

CRediT authorship contribution statement

Mohsen Ghofrani-Jahromi: Writing – review & editing, Writing – original draft, Visualization, Methodology, Formal analysis, Conceptualization. **Susmita Saha:** Writing – review & editing, Validation, Methodology, Formal analysis, Conceptualization. **Adeel Razi:** Writing – review & editing, Supervision, Methodology, Conceptualization. **Pubu M. Abeyasinghe:** Writing – review & editing, Validation, Data curation. **Govinda R. Poudel:** Writing – review & editing, Methodology, Conceptualization. **Jane S. Paulsen:** Writing – review & editing, Validation, Funding acquisition, Data curation. **Sarah J. Tabrizi:** Writing – review & editing, Validation, Funding acquisition, Data curation, Conceptualization. **Nellie Georgiou-Karistianis:** Writing – review & editing, Validation, Supervision, Resources, Project administration, Investigation, Funding acquisition, Data curation, Conceptualization.

Declaration of competing interest

The authors declare no competing interests.

Acknowledgements

IMAGE-HD data used in this work was generously provided by the participants in the IMAGE-HD study and made available by Professor Nellie Georgiou-Karistianis, Principal Investigator. TRACK-HD and TrackON-HD data used in this work were generously provided by the participants in these studies and made available by Dr Sarah Tabrizi, Principal Investigator. PREDICT-HD data used in this work was

generously provided by the participants in PREDICT-HD study and made available by the PREDICT-HD investigators and coordinators of the Huntington study group, Dr Jane S. Paulsen, Principal Investigator.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.cmpb.2026.109250.

References

[1] A. Videnovic, HC V Pfeiffer, A. Tylki-Szymańska, E. Berry-Kravis, F. Ezgü, J. Ganju, et al., Study design challenges and strategies in clinical trials for rare diseases: lessons learned from pantothenate kinase-associated neurodegeneration, Front. Neurol. 14 (2023) 1098454, https://doi.org/10.3389/fneur.2023.1098454.

[2] C. Sampaio, HA. Wilkinson, Facilitating Huntington's disease research: plasma neurofilament levels as a promising enrichment biomarker for HD-ISS stage 1, EBioMedicine 94 (2023) 104710, https://doi.org/10.1016/j.ebiom.2023.104710.

[3] G.M. Parkin, J. Corey-Bloom, in: E.A. Thomas, G.M. Parkin (Eds.), Considerations and Advances in Huntington's Disease Clinical Trial Design, Springer International Publishing, Cham, 2023, pp. 405–431, https://doi.org/10.1007/978-3-031-32815-2_17, editors.

[4] M. Bordukova, N. Makarov, R. Rodriguez-Esteban, F. Schmich, MP. Menden, Generative artificial intelligence empowers digital twins in drug discovery and clinical trials, Expert. Opin. Drug Discov. 19 (2024) 33–42, https://doi.org/10.1080/17460441.2023.2273839.

[5] B. Zhang, L. Zhang, Q. Chen, Z. Jin, S. Liu, S. Zhang, Harnessing artificial intelligence to improve clinical trial design, Commun. Med. 3 (2023) 8–10, https://doi.org/10.1038/s43856-023-00425-3.

[6] Z. Sun, S. Ghosh, Y. Li, Y. Cheng, A. Mohan, C. Sampaio, et al., A probabilistic disease progression modeling approach and its application to integrated Huntington's disease observational data, JAMIA Open 2 (2019) 123–130, https://doi.org/10.1093/jamiaopen/ooy060.

[7] Wijeratne P. The temporal event-based model: learning event timelines in progressive diseases 2023;1:1–19. https://doi.org/10.1162/imag.

[8] E. Castro, P. Polosecki, D. Pustina, A. Wood, C. Sampaio, GA. Cecchi, Predictive modeling of huntington's disease unfolds thalamic and caudate atrophy dissociation, Mov. Disord. 37 (2022) 2407–2416, https://doi.org/10.1002/mds.29219.

[9] P.A. Wijeratne, E.B. Johnson, S. Gregory, N. Georgiou-Karistianis, J.S. Paulsen, R. I. Scahill, et al., A Multi-study model-based evaluation of the sequence of imaging and clinical biomarker changes in Huntington's disease, Front. Big. Data 4 (2021) 1–8, https://doi.org/10.3389/fdata.2021.662200.

[10] Wijeratne PA, Young AL, Oxtoby NP, Marinescu RV, Firth NC, Johnson EB, et al. An image-based model of brain volume biomarker changes in Huntington's disease 2018. https://doi.org/10.1002/acn3.558.

[11] P.A. Wijeratne, S. Garbarino, S. Gregory, E.B. Johnson, R.I. Scahill, J.S. Paulsen, et al., Revealing the timeline of structural MRI changes in premanifest to manifest Huntington disease, Neurol. Genet. 7 (2021) e617, https://doi.org/10.1212/NXG.0000000000000617.

[12] P.M. Abeyasinghe, J.D. Long, A. Razi, D. Pustina, J.S. Paulsen, S.J. Tabrizi, et al., Tracking Huntington's disease progression using motor, functional, cognitive, and imaging markers, Mov. Disord. 36 (2021) 2282–2292, https://doi.org/10.1002/mds.28650.

[13] J. Warner, C. Sampaio, Modeling variability in the progression of Huntington's disease a novel modeling approach applied to structural imaging markers from TRACK-HD, CPT Pharmacomet. Syst. Pharmacol. 5 (2016) 437–445, https://doi.org/10.1002/psp4.12097.

[14] C-F Liu, L. Younes, X.J. Tong, J.T. Hinkle, M. Wang, S. Phatak, et al., Longitudinal imaging highlights preferential basal ganglia circuit atrophy in Huntington's disease, Brain Commun. 5 (2023) 1–16, https://doi.org/10.1093/braincomms/fcad214.

[15] I. Koval, T. Dighiero-Brecht, A.J. Tobin, S.J. Tabrizi, R.I. Scahill, S. Tezenas du Montcel, et al., Forecasting individual progression trajectories in Huntington disease enables more powered clinical trials, Sci. Rep. 12 (2022) 18928, https://doi.org/10.1038/s41598-022-18848-8.

[16] A. Mohan, Z. Sun, S. Ghosh, Y. Li, S. Sathe, J. Hu, et al., A Machine-learning derived huntington's disease progression model: insights for clinical trial design, Mov. Disord. 37 (2022) 553–562, https://doi.org/10.1002/mds.28866.

[17] P.A. Wijeratne, A.L. Young, N.P. Oxtoby, RV. Marinescu, N.C. Firth, E.B. Johnson, et al., An image-based model of brain volume biomarker changes in Huntington's disease, Ann. Clin. Transl. Neurol. 5 (2018) 570–582, https://doi.org/10.1002/acn3.558.

[18] S.J.A. van den Bogaard, E.M. Dumas, L. Ferrarini, J. Milles, M.A. van Buchem, J. van der Grond, et al., Shape analysis of subcortical nuclei in Huntington's disease, global versus local atrophy — Results from the TRACK-HD study, J. Neurol. Sci. 307 (2011) 60–68, https://doi.org/10.1016/j.jns.2011.05.015.

[19] L. Younes, J.T. Ratnanather, T. Brown, E. Aylward, P. Nopoulos, H. Johnson, et al., Regionally selective atrophy of subcortical structures in prodromal HD as revealed by statistical shape analysis, Hum. Brain Mapp. 35 (2014) 792–809, https://doi.org/10.1002/hbm.22214.

[20] X. Tang, C.A. Ross, H. Johnson, J.S. Paulsen, L. Younes, R.L. Albin, et al., Regional subcortical shape analysis in premanifest Huntington's disease, Hum. Brain Mapp. 40 (2019) 1419–1433, https://doi.org/10.1002/hbm.24456.

[21] F.A. Wilkes, Z. Abaryan, C.R.K. Ching, B.A. Gutman, S.K. Madsen, M. Walterfang, et al., Striatal morphology and neurocognitive dysfunction in Huntington disease: The IMAGE-HD study, Psychiatry Res. Neuroimaging 291 (2019) 1–8, https://doi.org/10.1016/j.pscychresns.2019.07.003.

[22] F.A. Wilkes, D. Jakabek, M. Walterfang, D. Velakoulis, G.R. Poudel, J.C. Stout, et al., Hippocampal morphology in Huntington's disease, implications for plasticity and pathogenesis: The IMAGE-HD study, Psychiatry Res. Neuroimaging 335 (2023), https://doi.org/10.1016/j.pscychresns.2023.111694.

[23] F.A. Wilkes, D. Jakabek, M. Walterfang, D. Velakoulis, G.R. Poudel, J.C. Stout, et al., The shape of things to come. Mapping spatiotemporal progression of striatal morphology in Huntington disease: The IMAGE-HD study, Psychiatry Res. Neuroimaging (2023) 111717, https://doi.org/10.1016/j.pscychresns.2023.111717.

[24] L.S. Furlong, D. Jakabek, B.D. Power, C. Owens-Walton, F.A. Wilkes, M. Walterfang, et al., Morphometric in vivo evidence of thalamic atrophy correlated with cognitive and motor dysfunction in Huntington's disease: The IMAGE-HD study, Psychiatry Res. Neuroimaging 298 (2020) 111048, https://doi.org/10.1016/j.pscychresns.2020.111048.

[25] A. Kopf, M. Claassen, Latent representation learning in biology and translational medicine, Patterns 2 (2021) 100198, https://doi.org/10.1016/j.patter.2021.100198.

[26] C.R. Qi, H. Su, K. Mo, LJ. Guibas, PointNet: deep learning on point sets for 3D classification and segmentation, in: 2017 IEEE Conf Comput Vis Pattern Recognit, 2016, pp. 77–85, https://doi.org/10.1109/CVPR.2017.16, 2017-Janua.

[27] B. Gutiérrez-Becker, I. Sarasua, C. Wachinger, Discriminative and generative models for anatomical shape analysis on point clouds with deep neural networks, Med. Image Anal. 67 (2021) 101852, https://doi.org/10.1016/j.media.2020.101852.

[28] S. Pölsterl, I. Sarasua, B. Gutiérrez-Becker, C. Wachinger, A Wide and deep neural network for survival analysis from anatomical shape and tabular clinical data, Commun. Comput. Inf. Sci. 1167 (2020) 453–464, https://doi.org/10.1007/978-3-030-43823-4_37. CCIS.

[29] W. Yang, X. Bai, X. Guan, C. Zhou, T. Guo, J. Wu, et al., The longitudinal volumetric and shape changes of subcortical nuclei in Parkinson's disease, Sci. Rep. 14 (2024) 7494, https://doi.org/10.1038/s41598-024-58187-4.

[30] M.A. Laansma, Y. Zhao, E.M. van Heese, J.K. Bright, C. Owens-Walton, S. Al-Bachari, et al., A worldwide study of subcortical shape as a marker for clinical staging in Parkinson's disease, Npj Park Dis. 10 (2024) 223, https://doi.org/10.1038/s41531-024-00825-9.

[31] D. Ravi, S.B. Blumberg, S. Ingala, F. Barkhof, D.C. Alexander, NP. Oxtoby, Degenerative adversarial neuroimage nets for brain scan simulations: Application in ageing and dementia, Med. Image Anal. 75 (2022) 102257, https://doi.org/10.1016/j.media.2021.102257.

[32] S. Dayarathna, K.T. Islam, S. Uribe, G. Yang, M. Hayat, Z. Chen, Deep learning based synthesis of MRI, CT and PET: review and analysis, Med. Image Anal. 92 (2024) 103046, https://doi.org/10.1016/j.media.2023.103046.

[33] E. Jung, M. Luna, SH. Park, Conditional GAN with 3D discriminator for MRI generation of Alzheimer's disease progression, Pattern. Recognit. 133 (2023) 109061, https://doi.org/10.1016/j.patcog.2022.109061.

[34] C. Gong, C. Jing, X. Chen, C.M. Pun, G. Huang, A. Saha, et al., Generative AI for brain image computing and brain network computing: a review, Front. Neurosci. 17 (2023), https://doi.org/10.3389/fnins.2023.1203104.

[35] A. Lawry Aguila, J. Chapman, M. Janahi, A. Altmann, Conditional VAEs for Confound Removal and Normative Modelling of Neurodegenerative Diseases, in: Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics), 13431, Springer International Publishing, 2022, pp. 430–440, https://doi.org/10.1007/978-3-031-16431-6_41. LNCS.

[36] H. Choi, H. Kang, DS. Lee, Predicting aging of brain metabolic topography using variational autoencoder, Front. Aging Neurosci. 10 (2018), https://doi.org/10.3389/fnagi.2018.00212.

[37] Y. Chai, M. Liu, B.A. Duffy, H. Kim, Learning to Synthesize Cortical Morphological Changes using Graph Conditional Variational Autoencoder, in: 2021 IEEE 18th Int. Symp. Biomed. Imaging, IEEE, 2021, pp. 1495–1499, https://doi.org/10.1109/ISBI48211.2021.9433837, vol. 2021- April.

[38] S.J. Tabrizi, R.I. Scahill, G. Owen, A. Durr, B.R. Leavitt, R.A. Roos, et al., Predictors of phenotypic progression and disease onset in premanifest and early-stage Huntington's disease in the TRACK-HD study: analysis of 36-month observational data, Lancet Neurol. 12 (2013) 637–649, https://doi.org/10.1016/S1474-4422(13)70088-7.

[39] S. Klöppel, S. Gregory, E. Scheller, L. Minkova, A. Razi, A. Durr, et al., Compensation in preclinical Huntington's disease: evidence from the track-on HD study, EBioMedicine 2 (2015) 1420–1429, https://doi.org/10.1016/j.ebiom.2015.08.002.

[40] S V Lobanov, B. McAllister, M. McDade-Kumar, G.B. Landwehrmeyer, M. Orth, A. E. Rosser, et al., Huntington's disease age at motor onset is modified by the tandem hexamer repeat in TCERG1, Npj Genom. Med. 7 (2022) 53, https://doi.org/10.1038/s41525-022-00317-w.

[41] J.S. Paulsen, J.D. Long, H.J. Johnson, E.H. Aylward, C.A. Ross, J.K. Williams, et al., Clinical and biomarker changes in premanifest Huntington disease show trial feasibility: A decade of the PREDICT-HD study, Front. Aging Neurosci. 6 (2014) 1–11, https://doi.org/10.3389/fnagi.2014.00078.

[42] N. Georgiou-Karistianis, J.C. Stout, D.JF Domínguez, S.P. Carron, A. Ando, A. Churchyard, et al., Functional magnetic resonance imaging of working memory

in Huntington's disease: Cross-sectional data from the IMAGE-HD study, Hum. Brain Mapp. 35 (2014) 1847–1864, https://doi.org/10.1002/hbm.22296.
[43] G.R. Poudel, J.C. Stout, D JF Domínguez, A. Churchyard, P. Chua, G.F. Egan, et al., Longitudinal change in white matter microstructure in Huntington's disease: the IMAGE-HD study, Neurobiol. Dis. 74 (2015) 406–412, https://doi.org/10.1016/j.nbd.2014.12.009.
[44] M. Reuter, N.J. Schmansky, H.D. Rosas, B. Fischl, Within-subject template estimation for unbiased longitudinal image analysis, Neuroimage 61 (2012) 1402–1418, https://doi.org/10.1016/j.neuroimage.2012.02.084.
[45] J.D. Long, D.R. Langbehn, S.J. Tabrizi, B.G. Landwehrmeyer, J.S. Paulsen, J. Warner, et al., Validation of a prognostic index for Huntington's disease, Mov. Disord. 32 (2017) 256–263, https://doi.org/10.1002/mds.26838.
[46] C. Estevez-Fraga, R.I. Scahill, A. Durr, B.R. Leavitt, R.A.C. Roos, D.R. Langbehn, et al., Composite UHDRS correlates with progression of imaging biomarkers in Huntington's Disease, Mov. Disord. 36 (2021) 1259–1264, https://doi.org/10.1002/mds.28489.
[47] M. Farag, S.J. Tabrizi, EJ. Wild, Huntington's disease clinical trials update: March 2025, J. Huntingt. Dis. 14 (2025) 191–206, https://doi.org/10.1177/18796397251337000.
[48] K.M. Kinnunen, A.J. Schwarz, E.C. Turner, D. Pustina, E.C. Gantman, M.F. Gordon, et al., Volumetric MRI-Based Biomarkers in Huntington's Disease: An Evidentiary Review, Front. Neurol. 12 (2021), https://doi.org/10.3389/fneur.2021.712555.
[49] S.M. Nestor, R. Rupsingh, M. Borrie, M. Smith, V. Accomazzi, J.L. Wells, et al., Ventricular enlargement as a possible measure of Alzheimer's disease progression validated using the Alzheimer's disease neuroimaging initiative database, Brain 131 (2008) 2443–2454, https://doi.org/10.1093/brain/awn146.
[50] J.H. Warner, J.D. Long, J.A. Mills, D.R. Langbehn, J. Ware, A. Mohan, et al., Standardizing the CAP Score in Huntington's Disease by Predicting Age-at-Onset, J. Huntingt. Dis. 11 (2022) 153–171, https://doi.org/10.3233/JHD-210475.
[51] S.J. Tabrizi, S. Schobel, E.C. Gantman, A. Mansbach, B. Borowsky, P. Konstantinova, et al., A biological classification of Huntington's disease: the Integrated Staging System, Lancet Neurol. 21 (2022) 632–644, https://doi.org/10.1016/S1474-4422(22)00120-X.
[52] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, et al., PyTorch: an imperative style, high-performance deep learning library, Adv. Neural Inf. Process. Syst. 32 (2019), https://doi.org/10.48550/arXiv.1912.01703.
[53] L. Minkova, S. Gregory, R.I. Scahill, A. Abdulkadir, C.P. Kaller, J. Peter, et al., Cross-sectional and longitudinal voxel-based grey matter asymmetries in Huntington's disease, Neuroimage Clin. 17 (2018) 312–324, https://doi.org/10.1016/j.nicl.2017.10.023.
[54] Wijeratne PA, Johnson EB, Eshaghi A, Aksman L, Gregory S, Johnson HJ, et al. Robust Markers and Sample Sizes for Multicenter Trials of Huntington Disease 2020:751–62. https://doi.org/10.1002/ana.25709.
[55] K. Sohn, X. Yan, H. Lee, Learning structured output representation using deep conditional generative models, in: Proc. 29th Int. Conf. Neural Inf. Process. Syst 2, MIT Press, Cambridge, MA, USA, 2015, pp. 3483–3491.
[56] R. He, G. Ang, D. Tward, A. Schroder, X. Li, T. Syeda-Mahmood, N.P. Oxtoby, A. Young, A. Hering, et al., Individualized Multi-horizon MRI Trajectory Prediction for Alzheimer's Disease, editors. Int. Conf. Med. Image Comput. Comput. Interv, Springer Nature Switzerland, Cham, 2025, pp. 26–37, https://doi.org/10.1007/978-3-031-84525-3_3.
[57] D.P. Kingma, M. Welling, Auto-encoding variational bayes. 2nd Int, in: Conf. Learn. Represent. ICLR 2014 - Conf. Track Proc, 2014, pp. 1–14, https://doi.org/10.61603/ceas.v2i1.33.
[58] W. Yuan, T. Khot, D. Held, C. Mertz, M. Hebert, PCN: Point Completion Network, in: 2018 Int. Conf. 3D Vis, IEEE, 2018, pp. 728–737, https://doi.org/10.1109/3DV.2018.00088.
[59] S.R. Bowman, L. Vilnis, O. Vinyals, A.M. Dai, R. Jozefowicz, S. Bengio, Generating Sentences from a Continuous Space, in: CoNLL 2016 - 20th SIGNLL Conf Comput Nat Lang Learn Proc, 2015, pp. 10–21, https://doi.org/10.18653/v1/k16-1002.
[60] T. Long, Y. Cao, JCK. Cheung, On Posterior Collapse and Encoder Feature Dispersion in Sequence VAEs, ArXiv Prepr ArXiv191103976 (2020), https://doi.org/10.48550/arXiv.1911.03976.
[61] GV. Roshchupkin, B.A. Gutman, M.W. Vernooij, N. Jahanshad, N.G. Martin, A. Hofman, et al., Heritability of the shape of subcortical brain structures in the general population, Nat. Commun. 7 (2016) 13738, https://doi.org/10.1038/ncomms13738.
[62] K.M. Kinnunen, A.P. Mullin, D. Pustina, E.C. Turner, J. Burton, M.F. Gordon, et al., Recommendations to Optimize the Use of Volumetric MRI in Huntington's, Dis. Clin. Trials 12 (2021), https://doi.org/10.3389/fneur.2021.712565.
[63] P. Muralidharan, J. Fishbaugh, H.J. Johnson, S. Durrleman, J.S. Paulsen, G. Gerig, et al., Diffeomorphic shape trajectories for improved longitudinal segmentation and statistics, Lect. Notes Comput. Sci. (Incl. Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinform.) 8675 (2014) 49–56, https://doi.org/10.1007/978-3-319-10443-0_7. LNCS.
[64] Z. Wang, DV. Rodriguez-Moreno, Y.M. Cycowicz, LV. Amsel, K. Cheslack-Postava, X. He, et al., Shapes of subcortical structures in adolescents with and without familial history of substance use disorder, Hum. Brain Mapp. 43 (2022) 2759–2770, https://doi.org/10.1002/hbm.25804.
[65] M. Shakeri, H. Lombaert, A.N. Datta, N. Oser, L. Létourneau-Guillon, L.V. Lapointe, et al., Statistical shape analysis of subcortical structures using spectral matching, Comput. Med. Imaging Graph. 52 (2016) 58–71, https://doi.org/10.1016/j.compmedimag.2016.03.001.
[66] B.A. Gutman, N. Jahanshad, C.R.K. Ching, Y. Wang, PV. Kochunov, T.E. Nichols, et al., Medial demons registration localizes the degree of genetic influence over subcortical shape variability: An N= 1480 meta-analysis, Proc. Int. Symp. Biomed Imaging (2015) 1402–1406, https://doi.org/10.1109/ISBI.2015.7164138, 2015-July.
[67] S.A. Primus, F. Hoffstaedter, F. Raimondo, S.B. Eickhoff, J. Winkelmann, K. Oexle, et al., Beyond volume: Unraveling the genetics of human brain geometry, Sci. Adv. 11 (2025) eadr1644, https://doi.org/10.1126/sciadv.adr1644.
[68] R. Scahill, M. Farag, M.J. Murphy, N.Z. Hobbs, M. Leocadi, K. Fayer, et al., J028 Preparing for secondary prevention treatment trials in HD-ISS Stages 0 and 1: Insights from the Huntington's disease young adult study (HD-YAS). J Clin. Ther, BMJ Publishing Group Ltd, 2024, https://doi.org/10.1136/jnnp-2024-EHDN.349.A171.2-A171.
[69] M. Nguyen, T. He, L. An, D.C. Alexander, J. Feng, B.T.T. Yeo, Predicting Alzheimer's disease progression using deep recurrent neural networks, Neuroimage 222 (2020) 7–10, https://doi.org/10.1016/j.neuroimage.2020.117203.
[70] J.C. Pang, K.M. Aquino, M. Oldehinkel, P.A. Robinson, B.D. Fulcher, M. Breakspear, et al., Geometric constraints on human brain function, Nature 618 (2023) 566–574, https://doi.org/10.1038/s41586-023-06098-1.
[71] X. Tang, D. Holland, A.M. Dale, L. Younes, MI. Miller, The diffeomorphometry of regional shape change rates and its relevance to cognitive deterioration in mild cognitive impairment and Alzheimer's disease, Hum. Brain Mapp. 36 (2015) 2093–2117, https://doi.org/10.1002/hbm.22758.
[72] I. Sarasua, S. Pölsterl, C. Wachinger, Hippocampal representations for deep learning on Alzheimer's disease, Sci. Rep. 12 (2022) 8619, https://doi.org/10.1038/s41598-022-12533-6.
[73] K. Wei, T. Tran, K. Chu, M.T. Borzage, M.N. Braskie, M.G. Harrington, et al., White matter hypointensities and hyperintensities have equivalent correlations with age and CSF β-amyloid in the nondemented elderly, Brain Behav. 9 (2019) e01457, https://doi.org/10.1002/brb3.1457.
[74] C.R. Qi, L. Yi, H. Su, LJ. Guibas, PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space, Adv. Neural Inf. Process. Syst. 30 (2017), https://doi.org/10.48550/arXiv.1706.02413.
[75] H. Knights, A. Coleman, N.Z. Hobbs, S.J. Tabrizi, RI. Scahill, Freesurfer Software Update Significantly Impacts Striatal Volumes in the Huntington's Disease Young Adult Study and Will Influence HD-ISS Staging, J Huntingt. Dis. 13 (2024) 77–90, https://doi.org/10.3233/JHD-231512.
[76] V. Kmetzsch, E. Becker, D. Saracino, D. Rinaldi, A. Camuzat, I. Le Ber, et al., Disease progression score estimation from multimodal imaging and MicroRNA data using supervised variational autoencoders, IEEE J. Biomed Heal Inform. 26 (2022) 6024–6035, https://doi.org/10.1109/JBHI.2022.3208517.
[77] B. Tan, R. Shishegar, A. Fornito, G. Poudel, N. Georgiou-Karistianis, Longitudinal mapping of cortical surface changes in Huntington's Disease, Brain Imaging Behav. 16 (2022) 1381–1391, https://doi.org/10.1007/s11682-021-00625-2.
[78] R. Shishegar, F. Pizzagalli, N. Georgiou-Karistianis, G.F. Egan, N. Jahanshad, LA. Johnston, A gyrification analysis approach based on Laplace Beltrami eigenfunction level sets, Neuroimage 229 (2021) 117751.
[79] B. Tan, R. Shishegar, G.R. Poudel, A. Fornito, N. Georgiou-Karistianis, Cortical morphometry and neural dysfunction in Huntington's disease: a review, Eur. J. Neurol. 28 (2021) 1406–1419, https://doi.org/10.1111/ene.14648.