

PROTECT: Protein circadian time prediction using unsupervised learning

Aram Ansary Ogholbake, Qiang Cheng*

Institute for Biomedical Informatics, Department of Internal Medicine and Department of Computer Science, University of Kentucky, Lexington, KY, US

Abstract

Circadian rhythms regulate the physiology and behavior of humans and animals. Despite advancements in understanding these rhythms and predicting circadian phases at the transcriptional level, predicting circadian phases from proteomic data remains elusive. This challenge is largely due to the scarcity of time labels in proteomic datasets, which are often characterized by small sample sizes, high dimensionality, and significant noise. Furthermore, existing methods for predicting circadian phases from transcriptomic data typically rely on prior knowledge of known rhythmic genes, making them unsuitable for proteomic datasets. To address this gap, we developed a novel computational method using unsupervised deep learning techniques to predict circadian sample phases from proteomic data without requiring time labels or prior knowledge of proteins or genes. Our model involves a two-stage training process optimized for robust circadian phase prediction: an initial greedy one-layer-at-a-time pre-training which generates informative initial parameters followed by fine-tuning. During fine-tuning, a specialized loss function guides the model to align protein expression levels with circadian patterns, enabling it to accurately capture the underlying rhythmic structure within the data. We tested our method on both time-labeled and unlabeled proteomic data. For labeled data, we compared our predictions to the known time labels, achieving high accuracy, while for unlabeled human datasets, including postmortem brain regions and urine samples, we explored circadian disruptions. Notably, our analysis identified disruptions in rhythmic proteins between Alzheimer’s disease and control subjects across these

*Corresponding author. Email: qiang.cheng@uky.edu

samples.

Keywords: Circadian Rhythms, Proteomic Data, Alzheimer’s Disease

1. Introduction

Circadian rhythms, driven by an internal molecular clock, regulate physiological and behavioral processes. Disruptions to these rhythms have been associated with various pathologies, including type 2 diabetes, obesity, and neurological disorders like Alzheimer’s disease (AD) [1, 2, 3, 4, 5, 6]. Hence, studying these rhythms is crucial due to their significant impact on health and disease.

In mammals, cell-autonomous circadian rhythms are driven by interconnected transcriptional-translational feedback loops. Additionally, substantial contributions to circadian processes come from posttranslational and post-transcriptional regulation [7, 8]. While the majority of research in the field of chronobiology has focused on the rhythmic expression at the mRNA level, understanding circadian rhythmicity at the protein level remains limited. This limitation is noteworthy given the acknowledged contribution of post-transcriptional mechanisms to circadian rhythms at the protein level [9, 10].

Obtaining precise sample times is currently essential for analyzing circadian rhythms in proteomic data. However, many datasets, particularly human samples, often lack labeled timestamps due to constraints imposed by health risks in collection protocols or the inability to collect precise sample times (e.g., time of death). This lack of time labels poses an urgent need for computational methods to predict the phase of each sample. Challenges arise in developing such methods due to several factors. Firstly, these datasets often contain small sample sizes with high dimensionality and significant noise. Moreover, the presence of proteins with periods shorter than 24 hours, i.e., ultradian rhythms, adds complexity. These challenges cannot be effectively tackled using conventional statistical or machine learning methods. The rapid advances in deep learning techniques offer promise for overcoming these challenges and estimating the circadian phase of each sample. While efforts have been made at the mRNA level [11, 12, 13, 14, 15, 16, 17], there is a notable gap in research concerning the prediction of circadian phases in proteomic data.

As mentioned earlier, many datasets lack labeled timestamps. Zeitzeiger [14], TimeSignature [12], TimeTeller [15], tauFisher [16], and PLSR [17] employ supervised learning approaches that require time labels and, therefore, cannot be used in practice when time information is unavailable. In contrast, we introduce a novel unsupervised learning approach which eliminates the need for time labeled samples to estimate circadian phases. This makes it a valuable alternative for circadian rhythm analysis, particularly in scenarios where accurate time annotations are lacking.

On the other hand, existing unsupervised or mathematical sample phase estimation (or temporal ordering) methods for gene expression data, such as CYCLOPS [11] and CIRCUST [18], were specifically developed for gene expression datasets¹. These methods explicitly require the use of expressions of a set of “seed rhythmic genes” to operate. The seed genes are set of known circadian rhythmic genes, including core clock genes, in animal tissues. Without these pre-selected seed rhythmic genes, these phase estimation methods may not function properly, limiting their applicability to datasets where such prior knowledge is available.

Proteomic data poses a significant challenge for applying these methods to study rhythmic proteins, as many proteins corresponding to the designated seed rhythmic genes may not be expressed in a tissue. This is because proteomic datasets typically measure the expression levels of 1,000 – 10,000 proteins, which is significantly fewer than the 15,000 - 50,000 genes typically measured in gene expression datasets. Furthermore, proteins corresponding to core clock genes are often expressed at low levels or not at all, making their measurements unavailable in proteomic datasets. Additionally, some genes known to be rhythmic at the mRNA level may not exhibit rhythmicity in their corresponding proteins, and vice versa. Previous studies have shown that only a small proportion of rhythmic proteins are rhythmic in their corresponding genes, while a large number of proteins exhibit rhythmicity even when their corresponding mRNAs do not, and vice versa [20].

To meet the urgent need and challenges, we introduce an unsupervised deep learning approach called PROTECT (PROTEin Circadian Time prediction). PROTECT is a rhythmicity-aware model designed to predict the circadian phase of each sample in proteomic data, without requiring time labels or prior knowledge of rhythmic markers. This method can effectively

¹We also tested a program of unpublished ESOCVD [19], which could not execute.

handle small sample size datasets and ultradian proteins. It does not require any known seed rhythmic proteins or genes and can handle considerable noise levels present in proteomic data. We demonstrate the efficacy, accuracy, and robustness of our approach using mouse, *Ostreococcus tauri* cell and human datasets, where time labels are available. Subsequently, we investigate circadian rhythms in human datasets obtained from different brain regions and urine samples from both control and AD subjects.

Proteins, along with the metabolic pathways they modulate, often serve as the ultimate biological effectors of AD genetic [21]. Despite the increasing body of work in recent years on discerning disparities in proteomic data between AD and control subjects [22, 23, 24, 25, 26], there is a lack of investigation in identifying differences between control and AD individuals based on circadian rhythms. Our developed approach can uncover AD-associated rhythmic proteins and distinguish differences in rhythmic patterns between control and AD subjects, filling this gap.

In brief, the novelty of our paper includes but is not limited to:

- We present a unique approach for accurately predicting circadian phases in un-labeled proteomic datasets, which, to our knowledge, no existing methods in the field address. It does not rely on prior information about circadian rhythmic genes or proteins, marking a significant advancement in the field of computational inference from proteomic data.
- Our method shows high prediction accuracy while effectively handling datasets with varying sample sizes, high dimensionality, and substantial noise. Additionally, it identifies ultradian rhythmic proteins, showcasing its versatility.
- Applying our method to un-labeled human proteomic datasets reveals circadian rhythmic differences between control and AD subjects across three postmortem brain regions and urine samples. This highlights the potential of our approach to uncover novel insights into circadian disruptions associated with AD in proteomic datasets.

2. Proposed PROTECT model and algorithm

We developed PROTECT, an unsupervised learning method, to predict the time of high-dimensional proteomic samples based on the data itself, without relying on any a priori information or time labels. Proteomic data

includes both rhythmic and non-rhythmic proteins. Rhythmic proteins typically exhibit periodic patterns, with values peaking at certain times of the day and dipping at others, while non-rhythmic proteins have no periodic patterns. Moreover, the peak times among rhythmic proteins are different. Without known sample times, rhythmic patterns are obscured, and PROTECT aims to recover them. PROTECT addresses key challenges in studying circadian rhythms using proteomic data, including small sample sizes, the presence of ultradian proteins, and limited knowledge of rhythmic proteins, especially in human datasets.

PROTECT utilizes a deep neural network (DNN) to predict the phase of each sample in proteomic data using a greedy layer-wise technique, inspired by Hinton et al.’s work [27]. The DNN architecture in PROTECT is well-suited for high-dimensional data. It consists of an input layer, multiple hidden layers, and an output layer with two neurons representing a single angular phase. Training the network involves pre-training the DNN and obtaining the weights of the corresponding DNN neurons, followed by fine-tuning the weights to predict sample phases by regressing with proteins’ cosine models.

Greedy layer-wise reconstruction has advantages shown in [27], which we adapt for pre-training our DNN architecture. The optimized weights obtained in the pre-training stage provide an effective initialization, facilitating effective optimization for learning the sample phases and fitting downstream cosine models in the fine-tuning stage. This approach is particularly effective for small datasets because of its robust initialization of the network’s weights, which can be crucial when data is limited. Additionally, by training each layer independently, the network captures features at different levels of abstraction, reducing the risk of overfitting that often accompanies small sample sizes. The hierarchical representation further augments the network’s capability to capture intricate patterns within the high-dimensional proteomic data.

PROTECT’s methodology is structured into three main steps, as depicted in Figure 1: Normalizing data, Greedy layer-wise pre-training, and Fine-tuning.

Normalizing data: In the first step, the proteomic data, which includes m samples with n protein measurements per sample, undergoes z-score normalization. This involves subtracting the mean and dividing by the standard deviation for each protein across all samples. This standardization ensures that all proteins are on a comparable scale, improving the stability and per-




Figure 1: Illustration of the overall diagram of PROTECT.

formance during the training process. The resulting normalized data is then utilized in the subsequent step: Greedy layer-wise pre-training.

For datasets with a very large number of features (e.g., $n \geq 5000$), we recommend applying a feature selection or dimensionality reduction technique after normalizing the data to enhance computational efficiency and model performance. One effective approach is to use k-means clustering, where the features are grouped into clusters, and only the clusters with the highest variance are kept. Alternatively, high-variance features can be directly selected, where only the features with the highest variance across samples are chosen.

Greedy layer-wise pre-training: This step involves greedily training each layer of the DNN using a separate shallow auto-encoder (AE) network suitable for datasets with small sample sizes. The primary objective is to encode the input into lower dimensions and capture intricate patterns in the proteomic data using each AE, ultimately reaching the output layer where the features represent the encoded data in two dimensions (see an example of encoded data in supplementary Figure S1).

Each shallow auto-encoder consists of an input layer, a hidden layer, and an output layer. The hidden layer’s output from the AE at layer l of the DNN becomes the input for the AE at layer $l + 1$ in the DNN. The objective function for each AE is the mean squared error (MSE) between the input and output, providing a measure of how well it reconstructs the input.

After training the last AE, the values of the two nodes in its hidden layer, denoted as s_i and c_i for sample i , are extracted. These feature values compute the initial phase, ϕ_i^0 , for each sample i , as expressed by the equation:

$$\phi_i^0 = \arctan \left(\frac{s_i}{c_i} \right). \quad (1)$$

Subsequently, CosinorPy is employed to extract geometric information for each protein using the initial predicted phases (ϕ_i^0). This includes parameters such as amplitude (A_p^0), mesor (L_p^0), and acrophase (ϕ_p^0), achieved by fitting each protein p to a cosine curve using the initial predicted sample phases. This geometric information is then incorporated into the objective function of the DNN in the next step, which is fine-tuning.

In summary, the pre-training process begins with training the first AE, where normalized proteomic data serves as the input. From this, we extract the hidden layer values. Then, we proceed to train the second AE, utilizing the hidden layer output of the first AE as its input, and extract the hidden

values. This iterative process continues until we extract the features of the last AE, which yields the initial angular phase. The initial phase is used to calculate the initial geometrical information of each protein which will be used in fine-tuning. Moreover, the extracted features from each AE are then utilized to initialize the weights of the DNN in fine-tuning.

Fine-tuning: After pre-training the DNN using sequence of shallow AEs and obtaining the initial weights and geometrical information, we fine-tune the network to predict sample phases. In this stage, the network learns to align protein expressions with rhythmicity by fitting each protein to a cosine function. We organize the data into a table where each row corresponds to a sample, and each column represents a proteins’s expression; that is, x_{ip} represents the original observation for sample i and protein p . We use a parameterized cosine function to fit the observations, given by:

$$\hat{x}_{ip} = L_p + A_p \cos(\omega_p \hat{\phi}_i + \phi_p), \quad (2)$$

where $\hat{\phi}_i$ represents the sample phase, and L_p , A_p , and ϕ_p are the mesor, amplitude, and the acrophase of protein p , respectively. $2\pi/\omega_p$ represents the period of protein p . The objective function is defined as:

$$\mathfrak{L} = \frac{1}{m} \sum_{i=1}^m \frac{1}{n} \sum_{p=1}^n \|x_{ip} - \hat{x}_{ip}\|_q^q + \lambda R(\Theta). \quad (3)$$

Here, $\|\cdot\|_q$ is an l_q norm with respect to both i and p , with q as a positive value. Θ represents the set of all relevant parameters (including L_p , A_p , $\hat{\phi}_i$ and ϕ_p), and $R(\Theta)$ represents a regularization function on the parameters. λ is a non-negative hyperparameter balancing the fitting error and regularization.

The learnable parameters L_p , A_p , and ϕ_p are initialized with the values calculated during the pre-training step. ω_p allows the model to fit proteins with different rhythmicity such as circadian rhythms and ultradian rhythms. This flexibility ensures that the model accurately captures various rhythmic patterns in protein expression.

This objective function aims to find the optimal phase for each sample and determines the best geometrical parameters, ensuring an accurate and effective fitting of the protein data to the cosine curve. Moreover, by incorporating the unique geometric information of each protein in each sample, this objective function helps with handling small sample sizes (to see an example of a random protein recovering its rhythmicity during this stage, see supplementary Figure S2).

We have explored the use of multiple regularization functions including l_1 norm, l_2 norm and total variation (TV) regularization. The total variation term aims to reduce the abrupt shifts between consecutive sample phases and is defined as:

$$\sum_{k=2}^m |\hat{\phi}_{i_k} - \hat{\phi}_{i_{(k-1)}}|, \quad (4)$$

where sample phases are sorted in ascending order, and i_k represents the k -th sorted sample for $k = 1, \dots, m$.

Our experiments with various λ values for different regularization functions on labeled datasets showed no significant improvement in phase prediction accuracy when $\lambda > 0$. Therefore, we set the $\lambda = 0$ in our final model. Moreover, we employed $q = 1$ for the fitting error term.

This parametric objective function enables the predicted phases to take into account the protein geometrical information, such as amplitudes, acrophases, and periods. Moreover, it is applicable when the noise in the data is non-Gaussian [28, 29]. In addition, cosine curve fitting has been shown to be effective on data without replicates, containing outliers, irregularly spaced time intervals, and unbalanced data distributions where more samples are collected at certain times of the day [29, 30, 31]. To see the performance of the objective function, refer to the convergence results provided in supplementary Figures S3, S4, and S5).

2.1. Screening for potential outliers for samples and proteins

In the fine-tuning phase, we design our model to handle inherent noise and potential outliers in proteomic data. The process of detecting outliers is as follows:

- **Sample-level outlier detection:** For each sample, we calculate the averaged fitting error using:

$$E_i = \frac{1}{n} \sum_{p=1}^n (x_{ip} - \hat{x}_{ip}), \quad i = 1, \dots, m. \quad (5)$$

Thus, E_i represents the fitting residues averaged for all proteins $p = 1, \dots, n$. We then compute the mean and standard deviation of these E_i values across all samples, respectively denoted by m_s and σ_s . If a sample's deviation exceeds two standard deviations from the mean m_s , i.e., if $|E_i - m_s| > 2\sigma_s$, then sample i is considered an outlier.

- **Protein-level outlier detection:** For each protein, we obtain its average fitting residual over all samples:

$$D_p = \frac{1}{m} \sum_{i=1}^m (x_{ip} - \hat{x}_{ip}), \quad p = 1, \dots, n. \quad (6)$$

Thus, D_p represents the fitting quality averaged for all samples $i = 1, \dots, m$. We compute the mean m_p and standard deviation σ_p of D_p across all proteins. If a protein’s deviation exceeds two standard deviations from the mean m_p , i.e., if $|D_p - m_p| > 2\sigma_p$, then protein p is considered an outlier candidate. To ensure no significant cyclic proteins are removed, the outlier candidates undergo a second screening process. In this screening, if the predicted amplitude of a candidate falls below the 75th percentile of all predicted amplitudes, it is classified as an outlier protein.

Through this approach, we identify outliers (resp. outlier proteins) that significantly deviate from the majority of samples (resp. proteins). Subsequently, we remove the outlier samples and proteins to retrain the model. This process minimizes the influence of outlier or corrupted proteins and samples that have unusually large residuals.

3. Results

3.1. Comprehensive Evaluation of PROTECT

To assess the efficacy of PROTECT, we performed computational experiments on multiple public datasets. We verified the accuracy of our method using labeled proteomic datasets from human, mouse and cell models. Subsequently, we conducted experiments on unlabeled human datasets, including postmortem brains and urine. Our investigation on unlabeled human datasets focused on comparing differences between control and AD subjects. This comprehensive, multifaceted approach enabled us to evaluate PROTECT’s capabilities on different species and varying sample sizes datasets.

3.2. Hyperparameters and Model Details

Our model was trained using PyTorch Lightning, where the best results were obtained using 5 hidden layers. The neuron counts of these layers ranged from $2^{\lfloor \log_2(f) \rfloor}$ to 2, where f is the number of features in each dataset. During

pre-training, optimizers such as SGD, Adam, and learning rate-free methods (DAdapt-SGD and DAdapt-Adam [32]) were tested. We found that the specific optimizer did not notably impact the end results, allowing flexibility in optimizer choice for each layer. During pre-training, the first 5 AEs were trained using 7 epochs and the last AE was trained for 20 epochs. During fine-tuning the whole network was trained for 20 epochs. To handle noise, we retrained the network for 200 epochs using the datasets with outliers removed. Check table 1 for some details on the implementation.

Training Stage	Optimizer	Best Learning Rate	Momentum	Weight Initialization
Pre-training (1st-5th AE)	Adam SGD DAdapt-SGD	Adam: 0.001 SGD: 0.1 DAdapt-SGD: -	SGD: 0.85	Xavier Uniform
Pre-training (last AE)	DAdapt-SGD	learning rate-free	-	Xavier Uniform
Fine-tuning	DAdapt-SGD	learning rate-free	-	Transferred Weights

Table 1: Details for pre-training and fine-tuning stages.

3.3. Labeled proteomic data experiments

3.3.1. Data description

Several proteomic datasets were utilized, sourced from mouse hip articular cartilage (PXD019431) [33], two different mouse liver datasets [20, 34], *Ostreococcus tauri* cells [35] and human plasma [36]. Mouse hip articular cartilage was collected every four hours over two days, sampling 6 animals at each time point. We calculated the average of label-free quantification (LFQ) intensity across all six animals for each time point, resulting in a set of 12 samples. Both mouse liver datasets consist of 16 samples obtained from mice at 3-hour intervals over a 2-day span. These data quantify the relative protein abundance in each of the 16 samples against a common reference sample, which was labeled using the SILAC method. The dataset mentioned in Wang et al. (2017) [34] contains about 1000 more proteins than the one referenced in Mauvoisin et al. (2014) [20]. The *Ostreococcus tauri* cell dataset provides the mean of normalized abundance per time point, encompassing 6 samples. The human plasma data consists of samples from 6 healthy young males at




Figure 2: Accuracy of PROTECT on (a) *Ostreococcus tauri*, (b) Mouse hip articular cartilage, (c) Mouse liver, and (d) Human plasma. The top row shows ROC curves where the y-axis shows the fraction of correctly predicted samples, and the x-axis shows the size of errors. The bottom row shows the scatter plots of predictions vs ground truth.

different time points over two days. We utilized the time points with more than one subject and averaged them. This resulted in time points of 1, 5, 9, 13, 15, 17, and 21 for one day, and 1, 5, 9, 13, 17, and 21 for the second day.

3.3.2. Performance evaluation

We first demonstrate the time of day prediction accuracy of PROTECT on time-labeled datasets, employing ROC curves that show the fraction of correctly predicted sample phases relative to the size of errors [12]. Moreover, we show scatter plots of predicted sample phases with respect to the ground truths.

PROTECT shows excellent performance on various datasets, achieving a high normalized area under the curve (nAUC) of 94% on *Ostreococcus tauri* cell data (Figure 2(a)), and over 80% on mouse hip articular cartilage (Figure 2(b)), mouse liver [20] (Figure 2(c)), and human plasma (Figure 2(d)) proteomic datasets. The bottom row of Figure 2 demonstrates accurate phase predictions compared to ground truth across all datasets. In particular, all samples of the *Ostreococcus tauri* cell data show minimal phase prediction errors, while samples in other datasets typically deviate by no more than 4 hours (60 degrees in circadian phase).

We further evaluated the performance of PROTECT on an additional




Figure 3: Comparison of CYCLOPS and our method on mouse liver dataset. (a) CYCLOPS ROC curve and predicted sample phases vs ground truths without using proteins corresponding to seed genes. (b) CYCLOPS ROC curve and predicted sample phases vs ground truths after using proteins corresponding to seed genes. (c) our ROC curve and predicted sample phases vs ground truths.

mouse liver dataset [10] and mouse brown adipose tissue (BAT) tissue [37] with the corresponding results presented in the supplementary material (Figure S6). The results further demonstrate PROTECT’s ability to accurately predict sample phases.

We tested existing methods, originally designed for gene expression data, on the proteomic mouse liver dataset (PXD003818) [34], which includes 16 samples and 5,301 proteins. This dataset was used because it contains the most proteins corresponding to seed rhythmic genes compared to other datasets. Figure 3(a) shows the ROC curve and the sample phase estimations compared to the time labels using CYCLOPS [11] without utilizing proteins corresponding to the seed rhythmic genes. This plot shows that without using proteins corresponding to seed genes, the prediction is close to random guessing with an nAUC of 59%. Moreover, correct estimates should ideally be on the dotted diagonal gray lines (where lines in the corners account for




Figure 4: Plots of four core clock proteins in mouse liver using predicted phases by PROTECT. The y-axis represents protein expression levels, and the x-axis represents the predicted phases (in degrees) as determined by PROTECT.

circadian periodicity). However, most samples are predicted to have similar values. Using the available 1,076 proteins corresponding to the CYCLOPS-designated 8,504 seed rhythmic genes, the ROC curve and scatter plot in Figure 3(b) show that most sample phase predictions still have similar values, with the nAUC improving by only 5%, indicating CYCLOPS’s failure on proteomic data. CIRCUST failed to execute due to the lack of core clock genes for their “synchronization procedure,” as pointed out by a CIRCUST author via GitHub discussions. In contrast, our proposed approach (PROTECT) in Figure 3(c) shows high accuracy, with a high nAUC of 90% on this mouse liver dataset.

Moreover, we used this mouse liver data to evaluate the robustness of PROTECT when working with fewer samples. We randomly subsampled the data by removing 3, 6, 9, and 12 samples. This process was repeated 4 times for each subsampling scenario. PROTECT was then applied to predict the sample phases, and the average results across the 4 repetitions are shown in Figure S7. While the AUC of predictions decreases slightly compared to using the full dataset, it remains consistently above 80% across all sampling levels. This robustness highlights PROTECT’s ability to perform reliably,




Figure 5: Plots of four randomly chosen proteins known to be strongly regulated by circadian cycle in human plasma using predicted phases by PROTECT. The y-axis represents protein expression levels, and the x-axis represents the predicted phases (in degrees) as determined by PROTECT.

even when part of the data is used, underscoring its utility for sparse high-throughput proteomics data.

3.3.3. Circadian rhythmicity in known rhythmic proteins

In Figure 4, we demonstrate the results of plotting four of the core clock proteins' values versus the predicted sample phases using PROTECT across the mouse liver dataset [34] where the core clock proteins are available. These results highlight the rhythmicity of these proteins using PROTECT's predicted sample phases. Figure 5 also depicts four proteins, randomly selected from a set of proteins known to be strongly regulated by the circadian cycle in human plasma [36]. Our predicted sample phases clearly reveal the circadian rhythmic patterns in these proteins as well. Additionally, we evaluated PROTECT's performance on human thyroid tissue [38]. To validate the predicted phases, we examined some known rhythmic proteins in thyroid tissue [30]. Our results (supplementary Figure S8) confirmed the rhythmicity of these proteins, supported by significant p-values, further demonstrating PROTECT's reliability in detecting circadian patterns in human samples.

3.3.4. Outlier handling

Here, we show the counts of proteins and samples outliers, as defined in Section 2.1, on our time-labeled datasets. Our algorithm found no outlier samples, suggesting that all residuals for samples are similarly distributed. However, each dataset exhibited a number of outlier proteins, as shown in Table 2. We excluded the outliers from each dataset and retrained the network to handle these protein outliers. We then fitted Cosinor curves against the ground truth and verified that these proteins did not exhibit significant rhythmicity at a p-value threshold of $5e-2$.

Dataset	Number of outliers
Mouse liver (Wang et al.)	0
Mouse hip articular cartilage	13
Mouse liver (Mauvoisin et al.)	42
Ostreococcus tauri	15
Human plasma	21

Table 2: Number of protein outliers detected in each analyzed dataset.

3.4. Un-labeled proteomic data experiments

3.4.1. Data description

In our exploration of circadian rhythms in unlabeled human proteomic data, we examined three brain regions: the Temporal Cortex (TC) [39], the Parietal Association Cortex [40], and the Dorsolateral Prefrontal Cortex (DLPFC) [41, 39]. Additionally, we incorporated a dataset derived from urine samples [42]. In the TC and DLPFC datasets, the LFQ intensity is used for protein quantitation. In the parietal association cortex dataset, the calculation of protein intensities involves summing the TMT reporter ions corresponding to all peptides assigned to each protein. In the urine dataset, protein quantification is conducted utilizing the intensity-based absolute quantification (iBAQ) algorithm.

3.4.2. Data preparation for control and AD subjects comparison

In preparing data for comparing control and AD subjects in unlabeled brain datasets, we took several steps. Initially, we removed proteins with missing values from both control and AD datasets. To ensure a fair comparison, we then used the common set of proteins in both groups. Moreover, to

handle different sample sizes in control and AD groups, particularly in TC and DLPFC datasets, we made them equal by matching age and sex characteristics. In the urine dataset with numerous missing values, we removed proteins with missing values in over 50% of samples, following a similar approach as in the original paper [42]. Then, we found the intersection of all proteins in control, MCI, and AD subjects.

3.4.3. Disparities between AD and control subjects

To find rhythmic proteins for control and AD groups, we used stringent criteria which involved utilizing the Benjamini-Hochberg False Discovery Rate (FDR), relative amplitude (rAmp) (amplitude divided by the baseline), and coefficient of determination (R^2). We utilized CosinorPy [29] to fit each protein with a cosine curve, employing our predicted phases to determine these values for each protein. Subsequently, we applied predetermined thresholds for significance as follows: $\text{FDR} < 0.05$, $\text{rAmp} \geq 0.1$, and $R^2 \geq 0.1$. Proteins that met all these criteria were considered rhythmic.

Figures 6, S9, and S10 illustrate the disparities between control and AD subjects in the TC, parietal association cortex, and DLPFC brain regions. First, Venn diagrams are provided to visualize the overlap of rhythmic proteins between control and AD in each brain region (Figures 6(a), S9(a), and S10(a)). The results indicate a close similarity in the number of rhythmic proteins between AD and control subjects within each brain region, with a higher count in control subjects. In the TC and DLPFC regions, large proportions of proteins exhibit rhythmicity. In the TC region, out of 2425 proteins, 1718 and 1707 proteins demonstrate rhythmic patterns in control and AD, respectively and 80% of rhythmic proteins identified in control subjects maintain their rhythmic nature in AD. Also, in the DLPFC region, out of 2483 proteins, 1980 and 1759 proteins display rhythmicity in control and AD, respectively, with 1610 proteins shared between the two groups. However, in the parietal association cortex lower proportions of proteins show rhythmicity. Specifically, out of 3389 proteins, 717 exhibit rhythmicity in control and 692 proteins are rhythmic in AD and 255 rhythmic proteins are shared between the two groups.

In Figures 6(b), S9(b), and S10(b), the distribution of peak times of rhythmic proteins in control and AD subjects is depicted within 24 hours (360 degrees) of circadian cycle. To ensure consistency across datasets, we used EHD1 - a protein that exhibited rhythmicity in all three brain regions and urine samples on different groups (we used Venn diagrams to find overlapping




Figure 6: Disparities between control and AD subjects in temporal cortex using PROTECT predicted phases. (a) Venn diagram of numbers of rhythmic proteins in control (blue) and AD (pink) subjects. (b) Rose plots of distributions of peak times (i.e., acrophases) in rhythmic proteins of control and AD subjects within 24 hours (360 degrees) of the circadian cycle. The radial distance indicates protein counts. The max of radial distance differs between control and AD plots. (c) Plots of 5 example rhythmic proteins in control subjects that lose rhythmicity in AD.

proteins between all different datasets, similar to Figure 6(a)) - as a reference point. We set the phase of EHD1 to zero and shifted the acrophases of other proteins relative to this one, thereby providing a common baseline for comparing peak times across regions and conditions. Dispersed protein peak distributions arise in the parietal association cortex unlike the TC and DLPFC. In the latter two regions, the peak times of rhythmic proteins are more concentrated in smaller areas and many proteins peak around same times. Additionally, peak times show rotational shifts between control and AD subjects in both the parietal association cortex and TC. In the TC region, peak times for control subjects are mostly between 90-135 degrees, whereas in AD subjects, they are concentrated between 0-45 degrees. In control subjects within the parietal association cortex region, the majority of peak times occur between 0-45 and 180-225 degrees, whereas in AD subjects, they are observed between 135-180 and 315-360 degrees.

We then illustrate five top proteins that demonstrate significant rhythmicity in control subjects but lose rhythmicity in AD subjects, based on the three criteria mentioned earlier, for each brain region. This is depicted in Figures 6(c), S9(c), and S10(c) on TC, parietal association cortex, and DLPFC regions, respectively. Importantly, these proteins not only show significantly increased p-values in AD subjects but also undergo a noticeable reduction in amplitude.

These findings highlight the differences in the rhythmic protein profiles between control and AD subjects across distinct brain regions. Overall, while the TC and DLPFC show large overlapping circadian proteomic signatures between disease and control states, the parietal association cortex exhibits far fewer shared rhythmic proteins. The variability indicates regionally specific circadian disruptions arise in AD. Identifying the underlying mechanisms driving these cortical differences in rhythmicity may open new therapeutic avenues for future exploration.

We also investigated the number of rhythmic proteins shared across all three brain regions, categorizing them into four groups: proteins rhythmic in all control subjects across different brain regions, proteins rhythmic in all AD subjects across different brain regions, proteins shared in all controls but not in ADs, and proteins shared in all ADs but not in controls. Table S1 summarizes these findings, revealing 276 proteins exhibits rhythmicity in all control subjects, 271 rhythmic proteins common to all AD subjects, and 4 proteins rhythmic in all controls but not in ADs. No proteins found to be rhythmic in all ADs but not in controls.

3.4.4. *Enrichment analysis*

We conducted Gene Ontology (GO) enrichment analysis using GSEAPy [43], which serves as an interface between Python and Enrichr web services [44, 45, 46]. This analysis was performed on rhythmic proteins in the control group that lose rhythmicity in AD (control-specific rhythmic proteins), as well as rhythmic proteins in AD that are not rhythmic in the control group (AD-specific rhythmic proteins), for each brain region (Figures 7, S12, and S14).

Observing the enriched biological processes in TC region (Figure 7), both control-specific and AD-specific rhythmic proteins are enriched with respiration and energy metabolism. For control-specific rhythmic proteins, the most significant process is aerobic respiration, whereas for AD-specific rhythmic proteins it is dicarboxylic acid catabolic process that is also related to energy metabolism. As intermediates, in the tricarboxylic acid cycle (TCA; also called citric acid cycle or Krebs cycle) such as succinate, fumarate, and malate, dicarboxylic acids are key substrates for cellular respiration. Their oxidation provides a critical source of electrons in the electron transport chain for oxidative phosphorylation and subsequent ATP synthesis through this integrated network of pathways.

The enrichment in AD-specific rhythmic proteins has distinctive processes such as response to endoplasmic reticulum stress, glutamate catabolic process and energy derivation by oxidation of organic compounds, which are not enriched in the control-specific rhythmic proteins. The response process is related to the activation of pathways managing stress affecting the endoplasmic reticulum (ER) in cells. Such stress is triggered by the accumulation of misfolded proteins in the ER lumen. The ER stress response involves signaling pathways that work to restore proper protein folding in the ER, manage the protein load, and degrade misfolded proteins. The unfolded protein response is one of the main pathways managing the cellular response to ER stress. It acts to restore ER homeostasis. If ER stress is prolonged, apoptotic cell death pathways may be activated. This enrichment aligns with literature on misfolded proteins in AD [47].

The other processes enriched in AD-specific rhythmic proteins processes are linked to the energy metabolic alterations in AD. For example, glutamate catabolism connects tightly to multiple pathways related to fueling respiration, driving aerobic/anaerobic energy production, and contributing to nitrogen balance. Glutaminolysis pathway metabolically breaks down gluta-

mate for substrates like lactate that is important for energy homeostasis, e.g., in cancer cells relying more on anaerobic glycolysis pathways. Our findings align with the results from the literature [48], which point out that circadian affects energy homeostasis and has significant effects on AD.

It is important to note that there is often an inherent biases in functional enrichment analysis. Our primary objective in using GSEA is not to establish definitive biological proof but rather to explore potential functional relevance of the rhythmic proteins. In this regard, the pathway analysis results should be viewed as a complementary and exploratory perspective to our modeling process, offering additional insights into the biological context of rhythmic proteins.

3.4.5. Disease associated with control and AD specific rhythmic proteins

We used DisGeNET [49] library to find diseases associated with control-specific and AD-specific rhythmic proteins (see Figures S11(a), S13(a), and S15(a)). Schizophrenia, AD, and neurodegenerative disorders are shared among all rhythmic proteins in control, which lose rhythmicity in AD in all brain regions. This shows that disrupted proteins are highly associated with AD, as well as schizophrenia, which has a high correlation with AD [50]. We also analyzed these proteins using additional annotation libraries (see Figures S11(b), S13(b), and S15(b)).

3.4.6. Hub proteins and related drugs

In addition, we present the top 20 hub proteins along with their directly connected first-stage nodes in the proteomic data of control-specific rhythmic proteins within the TC (Figure 8(a)). Similar analyses for the parietal association cortex, the DLPFC, and urine datasets are provided in Table S2. To achieve this, we initially constructed a protein coexpression network using the WGCNA R package [51], focusing on proteins that display rhythmicity in the control group but lose rhythmicity in AD. We applied the cytoHubba Cytoscape plugin [52, 53] and utilized the degree algorithm to identify and highlight the top-ranked proteins based on their connectivity with other proteins. Through the use of DrugBank [54], we discovered that six of these hub proteins represent potential drug targets. The corresponding drugs for each target are shown in Figure 8(b), categorized based on their development status as approved, investigational, or experimental.





Figure 7: Gene Ontology enrichment analysis of proteins rhythmic in control subjects but not in AD (blue), and in AD subjects but not in control (pink) on temporal cortex.



(a)

Hub Protein	Related Drug(s)
LTA4H	Captopril [A], Ubenimex [I], DG051 [I], Tosedostat [I], Imidazole [E&I]
SRSF1	Copper [A&I]
HNRNPK	Artenimol [A&E&I], Phenethyl Isothiocyanate [I]
CYB5A	Chromium[A], Chromic nitrate [A], Chromium gluconate [A], Chromous sulfate [A], Chromium nicotinate [A&E]
CALB2	Calcium Phosphate [A], Calcium phosphate dihydrate [A], Calcium citrate [A&I], Calcium levulinate [A&E]
APEX1	Lucanthone [I]

(b)

Figure 8: Hub proteins in temporal cortex. (a) Network depiction of the first twenty hub proteins based on degrees and proteins directly interacting with these hub proteins. The colorbar indicates the connection degrees of the hub proteins. (b) Hub proteins and related drugs. The status of each drug is shown in the brackets next to it: A, I, and E indicate that the drug is approved, investigational, and experimental, respectively.

4. Discussion

To study circadian rhythms at the protein level, sample times are needed. However, proteomic datasets, especially in humans, often lack explicit time labels. Other challenges include: the sample sizes are typically very small; ultradian proteins typically exist; and knowledge of rhythmic proteins are limited. To address these issues, PROTECT has been designed to predict sample phases in proteomic data in an unsupervised manner. This method is appropriate for different sample sizes and does not require a priori information. The effectiveness of PROTECT was validated through testing on time-labeled datasets, achieving remarkable results with over 80% nAUC on all instances. Our exploration of un-labeled human datasets using PROTECT’s time predictions revealed differences between control and AD subjects in three brain regions and urine samples. Additionally, the study provides potential drug targets and identifies some ultradian rhythmic proteins in TC.

PROTECT is generalizable to other circadian omics data types, such as transcriptomic and metabolomic datasets [20, 55, 56, 57, 58, 59, 60, 61, 62], as these data types share similar periodic structures. Unlike existing methods tailored for transcriptomics, PROTECT does not rely on time information or predefined rhythmic features, such as known rhythmic genes or metabolites, making it broadly applicable to these types of datasets.

To validate this, we applied PROTECT to some publicly available circadian transcriptomic and metabolomic datasets. Specifically, we analyzed transcriptomic data from a mouse liver study [20], which includes both proteomic and transcriptomic measurements, as well as a baboon amygdala dataset [55] and a mouse kidney dataset [56]. Additionally, we evaluated metabolomic data from a mouse liver study [57]. PROTECT successfully predicted sample phases across these datasets, with detailed results provided in the Supplementary Material (Figure S19). These findings further support the robustness of our method and its potential for broader applications in circadian and diurnal research.

4.1. Rhythmic proteins in human dataset

On the TC and DLPFC, large portions of proteins are rhythmic in both control and AD subjects. However, in the parietal association cortex, only about 21% of proteins show rhythmicity patterns. Moreover, in the urine dataset, 458, 434, and 387 out of 555 proteins show rhythmic patterns in control, MCI, and AD subjects, respectively (Figure S16). This indicates

a significant portion of proteins exhibiting rhythmicity across all groups. However, it’s important to note that the urine dataset had numerous missing values and after the removal of proteins with missing values in more than half of the samples, we lost many proteins. Therefore, we cannot determine whether they exhibit rhythmic patterns or not. Also, we observed different distributions of peak times in brain regions between control and AD subjects, which can also be seen in urine datasets (Figure S17).

4.2. Associated diseases

Our findings, focusing on proteins exhibiting rhythmicity in control subjects but losing this pattern in AD within the temporal cortex, identified mitochondrial disease as the primary associated condition. This observation is in line with previous studies [63, 64] emphasizing the involvement of mitochondrial dysfunction in the progression of AD. Notably, our results also reveal schizophrenia as the second most associated disease across all brain regions. A study by [50] supports our findings, demonstrating a robust genetic correlation between AD and schizophrenia.

4.3. Hub proteins and drug targets

In our study of proteins exhibiting rhythmic behavior in control subjects but not in individuals with AD, we present the top 20 hub proteins in the TC. Notably, six of these proteins, namely LTA4H, SRSF1, HNRNPK, CYB5A, CALB2, and APEX1, emerge as potential drug targets. Recent studies have linked some of these proteins to be therapeutic targets for AD. According to Adams et al. (2023) [65], LTA4H has been recognized not only as a therapeutic target but also as a plasma biomarker for cognitive impairment associated with aging and AD. Additionally, SRSF1 and PTPB1 have been implicated in suppressing the formation of a CD33 splicing isoform associated with AD [66]. Furthermore, there is evidence suggesting that CALB2, specifically in hippocampal interneurons, serves as an early target in a transgenic model exhibiting AD-like pathology [67].

4.4. Ultradian proteins

In the TC dataset, we utilized predicted sample times to identify potential ultradian rhythmic proteins. Our selection criteria included $\text{FDR} < 5e - 4$, $\text{rAMP} \geq 0.2$, and $R^2 \geq 0.6$ to select the proteins with a period of 12 hours. Through this analysis, we identified six proteins (ACOT7, CLTA, ELMO1, MDP1, NIT1, SH3GL2) exhibiting these characteristics, as illustrated in Figure S18.

4.5. *Potential limitations of our approach*

Despite its excellent performance, PROTECT requires further exploration. In our study, we did not address the impact of the postmortem interval (PMI) on postmortem human brain datasets. PMI refers to the duration between the time of death and sample collection, which can influence proteomic values. While most samples used in this paper have a PMI of less than 10 hours, thus minimally impacting the study, future work should investigate the effect of PMI on predicted times.

One limitation of PROTECT lies in its strategy for handling missing values within the proteomic dataset. The current practice of excluding proteins with missing values may lead to the exclusion of valuable data that could contribute to a more comprehensive understanding of our investigation. As advances in data collection lead to fewer missing values, we expect that more proteins can be incorporated into rhythmicity analysis.

Also, to determine which protein should exhibit its peak at time 0 and subsequently arrange the remaining proteins in relation to this target protein, we employed EHD1. This protein showed rhythmicity in all datasets for both control and AD subjects. We established the peak value of EHD1 at time 0 and organized the other proteins relative to EHD1. In the future, with more studies on proteins' peak times, we can use some known proteins as the target protein.

Besides potential future enhancements to PROTECT itself, systematically investigating these limitations provides broader understanding of the data dependencies and robustness in circadian proteomic modeling. In turn, with further investigation into these practical data constraints alongside algorithmic advances, PROTECT may significantly influence standardized protocols for circadian omics analysis, thus enabling more consistent and accurate biological findings.

Data availability statement

All data used in this paper are publicly available and cited in the paper.

CRedit authorship contribution statement

Aram Ansary Ogholbake: Conceptualization, Formal Analysis, Investigation, Methodology, Software, Visualization, Writing – original draft
Qiang Cheng: Conceptualization, Funding acquisition, Investigation, Methodology, Supervision, Writing – review & editing

Declaration of competing interest

None Declared

Funding sources

This work was supported by the by the NSF under Grants IIS 2327113 and ITE 2433190; and the NIH under Grants R21AG070909, P30AG072946, and R01HD101508-01.

References

- [1] Y. Leng, E. S. Musiek, K. Hu, F. P. Cappuccio, K. Yaffe, Association between circadian rhythms and neurodegenerative diseases, *The Lancet Neurology* 18 (3) (2019) 307–318.
- [2] R. W. Logan, C. A. McClung, Rhythms of life: circadian disruption and brain disorders across the lifespan, *Nature Reviews Neuroscience* 20 (1) (2019) 49–65.
- [3] D. J. Stenvers, F. A. Scheer, P. Schrauwen, S. E. la Fleur, A. Kalsbeek, Circadian clocks and insulin resistance, *Nature Reviews Endocrinology* 15 (2) (2019) 75–89.
- [4] K. Teeple, P. Rajput, M. Gonzalez, Y. Han-Hallett, E. Fernández-Juricic, T. Casey, High fat diet induces obesity, alters eating pattern and disrupts corticosterone circadian rhythms in female icr mice, *Plos One* 18 (1) (2023) e0279209.
- [5] N. Javeed, A. V. Matveyenko, Circadian etiology of type 2 diabetes mellitus, *Physiology* 33 (2) (2018) 138–150.
- [6] R. Jarrett, H. Keen, Diurnal variation of oral glucose tolerance: a possible pointer to the evolution of diabetes mellitus, *British Medical Journal* 2 (5653) (1969) 341.
- [7] D. Mauvoisin, F. Gachon, Proteomics in circadian biology, *Journal of Molecular Biology* 432 (12) (2020) 3565–3577.
- [8] D. Feng, M. A. Lazar, Clocks, metabolism, and the epigenome, *Molecular Cell* 47 (2) (2012) 158–167.

- [9] S. Kojima, D. L. Shingle, C. B. Green, Post-transcriptional control of circadian rhythms, *Journal of Cell Science* 124 (3) (2011) 311–320.
- [10] M. S. Robles, J. Cox, M. Mann, In-vivo quantitative proteomics reveals a key contribution of post-transcriptional mechanisms to the circadian regulation of liver metabolism, *PLoS genetics* 10 (1) (2014) e1004047.
- [11] R. C. Anafi, L. J. Francey, J. B. Hogenesch, J. Kim, Cyclops reveals human transcriptional rhythms in health and disease, *Proceedings of the National Academy of Sciences* 114 (20) (2017) 5312–5317.
- [12] R. Braun, W. L. Kath, M. Iwanaszko, E. Kula-Eversole, S. M. Abbott, K. J. Reid, P. C. Zee, R. Allada, Universal method for robust detection of circadian state from gene expression, *Proceedings of the National Academy of Sciences* 115 (39) (2018) E9247–E9256.
- [13] N. Leng, L.-F. Chu, C. Barry, Y. Li, J. Choi, X. Li, P. Jiang, R. M. Stewart, J. A. Thomson, C. Kendzierski, Oscope identifies oscillatory genes in unsynchronized single-cell rna-seq experiments, *Nature Methods* 12 (10) (2015) 947–950.
- [14] J. J. Hughey, T. Hastie, A. J. Butte, Zeitzeiger: supervised learning for high-dimensional data from an oscillatory system, *Nucleic Acids Research* 44 (8) (2016) e80–e80.
- [15] D. Vlachou, M. Veretennikova, L. Usselman, V. Vasilyev, S. Ott, G. A. Bjarnason, R. Dallmann, F. Levi, D. A. Rand, Timeteller: a tool to probe the circadian clock as a multigene dynamical system, *PLOS Computational Biology* 20 (2) (2024) e1011779.
- [16] J. Duan, M. N. Ngo, S. S. Karri, L. C. Tsoi, J. E. Gudjonsson, B. Shahbaba, J. Lowengrub, B. Andersen, taufisher predicts circadian time from a single sample of bulk and single-cell pseudobulk transcriptomic data, *Nature Communications* 15 (1) (2024) 3840.
- [17] T. Woelders, V. L. Revell, B. Middleton, K. Ackermann, M. Kayser, F. I. Raynaud, D. J. Skene, R. A. Hut, Machine learning estimation of human body time using metabolomic profiling, *Proceedings of the National Academy of Sciences* 120 (18) (2023) e2212685120.

- [18] Y. Larriba, I. C. Mason, R. Saxena, F. A. Scheer, C. Rueda, Circust: A novel methodology for temporal order reconstruction of molecular rhythms; validation and application towards a daily rhythm gene expression atlas in humans, *PLOS Computational Biology* 19 (9) (2023) e1011510.
- [19] Z. Ji, D. Shang, P. Kim, M. Yang, S. Wu, W. Zhao, E. Kim, M. Wirianto, Z. Chen, S.-H. Yoo, et al., Esocvd reveals the oscillatory patterns of gene expression in alzheimer’s disease, Available at SSRN 4079104.
- [20] D. Mauvoisin, J. Wang, C. Jouffe, E. Martin, F. Atger, P. Waridel, M. Quadroni, F. Gachon, F. Naef, Circadian clock-dependent and-independent rhythmic proteomes implement distinct diurnal functions in mouse liver, *Proceedings of the National Academy of Sciences* 111 (1) (2014) 167–172.
- [21] E. C. Johnson, E. K. Carter, E. B. Dammer, D. M. Duong, E. S. Gerasimov, Y. Liu, J. Liu, R. Betarbet, L. Ping, L. Yin, et al., Large-scale deep multi-layer analysis of alzheimer’s disease brain reveals strong proteomic disease-related changes not observed at the rna level, *Nature Neuroscience* 25 (2) (2022) 213–225.
- [22] N. T. Seyfried, E. B. Dammer, V. Swarup, D. Nandakumar, D. M. Duong, L. Yin, Q. Deng, T. Nguyen, C. M. Hales, T. Wingo, et al., A multi-network approach identifies protein-specific co-expression in asymptomatic and symptomatic alzheimer’s disease, *Cell Systems* 4 (1) (2017) 60–72.
- [23] V. P. Andreev, V. A. Petyuk, H. M. Brewer, Y. V. Karpievitch, F. Xie, J. Clarke, D. Camp, R. D. Smith, A. P. Lieberman, R. L. Albin, et al., Label-free quantitative lc–ms proteomics of alzheimer’s disease and normally aged human brains, *Journal of Proteome Research* 11 (6) (2012) 3053–3067.
- [24] D. C. Hondius, P. van Nierop, K. W. Li, J. J. Hoozemans, R. C. van der Schors, E. S. van Haastert, S. M. van der Vies, A. J. Rozemuller, A. B. Smit, Profiling the human hippocampal proteome at all pathologic stages of alzheimer’s disease, *Alzheimer’s & dementia* 12 (6) (2016) 654–668.

- [25] E. C. Johnson, E. B. Dammer, D. M. Duong, L. Yin, M. Thambisetty, J. C. Troncoso, J. J. Lah, A. I. Levey, N. T. Seyfried, Deep proteomic network analysis of alzheimer’s disease brain reveals alterations in rna binding proteins and rna splicing associated with disease, *Molecular Neurodegeneration* 13 (2018) 1–22.
- [26] H. Haytural, R. Benfeitas, S. Schedin-Weiss, E. Bereczki, M. Rezeli, R. D. Unwin, X. Wang, E. B. Dammer, E. C. Johnson, N. T. Seyfried, et al., Insights into the changes in the proteome of alzheimer disease elucidated by a meta-analysis, *Scientific Data* 8 (1) (2021) 312.
- [27] G. E. Hinton, S. Osindero, Y.-W. Teh, A fast learning algorithm for deep belief nets, *Neural computation* 18 (7) (2006) 1527–1554.
- [28] P. F. Thaben, P. O. Westermark, Detecting rhythms in time series with rain, *Journal of Biological Rhythms* 29 (6) (2014) 391–400.
- [29] M. Moškon, Cosinorpy: a python package for cosinor-based rhythmometry, *BMC Bioinformatics* 21 (2020) 1–12.
- [30] M. D. Ruben, G. Wu, D. F. Smith, R. E. Schmidt, L. J. Francey, Y. Y. Lee, R. C. Anafi, J. B. Hogenesch, A database of tissue-specific rhythmically expressed human genes has potential applications in circadian medicine, *Science translational medicine* 10 (458) (2018) eaat8806.
- [31] M. D. Ruben, L. J. Francey, Y. Guo, G. Wu, E. B. Cooper, A. S. Shah, J. B. Hogenesch, D. F. Smith, A large-scale study reveals 24-h operational rhythms in hospital treatment, *Proceedings of the National Academy of Sciences* 116 (42) (2019) 20953–20958.
- [32] A. Defazio, K. Mishchenko, Learning-rate-free learning by d-adaptation, in: *International Conference on Machine Learning*, PMLR, 2023, pp. 7449–7479.
- [33] M. Dudek, C. Angelucci, D. Pathirana, P. Wang, V. Mallikarjun, C. Lawless, J. Swift, K. E. Kadler, R. Boot-Handford, J. Hoyland, et al., Circadian time series proteomics reveals daily dynamics in cartilage physiology, *Osteoarthritis and Cartilage* 29 (5) (2021) 739–749.

- [34] J. Wang, D. Mauvoisin, E. Martin, F. Atger, A. N. Galindo, L. Dayon, F. Sizzano, A. Palini, M. Kussmann, P. Waridel, et al., Nuclear proteomics uncovers diurnal regulatory landscapes in mouse liver, *Cell Metabolism* 25 (1) (2017) 102–117.
- [35] Z. B. Noordally, M. M. Hindle, S. F. Martin, D. D. Seaton, T. I. Simpson, T. Le Bihan, A. J. Millar, A phospho-dawn of protein modification anticipates light onset in the picoeukaryote *ostreococcus tauri*, *Journal of Experimental Botany* 74 (18) (2023) 5514–5531.
- [36] C. M. Depner, E. L. Melanson, A. W. McHill, K. P. Wright Jr, Mistimed food intake and sleep alters 24-hour time-of-day patterns of the human plasma proteome, *Proceedings of the National Academy of Sciences* 115 (23) (2018) E5390–E5399.
- [37] L. Qian, Y. Gu, Q. Zhai, Z. Xue, Y. Liu, S. Li, Y. Zeng, R. Sun, Q. Zhang, X. Cai, et al., Multitissue circadian proteome atlas of wt and *per1*^{-/-}/*per2*^{-/-} mice, *Molecular & Cellular Proteomics* 22 (12) (2023).
- [38] L. Jiang, M. Wang, S. Lin, R. Jian, X. Li, J. Chan, G. Dong, H. Fang, A. E. Robinson, F. Aguet, et al., A quantitative proteome map of the human body, *Cell* 183 (1) (2020) 269–283.
- [39] E. C. Johnson, E. B. Dammer, D. M. Duong, L. Ping, M. Zhou, L. Yin, L. A. Higginbotham, A. Guajardo, B. White, J. C. Troncoso, et al., Large-scale proteomic analysis of alzheimer’s disease brain and cerebrospinal fluid reveals early changes in energy metabolism associated with microglia and astrocyte activation, *Nature Medicine* 26 (5) (2020) 769–780.
- [40] B. C. Carlyle, S. E. Kandigian, J. Kreuzer, S. Das, B. A. Trombetta, Y. Kuo, D. A. Bennett, J. A. Schneider, V. A. Petyuk, R. R. Kitchen, et al., Synaptic proteins associated with cognitive performance and neuropathology in older humans revealed by multiplexed fractionated proteomics, *Neurobiology of Aging* 105 (2021) 99–114.
- [41] T. G. Beach, C. H. Adler, L. I. Sue, G. Serrano, H. A. Shill, D. G. Walker, L. Lue, A. E. Roher, B. N. Dugger, C. Maarouf, et al., Arizona study of aging and neurodegenerative disorders and brain and body donation

- program, *Neuropathology: Official Journal of the Japanese Society of Neuropathology* 35 (4) (2015) 354.
- [42] Y. Wang, Y. Sun, Y. Wang, S. Jia, Y. Qiao, Z. Zhou, W. Shao, X. Zhang, J. Guo, B. Zhang, et al., Identification of novel diagnostic panel for mild cognitive impairment and alzheimer’s disease: findings based on urine proteomics and machine learning, *Alzheimer’s Research & Therapy* 15 (1) (2023) 1–14.
 - [43] Z. Fang, X. Liu, G. Peltz, Gseapy: a comprehensive package for performing gene set enrichment analysis in python, *Bioinformatics* 39 (1) (2023) btac757.
 - [44] E. Y. Chen, C. M. Tan, Y. Kou, Q. Duan, Z. Wang, G. V. Meirelles, N. R. Clark, A. Ma’ayan, Enrichr: interactive and collaborative html5 gene list enrichment analysis tool, *BMC Bioinformatics* 14 (1) (2013) 1–14.
 - [45] M. V. Kuleshov, M. R. Jones, A. D. Rouillard, N. F. Fernandez, Q. Duan, Z. Wang, S. Koplev, S. L. Jenkins, K. M. Jagodnik, A. Lachmann, et al., Enrichr: a comprehensive gene set enrichment analysis web server 2016 update, *Nucleic Acids Research* 44 (W1) (2016) W90–W97.
 - [46] Z. Xie, A. Bailey, M. V. Kuleshov, D. J. Clarke, J. E. Evangelista, S. L. Jenkins, A. Lachmann, M. L. Wojciechowicz, E. Kropiwnicki, K. M. Jagodnik, et al., Gene set knowledge discovery with enrichr, *Current Protocols* 1 (3) (2021) e90.
 - [47] P. Sweeney, H. Park, M. Baumann, J. Dunlop, J. Frydman, R. Kopito, A. McCampbell, G. Leblanc, A. Venkateswaran, A. Nurmi, et al., Protein misfolding in neurodegenerative diseases: implications and strategies, *Translational Neurodegeneration* 6 (2017) 1–13.
 - [48] A. Loehfelm, A. Boucsein, D. Pretz, A. Tups, Timing matters: Circadian effects on energy homeostasis and alzheimer’s disease, *Trends in Endocrinology & Metabolism* 30 (2) (2019) 132–143.
 - [49] J. Piñero, N. Queralt-Rosinach, A. Bravo, J. Deu-Pons, A. Bauer-Mehren, M. Baron, F. Sanz, L. I. Furlong, Disgenet: a discovery platform for the dynamical exploration of human diseases and their genes, *Database* 2015 (2015) bav028.

- [50] J. Chen, Y. Lu, J. M. Cue, N. Patel, J. J. Zheng, M. J. Cummings, J. Do, Genetic relationship between alzheimer’s disease and schizophrenia, *Alzheimer’s & Dementia* 18 (2022) e065861.
- [51] P. Langfelder, S. Horvath, Wgcna: an r package for weighted correlation network analysis, *BMC Bioinformatics* 9 (1) (2008) 1–13.
- [52] P. Shannon, A. Markiel, O. Ozier, N. S. Baliga, J. T. Wang, D. Ramage, N. Amin, B. Schwikowski, T. Ideker, Cytoscape: a software environment for integrated models of biomolecular interaction networks, *Genome Research* 13 (11) (2003) 2498–2504.
- [53] C.-H. Chin, S.-H. Chen, H.-H. Wu, C.-W. Ho, M.-T. Ko, C.-Y. Lin, cytohubba: identifying hub objects and sub-networks from complex interactome, *BMC Systems Biology* 8 (4) (2014) 1–7.
- [54] D. S. Wishart, C. Knox, A. C. Guo, S. Shrivastava, M. Hassanali, P. Stothard, Z. Chang, J. Woolsey, Drugbank: a comprehensive resource for in silico drug discovery and exploration, *Nucleic Acids Research* 34 (suppl_1) (2006) D668–D672.
- [55] L. S. Mure, H. D. Le, G. Benegiamo, M. W. Chang, L. Rios, N. Jillani, M. Ngotho, T. Kariuki, O. Dkhissi-Benyahya, H. M. Cooper, et al., Diurnal transcriptome atlas of a primate across major neural and peripheral tissues, *Science* 359 (6381) (2018) eaao0318.
- [56] R. Zhang, N. F. Lahens, H. I. Ballance, M. E. Hughes, J. B. Hogenesch, A circadian gene expression atlas in mammals: implications for biology and medicine, *Proceedings of the National Academy of Sciences* 111 (45) (2014) 16219–16224.
- [57] S. Y. Krishnaiah, G. Wu, B. J. Altman, J. Growe, S. D. Rhoades, F. Coldren, A. Venkataraman, A. O. Olarerin-George, L. J. Francey, S. Mukherjee, et al., Clock regulation of metabolites reveals coupling between transcription and metabolism, *Cell Metabolism* 25 (4) (2017) 961–974.
- [58] C. Vollmers, S. Gill, L. DiTacchio, S. R. Pulivarthy, H. D. Le, S. Panda, Time of feeding and the intrinsic circadian clock drive rhythms in hepatic gene expression, *Proceedings of the National Academy of Sciences* 106 (50) (2009) 21453–21458.

- [59] M. E. Hughes, L. DiTacchio, K. R. Hayes, C. Vollmers, S. Pulivarthi, J. E. Baggs, S. Panda, J. B. Hogenesch, Harmonics of circadian gene transcription in mammals, *PLoS genetics* 5 (4) (2009) e1000442.
- [60] L. Talamanca, C. Gobet, F. Naef, Sex-dimorphic and age-dependent organization of 24-hour gene expression rhythms in humans, *Science* 379 (6631) (2023) 478–483.
- [61] F. Atger, C. Gobet, J. Marquis, E. Martin, J. Wang, B. Weger, G. Lefebvre, P. Descombes, F. Naef, F. Gachon, Circadian and feeding rhythms differentially affect rhythmic mrna transcription and translation in mouse liver, *Proceedings of the National Academy of Sciences* 112 (47) (2015) E6579–E6588.
- [62] B. D. Weger, C. Gobet, F. P. David, F. Atger, E. Martin, N. E. Phillips, A. Charpagne, M. Weger, F. Naef, F. Gachon, Systematic analysis of differential rhythmic liver gene expression mediated by the circadian clock and feeding rhythms, *Proceedings of the National Academy of Sciences* 118 (3) (2021) e2015803118.
- [63] S. Bhatia, R. Rawal, P. Sharma, T. Singh, M. Singh, V. Singh, Mitochondrial dysfunction in alzheimer’s disease: opportunities for drug development, *Current Neuropharmacology* 20 (4) (2022) 675.
- [64] R. H. Swerdlow, S. M. Khan, A “mitochondrial cascade hypothesis” for sporadic alzheimer’s disease, *Medical Hypotheses* 63 (1) (2004) 8–20.
- [65] J. M. Adams, S. V. Rege, A. T. Liu, N. V. Vu, S. Raina, D. Y. Kirsher, A. L. Nguyen, R. Harish, B. Szoke, D. P. Leone, et al., Leukotriene a4 hydrolase inhibition improves age-related cognitive decline via modulation of synaptic function, *Science Advances* 9 (46) (2023) eadf8764.
- [66] P. van Bergeijk, U. Seneviratne, E. Aparicio-Prat, R. Stanton, S. A. Hasson, *Srsf1* and *ptbp1* are trans-acting factors that suppress the formation of a *cd33* splicing isoform linked to alzheimer’s disease risk, *Molecular and Cellular Biology* 39 (18) (2019) e00568–18.
- [67] D. Baglietto-Vargas, I. Moreno-Gonzalez, R. Sanchez-Varo, S. Jimenez, L. Trujillo-Estrada, E. Sanchez-Mejias, M. Torres, M. Romero-Acebal, D. Ruano, M. Vizuite, et al., Calretinin interneurons are early targets

of extracellular amyloid- β pathology in ps1/a β pp alzheimer mice hippocampus, Journal of Alzheimer's Disease 21 (1) (2010) 119–132.






Figure S1: Mouse liver encoded data in pre-training stage.



(a) First epoch



(b) Last epoch

Figure S2: Training progress on a random protein in mouse liver data during fine-tuning stage.




Figure S3: Convergence results on time labeled datasets.




Figure S4: Convergence results on brain datasets.




Figure S5: Convergence results on urine dataset.







Figure S6: Accuracy of PROTECT on (a) mouse liver and (b) mouse BAT tissues. The top row shows ROC curves where the y-axis shows the fraction of correctly predicted samples, and the x-axis shows the size of errors. The bottom row shows the scatter plots of predictions vs ground truth.




(a)



(b)



(c)



(d)

Figure S7: Results on mouse liver dataset of Wang et al. [34] using less number of samples: (a) using 4 samples, (b) using 7 samples, (c) using 10 samples, and (d) using 13 samples.




Figure S8: Plots of four randomly chosen proteins known to be rhythmic in thyroid tissue using predicted phases by PROTECT. The y-axis represents protein expression levels, and the x-axis represents the predicted phases (in degrees) as determined by PROTECT.

Table S1: Intersection of rhythmic proteins in 3 brain regions

Intersection Group	Number of Proteins
All 3 Control Sets	276
All 3 AD sets	271
Control $\ominus$ AD	4
AD $\ominus$ Control	0

DLFPC Region	Parital Association Cortex Region	Urine
SEPT11	ANLN	SERPINA7
CTNNB1	HIP1	KRT7
TARDBP	LIPS	HLA-A
AAK1	MBP	CRNN
TSNAX	SIR2	EVPL
CORO1A	CN37	MPO
DNAJA3	MYO1D	TUBB3
PCLO	CD9	C9
PIGN	DOCK1	KRT75
PPP1R7	HSPA1L	ANXA3
FGA	CRNKL1	SERPINA1
SDHA	HSPD1	SERPINA3
FGG	SHROOM4	YWHAZ
VDAC1	CA14	A1BG
ACTN2	FMNL2	KRT79
ACTR2	MYO1E	FGG
NCAM2	ADA10	HPX
SEPT9	GAPR1	SERPIND1
F13A1	DAAM2	HSP90AB1
PLCL1	CSPG2	GATM

Table S2: First 20 Hub Proteins in DLFPC, parital association cortex, and urine datasets




Figure S9: Disparities between control and AD subjects in parietal association cortex using PROTECT predicted phases. (a) Venn diagram of numbers of rhythmic proteins in control (blue) and AD (pink) subjects. (b) Rose plots of distributions of peak times (i.e., acrophases) in rhythmic proteins of control and AD subjects within 24 hours (360 degrees) of the circadian cycle. The radial distance indicates protein counts. (c) Plots of 5 example rhythmic proteins in control subjects that lose rhythmicity in AD.




Figure S10: Disparities between control and AD subjects in DLFP using PROTECT predicted phases. (a) Venn diagram of numbers of rhythmic proteins in control (blue) and AD (pink) subjects. (b) Rose plots of distributions of peak times (i.e., acrophases) in rhythmic proteins of control and AD subjects within 24 hours (360 degrees) of the circadian cycle. The radial distance indicates protein counts. It is noted that the max of radial distance differs between control and AD plots. (c) Plots of 5 example rhythmic proteins in control subjects that lose rhythmicity in AD.




Figure S11: Disease and pathways associated with temporal cortex using DisGeNET, Jensen and KEGG libraries in proteins rhythmic in control subjects but not AD (blue) and proteins rhythmic in AD subject but not in control (pink).




Figure S12: Gene Ontology enrichment analysis of proteins rhythmic in control subjects but not AD (blue) and AD subject but not in control (pink) in parietal association cortex.




Figure S13: Disease and pathways associated with parietal association cortex using DisGeNET, Jensen and KEGG libraries in proteins rhythmic in control subjects but not AD (blue) and proteins rhythmic in AD subject but not in control (pink).




Figure S14: Gene Ontology enrichment analysis of proteins rhythmic in control subjects but not AD (blue) and AD subject but not in control (pink) in DLFPC brain region.




Figure S15: Disease and pathways associated with DLFPFC region using DisGeNET, Jensen and KEGG libraries in proteins rhythmic in control subjects but not AD (blue) and proteins rhythmic in AD subject but not in control (pink).




Figure S16: Venn diagram of numbers of rhythmic proteins in control (blue), AD (pink) and MCI (green) subjects in Urine dataset.




Figure S17: Rose plots of distributions of peak times (i.e., acrophases) in control, AD and MCI subjects in Urine dataset. Each radial distance indicates protein counts.




Figure S18: Ultradian proteins found in temporal cortex with period of 12 hours.




Figure S19: Accuracy of PROTECT on transcriptomic and metabolomic data. (a) mouse liver transcriptomic, (b) baboon amygdala transcriptomic, (c) mouse kidney transcriptomic, and (d) mouse liver metabolomic. The top row shows ROC curves where the y-axis shows the fraction of correctly predicted samples, and the x-axis shows the size of errors. The bottom row shows the scatter plots of predictions vs ground truth.