

RESEARCH

Open Access



Leveraging explainable deep learning methodologies to elucidate the biological underpinnings of Huntington's disease using single-cell RNA sequencing data

Shichen Gao^{1,2†}, Yadong Wang^{1,2†}, Jiajia Wang² and Yan Dong^{2*}

Abstract

Background Huntington's disease (HD) is a hereditary neurological disorder caused by mutations in *HTT*, leading to neuronal degeneration. Traditionally, HD is associated with the misfolding and aggregation of mutant huntingtin due to an extended polyglutamine domain encoded by an expanded CAG tract. However, recent research has also highlighted the role of global transcriptional dysregulation in HD pathology. However, understanding the intricate relationship between mRNA expression and HD at the cellular level remains challenging. Our study aimed to elucidate the underlying mechanisms of HD pathology using single-cell sequencing data.

Results We used single-cell RNA sequencing analysis to determine differential gene expression patterns between healthy and HD cells. HD cells were effectively modeled using a residual neural network (ResNet), which outperformed traditional and convolutional neural networks. Despite the efficacy of our approach, the F1 score for the test set was 96.53%. Using the SHapley Additive exPlanations (SHAP) algorithm, we identified genes influencing HD prediction and revealed their roles in HD pathobiology, such as in the regulation of cellular iron metabolism and mitochondrial function. SHAP analysis also revealed low-abundance genes that were overlooked by traditional differential expression analysis, emphasizing its effectiveness in identifying biologically relevant genes for distinguishing between healthy and HD cells. Overall, the integration of single-cell RNA sequencing data and deep learning models provides valuable insights into HD pathology.

Conclusion We developed the model capable of analyzing HD at single-cell transcriptomic level.

Keywords Deep residual model, Huntington's disease, Single-cell RNA-seq, Unveiling potentially key genes

[†]Shichen Gao and Yadong Wang contributed equally to this work.

*Correspondence:

Yan Dong
dongyan_bio@126.com

¹School of Life Sciences, Anhui University, Hefei 230601, China

²College of Biology and Food Engineering, Chuzhou University, Chuzhou 239000, China



Background

Huntington's disease (HD) is a neurodegenerative disorder caused by a CAG-repeat expansion in the HTT gene, which leads to the progressive degeneration of nerve cells. The CAG mutation causes a polyglutamine expansion in huntingtin, which results in misfolding and aggregation of polyglutamine (polyQ)-huntingtin. These soluble aggregates disrupt various cellular processes through gain and loss of function, leading to transcriptional dysregulation and other pathologies. Despite the crucial role of *HTT* in neurodevelopment, its precise function remains elusive [1]. *HTT* is intricately involved in various cellular activities associated with cytoplasmic and nuclear functions, including gene expression, intracellular signaling, transport, metabolism, neurogenesis, and DNA repair [2]. The expansion of a CAG trinucleotide DNA sequence, encoding a polyglutamine motif in *HTT*, is implicated in the pathogenesis of HD. Although the CAG-repeat expansion mutation within the Huntington locus (4p16.3) in Exon1 of the *HTT* gene is well-documented, the mechanisms underlying neuron apoptosis and brain susceptibility remain unclear [1–3].

Transcriptional dysregulation alters downstream cellular processes, as inferred from sequencing data [4]. Identifying the roles of various genes and their expression patterns over the course of disease onset and progression remains challenging. Nonetheless, evidence suggests that the polyQ protein in *HTT* exon 1 tends to form amyloid-like fibrillar aggregates, potentially mediating toxicity [5]. Protein aggregation is common in neurodegenerative disorders like Alzheimer's disease, Parkinson's disease, and amyotrophic lateral sclerosis, but in Huntington's disease, this aggregation particularly involves polyQ *HTT* protein.

Recent research has revealed significant findings about HD. For instance, Snell and Handley discovered that urea metabolism is disrupted in both HD patients and a sheep model, detecting elevated urea levels early in the disease. This suggests a potential systemic issue and indicates that targeting urea metabolism could be a viable therapeutic strategy for HD [6]. Similarly, Carroll and Bragg's research on *Htt*Q111/+ mice showed early, progressive, striatal-specific molecular changes and subtle behavioral alterations, emphasizing the importance of knock-in models for studying early therapeutic interventions [7]. Hood and Ament also highlighted the value of *Htt*Q111/+ mice in reflecting human HD pathology and testing neuroprotective therapies [8]. Furthermore, Carroll and Coffey demonstrated that these mice develop early molecular changes and behavioral alterations similar to early HD, underscoring their potential for early therapeutic intervention studies [9]. Price and Ament characterized *Htt*Q111/+ mice as generally healthy during their first year, with early striatal-specific changes, supporting their use in studying early therapeutic

strategies and their better translatability to human HD compared to transgenic models [4]. Additionally, Ament and Malaiya identified core transcription factors involved in HD by reconstructing transcriptional regulatory networks in the mouse and human striatum, which could contribute to understanding and targeting HD [10]. Recent advances in single-cell transcriptomic data analysis have elucidated some pathogenic mechanisms underlying HD. Transcriptomic profiling revealed cell cycle-related signaling predominantly in the mitotic cell population, namely, neuronal stem cells (NSCs), rather than in coexisting striatal neurons [11]. Charlene et al. identified a sustained presence of NSCs due to aberrant WNT signaling, which could be ameliorated by WNT inhibition, thereby establishing a potential therapeutic target for future exploration [11].

These findings highlight the challenges associated with transcriptional data analysis and provide a rationale for implementing machine learning algorithms to improve the interpretation of sequencing data. State-of-the-art machine learning techniques, such as deep learning algorithms, are increasingly used to map transcriptomic profiles to phenotypic variations, such as those in tissue types, cancer staging and grading, drug responses, disease outcomes, and glycan features, thereby facilitating a molecular system-level understanding and interpretation of phenotypic outcomes [12–15]. In terms of model interpretation, the use of SHapley Additive exPlanations (SHAP) to elucidate deep learning models aids in addressing biological questions [16, 17]. Traditional differential expression analysis (DEA) exhibits two main drawbacks: (1) susceptibility to information loss due to the arbitrariness of p-values and fold change thresholds [18, 19] and (2) a bias toward highly expressed genes [18–20]. Using SHAP, Yap et al. discovered genes that were overlooked using DEA [14], thereby underscoring the ability of SHAP to detect subtle yet significant differences.

Herein, we established a deep residual neural network (ResNet) to differentiate between healthy and HD cells using single-cell transcriptomic data and applied SHAP to obtain biological insights into the early processes involved in HD. To the best of our knowledge, deep learning algorithms using transcriptomes as input have not been previously employed in HD research. Our approach highlights the potential of interpretable deep learning models in revealing novel regulatory mechanisms underlying HD, providing a new avenue for understanding the disease at a molecular level.

Methods

scRNA-seq data processing

Single-cell RNA sequencing (scRNA-seq) data were sourced from the Cellular Drug Response Atlas database

(Project ID is GEND000121) and the study “Single-nucleus RNA-seq identifies Huntington disease astrocyte states” [21, 22]. Sequencing data preprocessing for each sample was conducted using Cell Ranger (version 7.1.0; 10X Genomics) with default parameters and Seurat R package (version 4.1.3), providing gene expression matrices [23]. Initial processing of the gene expression matrices involved the following steps: (1) genes expressed in ≤ 3 cells and cells expressing ≤ 200 genes were removed. The percentage of mitochondrial reads tolerated per cell was set at 15%. (2) Normalization was performed using SCTransform function. (3) Principal component analysis (PCA) on the 2000 most variable features was conducted using the FindVariableFeatures function, retaining 10 dimensions. (4) Cells were clustered using the FindClusters function, employing the Louvain algorithm with a resolution of $r=1.7$. The neighbors used for clustering were identified using the FindNeighbors function. (5) Marker genes were determined using the FindAllMarkers function, performing the Wilcoxon rank-sum test with `min.pct` and `logfc.threshold` parameters set to 0.1.

Model training

We designed a residual neural network (ResNet) named DeepHD, which incorporates additional residual

layers, enhancing its performance on imbalanced datasets (Fig. 1).

The input data was randomly divided into a training set (80%) and a validation set (20%). Utilizing Python 3.10.9, we constructed the ResNet, deep neural network, and convolutional neural network, training each model with the same training, validation, and test sets.

For the model architecture in ResNet, models were trained using the PyTorch framework with mini-batches of size 64. The best-performing model featured one hidden convolutional layer, three hidden residual layers, and two hidden fully connected layers. The convolutional layer (6 channels; kernel size: 1; stride: 1) was followed by a batch normalization layer. The first residual layer comprised two plain convolutional layers (the first: 6 channels; kernel size: 3; stride: 4; padding: 1, the second: 6 channels; kernel size: 3; stride: 1; padding: 1), each followed by batch normalization and leaky ReLU activation layers (negative slope 0.01). Shortcut connections (6 channels; kernel size: 1, stride: 4) were inserted to convert this plain convolutional network into its residual counterpart. Another residual layer comprised two plain convolutional layers (6 channels; kernel size: 3; stride: 1; padding: 1), each followed by batch normalization and leaky ReLU activation layers (negative slope 0.01). Furthermore, shortcut connections (6 channels; kernel

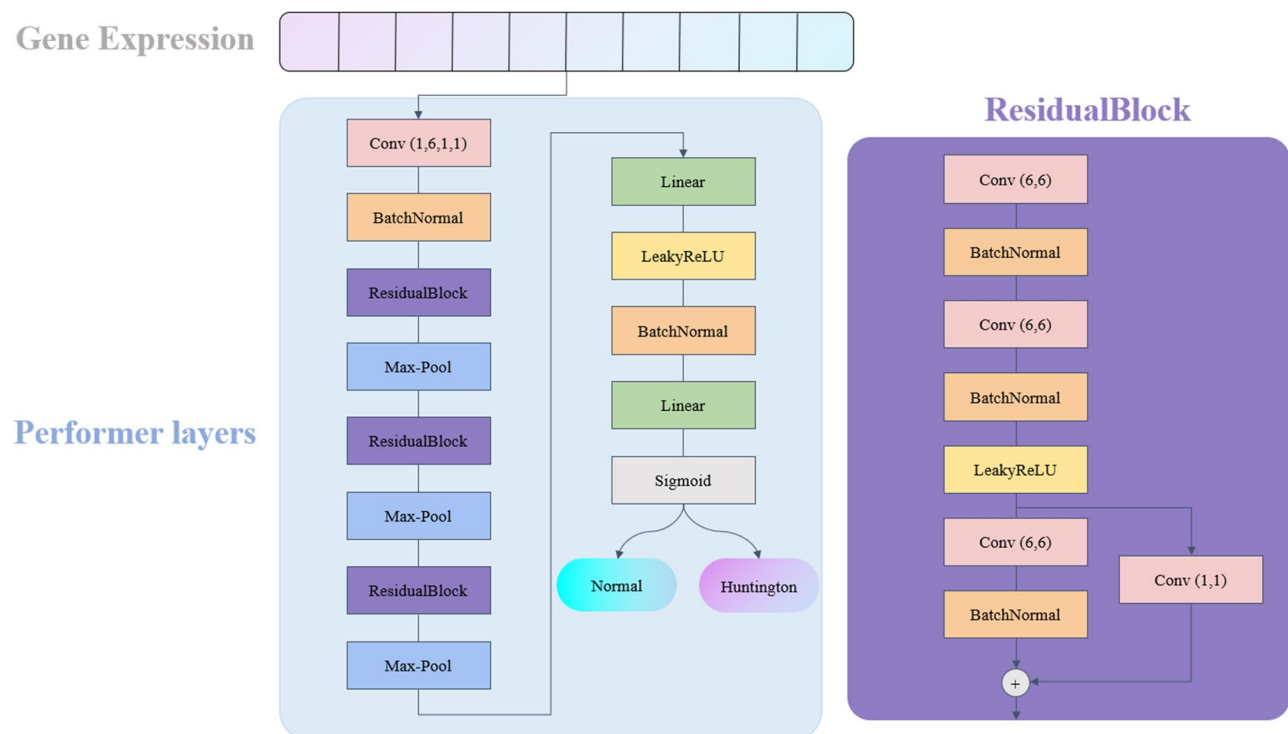


Fig. 1 Graphical presentation of the residual neural network structure, a deep learning architecture for analyzing single-cell gene expression data, focuses on distinguishing between normal and Huntington's disease samples. It consists of input layers, performer layers with CNN blocks, and residual blocks containing linear layers, activations, and convolutions. This network can classify samples as either normal or Huntington's disease. Residual connections help train deeper networks effectively

size: 1, stride: 1) were inserted to transform this into its residual version. Each residual layer was followed by max pooling layers (filter size: 2; stride: 2). The first fully connected layer (16 nodes) was followed by a leaky ReLU activation layer (negative slope 0.01) and a batch normalization layer. The second fully connected layer was followed by a sigmoid activation layer. Moreover, the initial learning rate was set to 0.001. The binary cross-entropy loss function and a cosine annealing learning rate scheduler were employed for model optimization.

The architectures for the deep neural network and convolutional neural network used for comparison were derived from Daniel Bojar's previous research [15].

Differential gene expression analysis

We used the Seurat functions FindMarkers and FindAllMarkers for differential gene expression analysis. Particularly, FindMarkers was used to identify marker genes between HD cells and normal cells, while FindAllMarkers was employed to identify marker genes between each cluster and other cells. We set the absolute log₂(fold change) threshold to the default value of 0.25, and the Wilcoxon ranked sum test was used to determine p-values. All p-values for each cell type and analysis test were adjusted for multiple testing using the Benjamini–Hochberg correction.

Model interpretation with SHAP

We use the DeepExplainer function of the SHAP library (version 0.40.0) implementation to calculate SHAP values [24]. Based on our computational resources, SHAP values were computed for 2000 cells randomly selected from the entire dataset. The SHAP values used for biological interpretation were the mean absolute values of the SHAP values for each feature.

Permutation feature importance analysis

Keeping the labels unchanged, each feature (gene) column was randomly permuted. The model made predictions using the permuted expression values, and prediction loss was calculated using the binary cross-entropy loss function. A total of 25 permutations were performed, and the average loss from all permutations was computed for each gene. Permutation feature importance for a gene was defined as the mean model loss minus the original model loss (i.e., the model loss when making predictions with the original expression values).

Analysis of individual joint feature contributions

A SHAP dependence plot was created using the dependence plot function from the SHAP library. This plot illustrated the relationship between the feature's value (x-axis) and its corresponding SHAP value (y-axis). The color of the plot represented an interaction feature. The

vertical dispersion of the data points indicated interaction effects. Gray ticks along the y-axis represented data points where the feature's value was NaN.

Pathway enrichment analysis and protein interaction network analysis

SHAP gene ontology pathway enrichment analysis was performed via the online portal of ShinyGo version 0.77 (<http://bioinformatics.sdstate.edu/go77/>) [25]. Genes with the top 10% (90th percentile) median absolute SHAP values were compared to all genes in the dataset (background genes; top 10% genes included). The false discovery rate cutoff was set at 0.05. DEA Gene Ontology (GO) pathway enrichment analysis was performed using the R tool “clusterProfiler” [23, 26].

Results

Single-cell RNA sequencing analysis revealed differential expression between healthy and adult-onset HD cells

The analysis of single-cell transcriptomic data can markedly enhance our understanding of the pathogenic mechanisms underlying HD. Therefore, we used recently accessible single-cell RNA sequencing data [11], which includes two datasets derived from induced pluripotent stem cells (iPSCs) of (i) unaffected control individuals and (ii) HD patients (Fig. 2, left panel). These iPSCs were differentiated into neuron cultures enriched with medium spiny neurons, the cell type predominantly impacted in HD. During differentiation, a persistently observed population of SCs positive for cell-cycle protein cyclin D1 (D1+) was selectively noted in adult-onset HD iPSC-derived neurons.

A list of upregulated marker genes was prepared for each cell type (Table S1). Using this list, we performed GO analysis and revealed specific enrichment of genes related to cell identity in different cell types (Fig. 3).

In line with the designated annotations, mature neurons exhibit dominant enrichment in processes related to macroautophagy [27], proteasome-mediated ubiquitin-dependent protein catabolic processes, Golgi vesicle transport [28, 29], proteasomal protein catabolic processes, and endosomal transport [30]. These processes suggest a robust system for cellular maintenance, protein turnover, and trafficking. NSCs demonstrate potent enrichment in cytoplasmic translation [31], ribonucleoprotein complex biogenesis [32], ribosome biogenesis, rRNA processing, and rRNA metabolic processes [33]. This profile indicates high activity in protein synthesis and regulation, necessary for future differentiation and specialization. Spiny neurons positive for cyclin D1 (D1+) and cyclin D2 (D2+) were notably enriched in sterol and cholesterol biosynthesis and metabolism as well as secondary alcohol metabolic processes [34]. This suggests an essential role of lipid metabolism in these neurons,

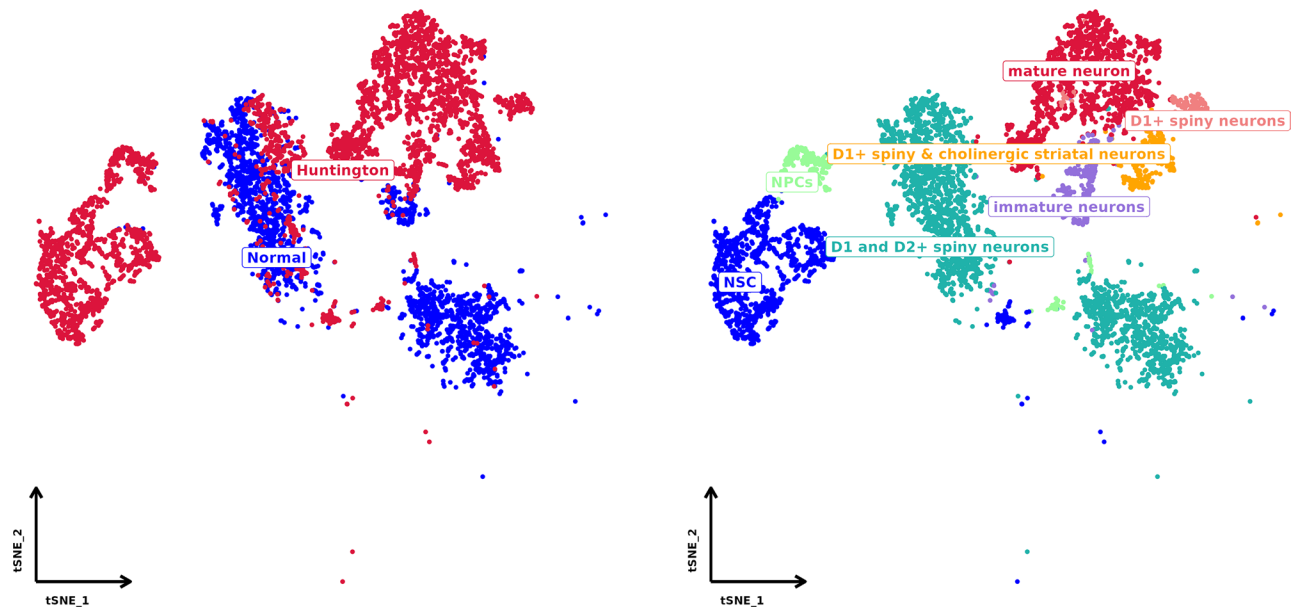


Fig. 2 t-SNE visualization of the integrated single-cell RNA sequencing data. The left panel depicts the distribution of healthy and Huntington's disease samples within the t-SNE plot, while the right panel illustrates the distribution of various cell types in the same t-SNE plot

which could be associated with their ability to modulate neuronal signals. Moreover, spiny and cholinergic striatal neurons positive for cyclin D1 (D1+) were involved in regulating transsynaptic signaling [35], modulating chemical synaptic transmission, protein localization to the cell periphery, and axon development. These neurons might also play roles in neuronal communication and neural circuit formation. Immature neurons exhibit specific enrichment in response to endoplasmic reticulum stress [36], proteasomal protein catabolic process, proteasome-mediated ubiquitin-dependent protein catabolic process, in response to topologically incorrect protein [37], and Golgi vesicle transport [38]. These enrichments reflect responses to protein misfolding, indicative of early-stage adaptation mechanisms in HD neurons. Neural progenitor cells (NPCs), the neuron precursors, exhibited enrichment in cytoplasmic translation, axon development [39], and translational initiation, reflecting their essential role in NPC development. Moreover, D1+spiny neurons exhibit a specific signature of aerobic respiration, oxidative phosphorylation, cellular respiration [40], ATP metabolism [41], and ATP synthesis coupled with electron transport, suggesting a high energy demand and an important role in maintaining the electrical activity of neural network. Overall, our preliminary analyses identified several biologically relevant cell clusters (Fig. 2, right panel).

Training a deep residual model to predict HD or healthy phenotypes from the transcriptome

We used ResNet to model HD cells based on transcriptome-wide gene expression data; Fig. 2 illustrates their

distribution in the t-distributed stochastic neighbor embedding (t-SNE) plot. Next, we assigned binary labels to cells based on cell-type data. In particular, we labeled healthy cells as “label 0” and HD cells as “label 1” to establish robust labels comparable between datasets. This robustly separated biologically distinct populations and enabled subsequent analyses. Accordingly, we performed a binary classification, using gene expression values as input and the probabilities for the positive phenotype (label 1) of the corresponding cells as output.

We developed an artificial neural network classifier, including three residual layers (Fig. 1), whose workflow is shown in Fig. 4. In the test set, among the three artificial neural network architectures, our designed architecture surpassed the other two in terms of both AUC and F1 score (Fig. 5a–d). This classifier achieved an accuracy of 98.87% for predicting the label 0 phenotype and 96.54% for predicting the label 1 phenotype (AUC: 0.9967, F1 score: 0.9683; Fig. 5a and d). In the existing standard neural network, the F1 score was lower, only 0.8245 (Fig. 5b). The prediction accuracies were 98.34% for the label 0 phenotype but only 73.07% for the label 1 phenotype, indicating a significant imbalance (Fig. 5e). The existing convolutional neural network performed even worse, with an accuracy of 98.39% for the label 0 phenotype and only 74.49% for the label 1 phenotype (Fig. 5e). This decline in performance was likely due to imbalanced training data, further demonstrating the robustness of our designed ResNet model in handling imbalanced data.

Figure 6 illustrates the training and validation accuracy (a–c) and loss curves (d–f) for the residual network based on three distinct datasets: iPSC-derived neurons,




Fig. 3 Enrichment in biological processes for marker genes in each cluster and associated p-values were generated using clusterProfiler. Each point in the plot corresponds to a specific biological process, and its size indicates the count of genes associated with that process. Each point is color-coded according to the significance level (p.adjust) involved in that particular biological process, with darker red points representing more statistically significant enrichments. The x-axis labels represent different cell types or neuron subpopulations, including D1 and D2 + spiny neurons, D1 + cholinergic striatal neurons, D1 + spiny neurons, immature neurons, mature neurons, NPCs (Neural Progenitor Cells), and NSCs (Neural Stem Cells). The y-axis lists the enriched biological processes. For a complete list of GO enrichment results, see Supplemental Table S1

astrocytes, and neurons. The training and validation accuracy curves for all three datasets rapidly stabilized around 98–99% after a few epochs, indicating high accuracy. The loss curves showed that the training loss steadily decreased and stabilized around 0.05, while the validation loss remained slightly higher and more variable, stabilizing around 0.1 for iPSC-derived neurons and astrocytes datasets, and around 0.2 for the neuron dataset. Despite fluctuations in validation loss, the overall performance of the residual network was consistent and robust across datasets.

Identification of genes with high Shapley value using SHAP

Subsequently, we used SHAP to identify significant genes for prediction. For every input sample (cell), the SHAP algorithm computes a SHAP value for each feature (gene), indicative of the feature's influence on the anticipated model output for that particular input [24]. SHAP values may be positive or negative, denoting additive or subtractive impacts on the model output. The median absolute SHAP value is often used to evaluate global feature importance.

Next, we computed the median absolute SHAP values for all genes (Table S2) to obtain gene rankings. In both healthy and Huntington models, only a small subset of genes had high importance. We selected genes

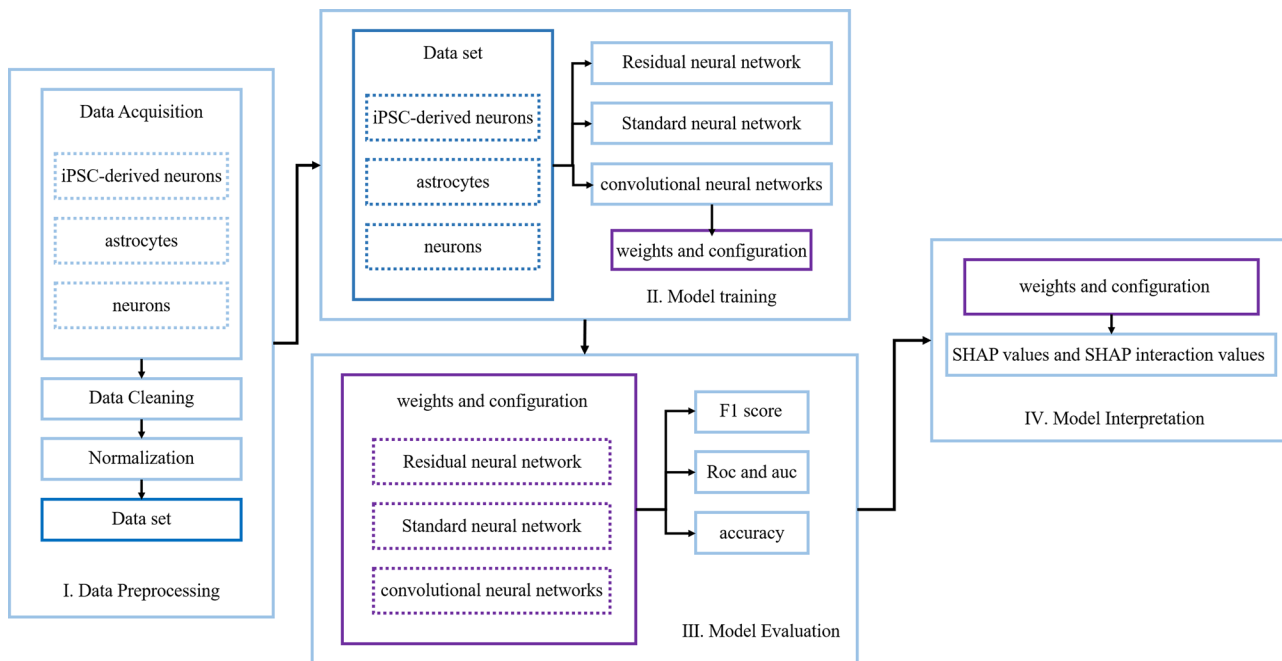


Fig. 4 The workflow of an interpretable deep learning framework includes data preprocessing, model training, model evaluation, and model interpretation

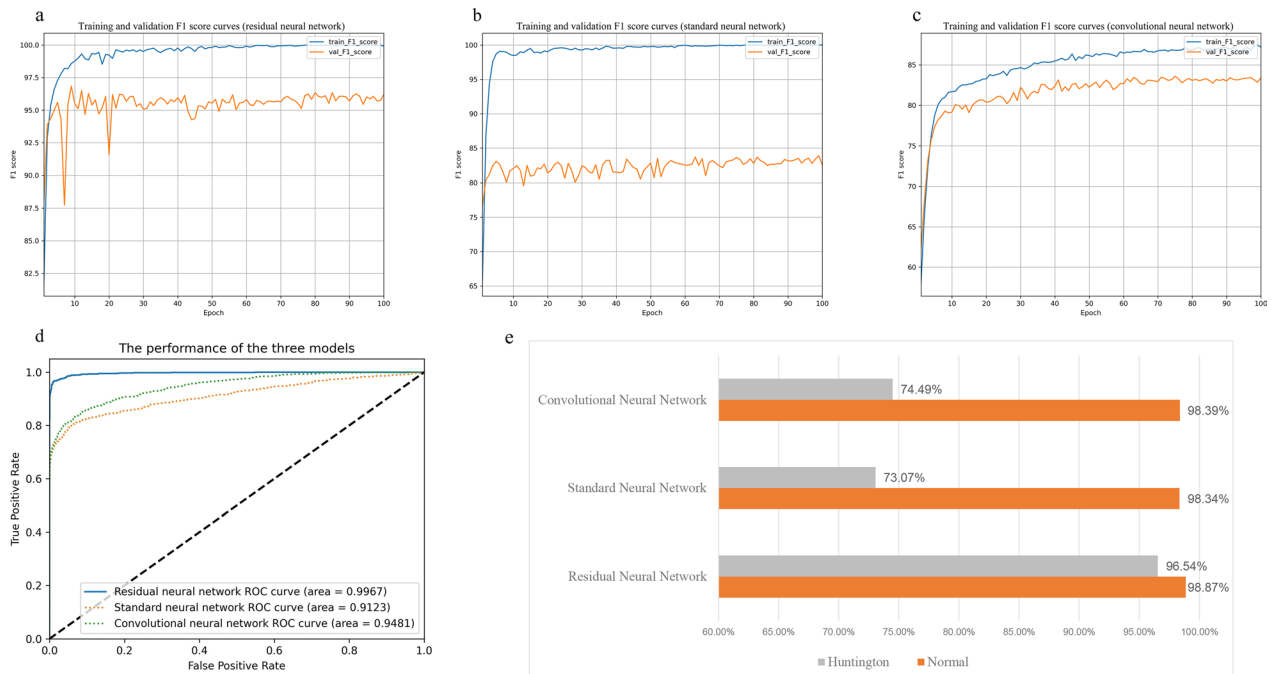


Fig. 5 Performance comparison of the three artificial neural network architectures (residual neural network, standard neural network, convolutional neural network). Figures a to c show the F1 score curves for the three architectures over epochs on the training set (blue curve) and validation set (orange curve). Figure d displays the ROC curves on the validation set, with the blue solid line representing the ROC curve of the residual neural network, the green dashed line representing the ROC curve of the convolutional neural network, and the orange dashed line representing the ROC curve of the standard neural network. Figure e illustrates the accuracy of the three architectures in predicting labels for Huntington and Normal on the same dataset

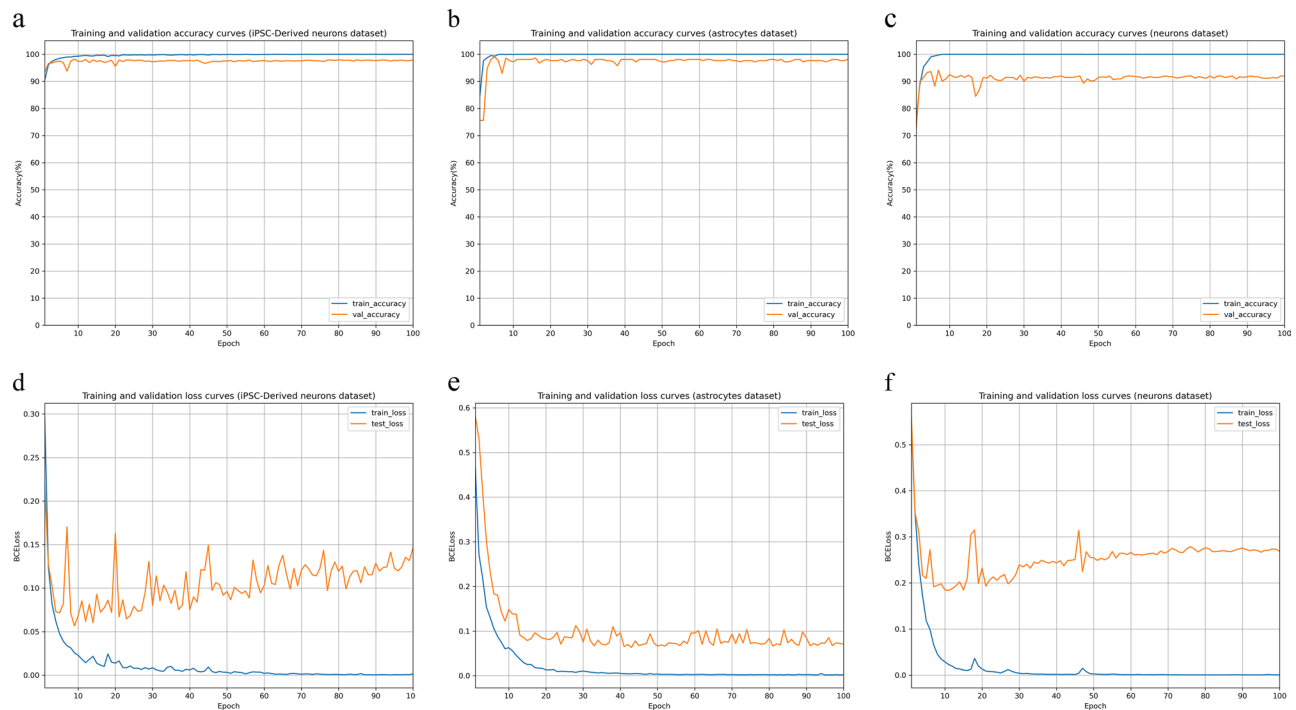


Fig. 6 Performance of the residual neural network across three datasets. Figures a to c show the accuracy curves over epochs for the iPSC-derived neurons dataset, astrocytes dataset, and neurons dataset, respectively, with the blue curve representing performance on the training set and the orange curve representing performance on the validation set. Figures d to f display the BCELoss curves over epochs for the iPSC-derived neurons dataset, astrocytes dataset, and neurons dataset, respectively, with the blue curve representing performance on the training set and the orange curve representing performance on the validation set

corresponding to the top 25 SHAP rankings (SHAP genes). For most SHAP genes (e.g., *MALAT1*, *RPS2*, and *RPS6*; Fig. 7), increased expression tended to hint toward HD prediction, while the reverse was true for a smaller subset (e.g., *FTH1*, *EEF1A1*, and *VAMP2*; Fig. 7).

Next, we performed pathway enrichment analysis on SHAP genes. The most enriched pathways (biological processes) included SRP-dependent cotranslational protein targeting to the membrane and protein targeting to the ER (Fig. 8, left panel), aligning with established roles in HD [42, 43].

SHAP genes explain the pathophysiology of HD

One of our goals was to determine whether the most predictive SHAP genes identified through this data-driven approach were involved in the pathophysiology of HD. A highly predictive gene can be biologically associated with HD in several ways, including the following: (i) alternative splicing, transcriptional regulation, and posttranscriptional regulation [44]; (ii) *HTT*-mediated ribosome stalling [45, 46]; and (iii) direct impairment of mitochondrial function by *mHTT* as well as indirect dysregulation of transcriptional processes [47]. Indeed, many highly predictive SHAP genes (top 25 SHAP genes) have been implicated in HD, suggesting their high biological relevance. These genes were also enriched in

ribosome-related and NADH-related functions, as previously shown [45, 48].

We first used a t-SNE plot to evaluate the expression levels of high-ranking SHAP genes and observed their consistent upregulation across HD cell populations (Fig. 9 and S1), which indicated their likely role in the cellular response to HD. Among these genes, the metastasis-associated lung adenocarcinoma transcript 1 (*MALAT1*) lncRNA is involved in alternative splicing, transcriptional regulation, and posttranscriptional regulation, acting as a miRNA sponge that interacts with RNA.

FTH1 encodes the ferritin heavy subunit, a major intracellular iron storage protein in both prokaryotic and eukaryotic cells. At the expression level, specific *FTH1* upregulation was observed in HD cells (Fig. 9), suggesting the importance of disrupted brain iron homeostasis in the pathogenesis of neurodegenerative diseases and providing a novel perspective on the role of iron in neurodegeneration.

In our experiments, mitochondrial cytochrome c oxidase genes (*MT-CO3*, *MT-CO2*, and *MT-CO1*) exhibited specific upregulation in HD cells (Fig. 9 and S1), as did genes encoding ribosomal proteins (RPs) (e.g., *RPS2*, *RPS6*, and *RPS3A*). Unlike *MT-CO3*, RP genes were predominantly expressed in NSCs and NPCs, with decreased



Fig. 7 SHAP summary plot containing the 25 features with the trained residual neural network across the three datasets. The SHAP contributions for each data point are summed for each computed gene score




Fig. 8 The left panel shows the Gene Ontology pathway enrichment analysis of the SHAP genes. The x-axis represents fold enrichment, indicating how much more likely the given biological process is represented among the genes of interest compared to the background population. The y-axis shows the biological processes themselves, sorted in decreasing order of significance. Each circle in the chart corresponds to a specific biological process, with the size of the circle representing the number of genes involved in each process, ranging from 100 to 350 genes. The color of the circles corresponds to the $\log_{10}(\text{FDR})$ values, which measure the statistical significance of the enrichment. Higher values indicate lower FDR, meaning greater confidence in the observed enrichment. The right panel displays a volcano plot showing differentially expressed genes between HD and healthy cells. Blue circles represent genes with significantly downregulated expression, red circles represent genes with significantly upregulated expression, and gray circles denote genes with nonsignificant differences. The significance threshold was set at $P < 0.05$

expression in mature neurons, consistent with previous research [49].

To examine the interaction of a prominent high-ranking SHAP gene with other features, we plotted the prominent feature value against its SHAP value across the entire dataset. In addition, we colored the value of

several other features with strong interactions on the prominent feature. Figure 10 and S2 show that, although most *RPS6* values were < 25 , the extent of *RPS6* impact on the prediction varied, as evidenced by the vertical dispersion of dots at values < 25 . This suggests that other features influence the contribution of *RPS6*. Conversely,

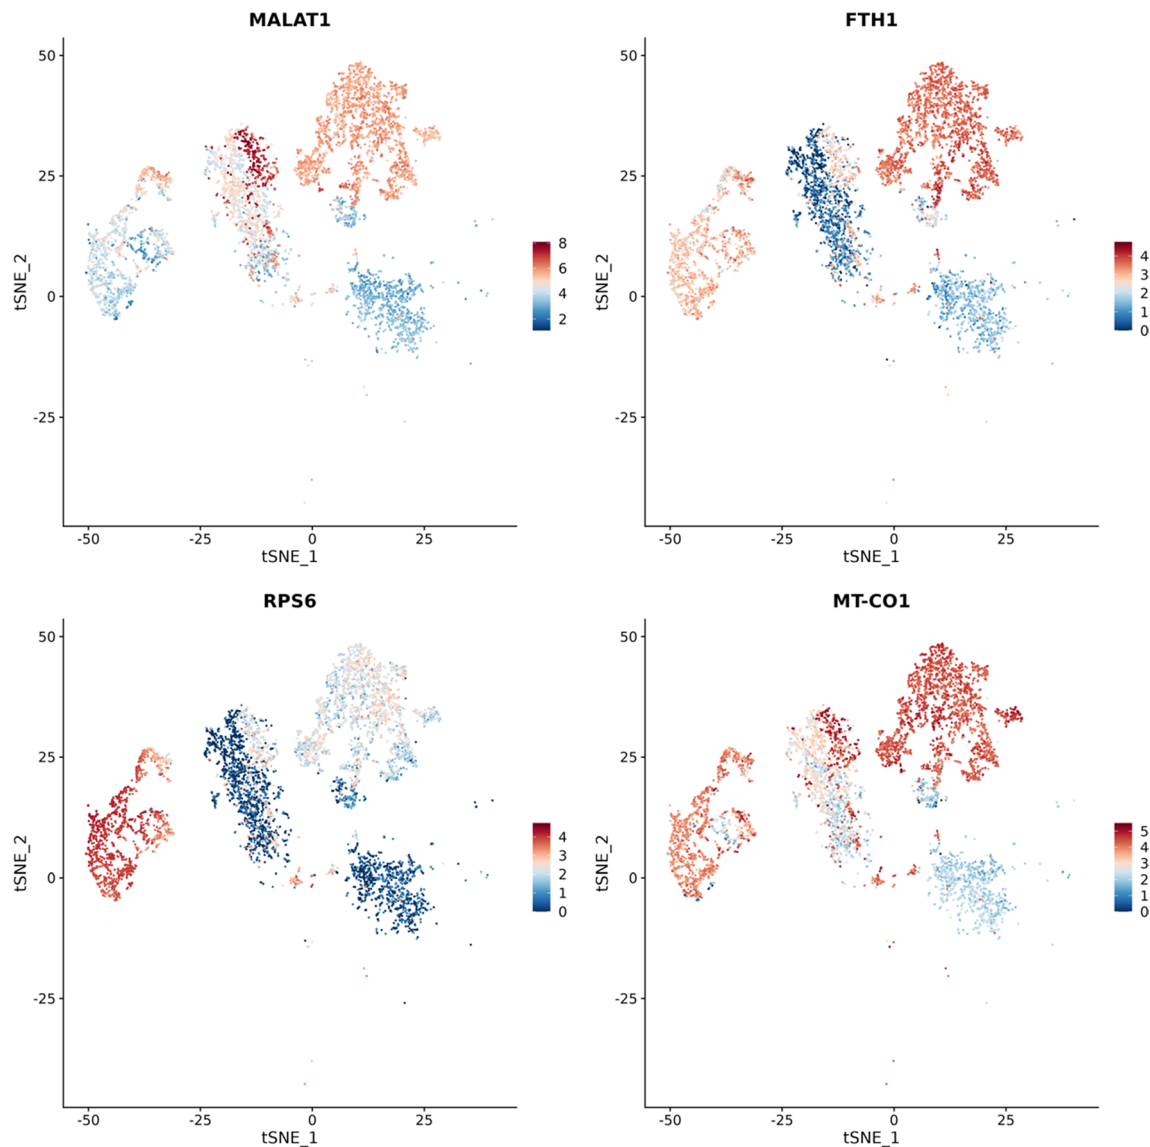


Fig. 9 Feature plot of cells that are positive for each of the four high-ranking SHAP genes, respectively, in the t-SNE projection

for values >25 , the interaction effect significantly diminished, as indicated by the lower vertical dispersion of the dots. The distribution of red dots above the y-axis 0.0 and the values of *RPS3A*, *RPL10*, *RPL3*, *RPL13*, and *RPS14* on the sample dots imply that higher *RPS3A*, *RPL10*, *RPL3*, *RPL13*, and *RPS14* expression levels correspond to a more positive contribution from *RPS6*. This indicates cooperative expression of RP genes. However, the color of dots representing *MALAT1* suggests that *MALAT1* attenuates the contribution of *RPS6* for values >25 . This implies that higher *MALAT1* expression levels likely decrease *RPS6* expression. Similarly, the colors of dots representing *FTH1*, *MT-CO1*, *MT-CO2*, and *MT-CO3* suggest that high expression levels of these genes positively influence the contribution of *RPS6*.

SHAP genes include low-abundance genes not captured through DEA

SHAP can reveal gene subsets overlooked by traditional DEA [14]. We tested this in our dataset as well as evaluated whether any SHAP-exclusive genes had biological relevance. In our data, DEA between HD and healthy populations applying a false discovery rate threshold of 0.05, indicated 762 DEGs (3.7% of all genes; Table S3), significantly fewer than the number of SHAP genes detected (2047 genes or 10% of all genes).

DEA tends to favor highly expressed genes [18–20]. Among DEA genes, only a small subset overlapped with top-ranking SHAP genes, such as *RPS3A*, *RPS2*, and *MALAT1* (Fig. 8, right panel). DEA predominantly captures highly abundant housekeeping genes, whereas SHAP exhibits less inclination toward high-abundance

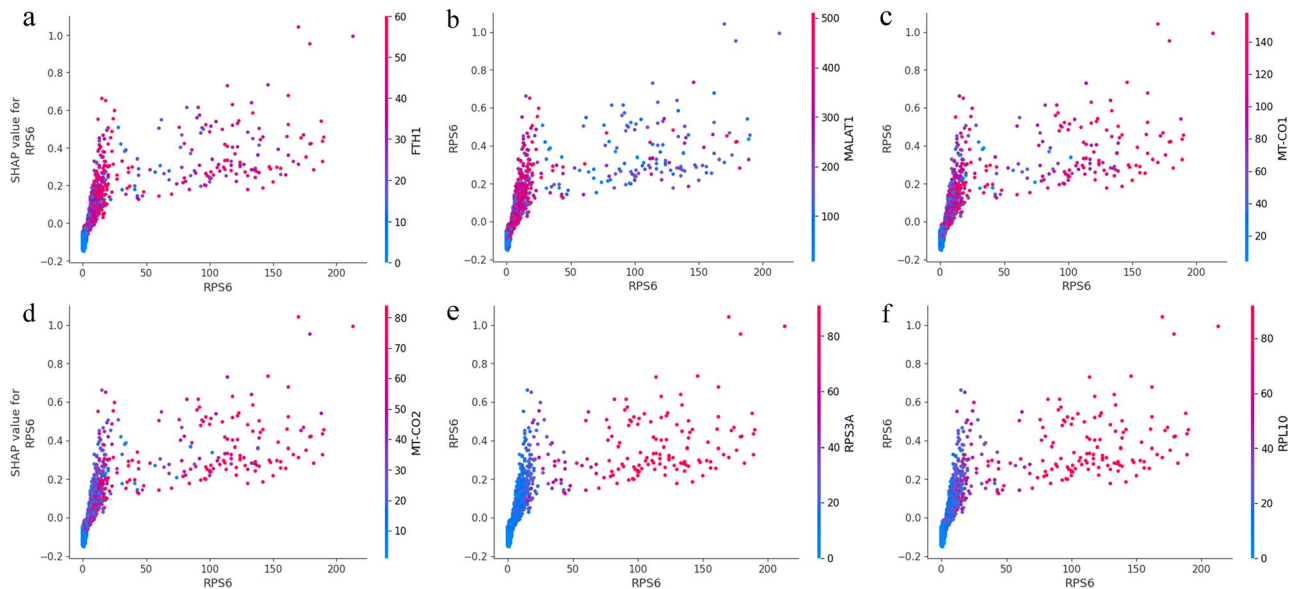


Fig. 10 The SHAP value to feature value plot for the *RPS6* gene, with the respective other gene (*FTH1*, *MALAT1*, *MT-CO1*, etc.) values color-coded. Positive SHAP values indicate an association with the high expression class, while negative SHAP values indicate an association with the low expression class of genes

genes. Notably, mitochondrial genes implicated in neurodegenerative diseases, such as the members of the cytochrome C oxidase gene family (*MT-CO1*, *MT-CO2*, and *MT-CO3*), were absent among DEA genes. Overall, our findings suggest that SHAP is more effective than traditional DEA in identifying low-abundance genes of significant biological relevance for distinguishing between healthy and HD cells.

Discussion

To the best of our knowledge, this is the first application of deep learning models in HD research at the transcriptomic level. Currently, all applications of deep learning models in HD research are focused on radiomics or to predict macroscopic behaviors [50, 51]. Consequently, research on the underlying mechanisms of HD pathogenesis has not yet significantly benefited from the rapid advancement of computational tools, which are used for predicting splicing from primary sequences [52]. Despite the wide use of state-of-the-art artificial intelligence technologies for single-cell sequencing data [15], they remain underutilized in HD research. Single-cell RNA has just started being used to investigate the pathogenesis of HD using deep learning, as demonstrated in this study.

Previous studies have explored the feasibility of predicting phenotypes using single-cell RNA data and deciphering biological processes through model interpretation, with model classification accuracy of >85% (AUC: 0.9123, F1 score: 0.8245) [15]. However, significant disparities in classification accuracy can arise with imbalanced training data. The accuracy of predicting healthy results in the standard neural network is 98.34%, whereas

that of predicting HD is only 73.07% (Fig. 5e). Therefore, we refined our model by constructing a ResNet model for high-precision prediction. The accuracy of predicting healthy cells improved to 98.87%, whereas that of predicting HD increased to 96.54 (Fig. 5e).

Using SHAP, we further demonstrated that, in addition to predicting cell phenotypes, our model could be utilized for large-scale extraction of biological correlations between transcriptomic data and disease mechanisms. Notably, the most crucial genes for predicting the HD phenotype significantly overlapped with those obtained through DEA. However, DEA tends to favor highly expressed genes [18–20]. SHAP allowed us to identify the role of mitochondrial genes, which was not detected with DEA. Existing evidence also indicates an increase in ribosomal occupancy of *MT-CO1*, *MT-CO2*, and *MT-CO3* [53]. Previous research suggested an increase in mitochondrial oxidative phosphorylation (mt-OXPHOS) transcript (*MT-CO3*, *MT-CO2*, and *MT-CO1*) occupancy in HD cells [54]. Furthermore, previous studies have hypothesized that the binding of mHTT to mitochondrial membranes can induce ribosomal occupancy in the matrix, thereby regulating mitochondrial protein synthesis through transmembrane mitochondrial signaling [54].

Herein, we employed joint feature contribution analysis based on SHAP values to investigate the collaborative contributions among RP genes, mitochondrial cytochrome c oxidase genes, *FTH1*, and *MALAT1*. High expression of mitochondrial cytochrome c oxidase genes and *FTH1* promotes *RPS6* expression, while *MALAT1* has the opposite effect. RP genes display similar

expression patterns. Recent studies have indicated that the lysosomal autophagy pathway plays a crucial role in iron distribution within cells, with defects in this pathway being often associated with the pathogenesis of neurodegenerative diseases [55]. MALAT1 is highly expressed in the brain and is involved in synaptic formation and other neurological pathways [44]. Furthermore, there is a close association between MALAT1 and the pathogenesis of Parkinson's disease, where it may serve as a diagnostic biomarker [44]. In this study, MALAT1 was particularly upregulated in HD spiny neurons (Fig. 9), suggesting its potential relevance in HD pathogenesis.

Our study demonstrated that (i) cell phenotypes can be modeled through artificial neural networks using transcriptomic data, (ii) biological processes associated with HD pathology can be deciphered through ResNet model interpretation using SHAP, and (iii) this combined approach may help in the discovery of new pathogenic mechanisms and therapeutic strategies for HD.

Our study had some limitations. Notably, our model was trained on a specific cell type. Thus, it cannot be generalized to other cell types or diseases. For new cell types, the model must be trained on new data. Although the same model architecture can be used for this purpose, some important genes with low expression may not always be measured. In addition, our joint contribution analysis did not reveal the actual regulatory relationships between genes. Future research should focus on collecting as much data as possible on a wide range of diseases. In addition to transcriptomic data, integrating more data modalities, such as proteomics or miRNA data, into ResNet can enhance the robustness of experimental conclusions. Second, our findings were validated by knocking out SHAP genes, for instance, using morpholino antisense oligos (MASOs) [56]. Finally, existing gene regulatory networks were explored to validate the results of the joint contribution analysis.

Conclusions

In summary, this study pioneers the application of deep learning models in HD research at the transcriptomic level, differing from the conventional focus on radiomics or macroscopic behavior prediction. Our findings highlight the potential of single-cell RNA data and advanced artificial intelligence techniques in elucidating the intricate mechanisms of HD pathogenesis. Notably, refinement of our model, specifically the construction of a ResNet model, substantially enhanced prediction accuracy, paving the way for precise understanding of cell phenotypes and disease processes. Through SHAP analysis, we demonstrated critical genes implicated in HD pathology, shedding light on previously undetected molecular pathways. Although our study contributes valuable insights, limitations such as cell-type specificity

and incomplete regulatory understanding highlight avenues for future research. Thus, by integrating diverse data modalities and expanding dataset breadth, we anticipate further advancements in unraveling HD pathogenic mechanisms and developing therapeutic strategies.

Abbreviations

HD	Huntington's disease
ResNet	Residual neural network
DEA	Differential expression analysis
NSCs	Neural stem cells
SHAP	Shapley additive explanations
scRNA-seq	Single-cell RNA sequencing
t-SNE	t-distributed stochastic neighbor embedding
PCA	Principal component analysis
iPSCs	Induced pluripotent stem cells
GO	Gene ontology
NPC	Neural progenitor cell
mt-OXPHOS	Mitochondrial oxidative phosphorylation
RPs	Ribosomal proteins
MASOs	Morpholino antisense oligos

Supplementary Information

The online version contains supplementary material available at <https://doi.org/10.1186/s12864-024-10855-5>.

Supplementary Material 1
Supplementary Material 2
Supplementary Material 3
Supplementary Material 4
Supplementary Material 5

Acknowledgements

Not applicable.

Author contributions

YD designed the framework. SCG conducted the experiments and wrote the manuscript. YDW and JJW modified the manuscript. All the authors revised and approved the manuscript.

Funding

This study was funded by the Outstanding Youth Science Foundation of the Higher Education Institutions of Anhui Province (2022AH030110).

Data availability

We acquired scRNA-seq data from the China National Center for Bioinformatics (<https://download.cncb.ac.cn/gen/GEND000121/GENDX000121/>) and from the study "Single-nucleus RNA-seq identifies Huntington disease astrocyte states". The codes, training data, and test samples are openly accessible on GitHub at <https://github.com/pwerner-web/DeepHD>.

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare no competing interests.

Received: 10 May 2024 / Accepted: 30 September 2024

Published online: 04 October 2024

References

- Sharma V, Sharma P, Deshmukh R. Huntington's disease: clinical complexities and therapeutic strategies. *J Adv Sci Res.* 2012;3(02):30–6.
- Yu MS, Tanese N. Huntingtin is required for neural but not cardiac/pancreatic progenitor differentiation of mouse embryonic stem cells in vitro. *Front Cell Neurosci.* 2017;11:33.
- Biglan KM, et al. Refining the diagnosis of Huntington disease: the PREDICT-HD study. *Front Aging Neurosci.* 2013;5:12.
- Ament SA, et al. Transcriptional regulatory networks underlying gene expression changes in Huntington's disease. *Mol Syst Biol.* 2018;14(3):e7435.
- Trepte P, Stempel N, Wanker EE. Spontaneous self-assembly of pathogenic huntingtin exon 1 protein into amyloid structures. *Essays Biochem.* 2014;56:167–80.
- Handley RR, et al. Brain urea increase is an early Huntington's disease pathogenic event observed in a prodromal transgenic sheep model and HD cases. *Proc Natl Acad Sci.* 2017;114(52):e11293–E11302.
- Bragg RM, et al. Motivational, proteostatic and transcriptional deficits precede synapse loss, gliosis and neurodegeneration in the B6. Htt Q111/+ model of Huntington's disease. *Sci Rep.* 2017;7(1):41570.
- Ament SA, et al. High resolution time-course mapping of early transcriptional, molecular and cellular phenotypes in Huntington's disease CAG knock-in mice across multiple genetic backgrounds. *Hum Mol Genet.* 2017;26(5):913–22.
- Coffey SR, et al. Peripheral huntingtin silencing does not ameliorate central signs of disease in the B6. HttQ111/+ mouse model of Huntington's disease. *PLoS ONE.* 2017;12(4):e0175968.
- Malaiya S, et al. Single-nucleus RNA-seq reveals dysregulation of striatal cell identity due to Huntington's disease mutations. *J Neurosci.* 2021;41(25):5534–52.
- Smith-Geater C, et al. Aberrant development corrected in adult-onset Huntington's disease iPSC-derived neuronal cultures via WNT signaling modulation. *Stem Cell Rep.* 2020;14(3):406–19.
- Hanczar B, et al. Biological interpretation of deep neural network for phenotype prediction based on gene expression. *BMC Bioinformatics.* 2020;21:1–18.
- Jia P, et al. Deep generative neural network for accurate drug response imputation. *Nat Commun.* 2021;12(1):1740.
- Yap M, et al. Verifying explainability of a deep learning tissue classifier trained on RNA-seq data. *Sci Rep.* 2021;11(1):2641.
- Qin R, Mahal LK, Bojar D. Deep Learn Explains Biology Branched Glycans single-cell Sequencing data. *IScience.* 2022. 25(10).
- Wang Y, et al. XGraphCDS: an explainable deep learning model for predicting drug sensitivity from gene pathways and chemical structures. *Comput Biol Med.* 2024;168:107746.
- Tang Y-C, Gottlieb A. Explainable drug sensitivity prediction through cancer pathway enrichment. *Sci Rep.* 2021;11(1):3128.
- Bui TT, Lee D, Selvarajoo K. ScatLay: utilizing transcriptome-wide noise for identifying and visualizing differentially expressed genes. *Sci Rep.* 2020;10(1):17483.
- Yang W, Rosenstiel P, Schulenburg H. Afold—using polynomial uncertainty modelling for differential gene expression estimation from RNA sequencing data. *BMC Genomics.* 2019;20:1–17.
- Oshlack A, Wakefield MJ. Transcript length bias in RNA-seq data confounds systems biology. *Biol Direct.* 2009;4:1–10.
- Wang Y-Y, et al. CeDR Atlas: a knowledgebase of cellular drug response. *Nucleic Acids Res.* 2022;50(D1):D1164–71.
- Al-Dalahmah O, et al. Single-nucleus RNA-seq identifies Huntington disease astrocyte states. *Acta Neuropathol Commun.* 2020;8:1–21.
- Wu T et al. clusterProfiler 4.0: a universal enrichment tool for interpreting omics data. *Innov.* 2021. 2(3).
- Lundberg, S.M. and S.-I. Lee. A unified approach to interpreting model predictions, in Proceedings of the 31st International Conference on Neural Information Processing Systems. 2017, Curran Associates Inc.: Long Beach, California, USA. p. 4768–4777.
- Ge SX, Jung D, Yao R. ShinyGO: a graphical gene-set enrichment tool for animals and plants. *Bioinformatics.* 2020;36(8):2628–9.
- Yu G, et al. clusterProfiler: an R package for comparing biological themes among gene clusters. *OMICS.* 2012;16(5):284–7.
- Boland B, Nixon RA. Neuronal macroautophagy: from development to degeneration. *Mol Aspects Med.* 2006;27(5–6):503–19.
- Hanus C, Ehlers MD. Secretory outposts for the local processing of membrane cargo in neuronal dendrites. *Traffic.* 2008;9(9):1437–45.
- Horton AC, Ehlers MD. Dual modes of endoplasmic reticulum-to-golgi transport in dendrites revealed by live-cell imaging. *J Neurosci.* 2003;23(15):6188–99.
- Estrada-Bernal A, et al. Functional complexity of the axonal growth cone: a proteomic analysis. *PLoS ONE.* 2012;7(2):e31858.
- Xu X, et al. Folate regulates RNA m5C modification and translation in neural stem cells. *BMC Biol.* 2022;20(1):261.
- Hetman M, Slomnicki LP. Ribosomal biogenesis as an emerging target of neurodevelopmental pathologies. *J Neurochem.* 2019;148(3):325–47.
- Corsini NS, et al. Coordinated control of mRNA and rRNA processing controls embryonic stem cell pluripotency and differentiation. *Cell Stem Cell.* 2018;22(4):543–58. e12.
- Biolini G, et al. Striatal infusion of cholesterol promotes dose-dependent behavioral benefits and exerts disease-modifying effects in Huntington's disease mice. *EMBO Mol Med.* 2020;12(10):e12519.
- Yellajoshyula D, et al. Genetic evidence of aberrant striatal synaptic maturation and secretory pathway alteration in a dystonia mouse model. *Dystonia.* 2022;1:10892.
- Merighi A, Lossi L. Endoplasmic reticulum stress signaling and neuronal cell death. *Int J Mol Sci.* 2022;23(23):15186.
- Oldham MC, et al. Functional organization of the transcriptome in human brain. *Nat Neurosci.* 2008;11(11):1271–82.
- Rosso S, et al. LIMK1 regulates golgi dynamics, traffic of golgi-derived vesicles, and process extension in primary cultured neurons. *Mol Biol Cell.* 2004;15(7):3433–49.
- Pfeifer K, et al. Adult neural progenitor cells provide a permissive guiding substrate for corticospinal axon growth following spinal cord injury. *Eur J Neurosci.* 2004;20(7):1695–704.
- Wu B, et al. 2, 4 DNP improves motor function, preserves medium spiny neuronal identity, and reduces oxidative stress in a mouse model of Huntington's disease. *Exp Neurol.* 2017;293:83–90.
- Seong IS, et al. HD CAG repeat implicates a dominant property of huntingtin in mitochondrial energy metabolism. *Hum Mol Genet.* 2005;14(19):2871–80.
- Xiang C, et al. Bioinformatic gene analysis for potential therapeutic targets of Huntington's disease in pre-symptomatic and symptomatic stage. *J Translational Med.* 2020;18:1–10.
- Vidal R, et al. Converging pathways in the occurrence of endoplasmic reticulum (ER) stress in Huntington's disease. *Curr Mol Med.* 2011;11(1):1–12.
- Abrishamdar M, Jalali M, Rashno M. MALAT1 lncRNA and Parkinson's Disease: the role in the pathophysiology and significance for diagnostic and therapeutic approaches. *Mol Neurobiol.* 2022;59(9):5253–62.
- Eshraghi M, et al. Mutant huntingtin stalls ribosomes and represses protein synthesis in a cellular model of Huntington disease. *Nat Commun.* 2021;12(1):1461.
- Subramaniam S. Ribosome traffic jam in neurodegeneration: decoding hurdles in Huntington disease. *Cell Stress.* 2021;5(6):86.
- Jin YN, Johnson GV. The interrelationship between mitochondrial dysfunction and transcriptional dysregulation in Huntington disease. *J Bioenerg Biomembr.* 2010;42:199–205.
- Parker WD Jr, et al. Evidence for a defect in NADH: ubiquinone oxidoreductase (complex I) in Huntington's disease. *Neurology.* 1990;40(8):1231–1231.
- Gabut M, Bourdelais F, Durand S. Ribosome and translational control in stem cells. *Cells.* 2020;9(2):497.
- Plis SM, et al. Deep learning for neuroimaging: a validation study. *Front Neurosci.* 2014;8:92071.
- Faisal MAA, et al. NDDNet: a deep learning model for predicting neurodegenerative diseases from gait pattern. *Appl Intell.* 2023;53(17):20034–46.
- Jaganathan K, et al. Predicting splicing from primary sequence with deep learning. *Cell.* 2019;176(3):535–48. e24.
- Subramaniam S, Shahani N. *Ribosome Profiling Reveals a Dichotomy Between Ribosome Occupancy of Nuclear-Encoded and Mitochondrial-Encoded OXPHOS mRNA Transcripts in a Striatal Cell Model of Huntington Disease.* bioRxiv, 2021: p. 2021.01. 30.428960.
- Dagar S, et al. Ribosome profiling and Mass Spectrometry reveal widespread mitochondrial translation defects in a Striatal Cell Model of Huntington Disease. *Molecular & Cellular Proteomics.* 2024. p. 100746.
- Biasiotto G, et al. Iron and neurodegeneration: is ferritinophagy the link? *Mol Neurobiol.* 2016;53:5542–74.

56. Davidson PL, et al. Recent reconfiguration of an ancient developmental gene regulatory network in *Heliocidaris* sea urchins. *Nat Ecol Evol.* 2022;6(12):1907–20.

Publisher's note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.