Characterizing collaborative transcription regulation with a graph-based deep learning approach

Zhenhao Zhang¹, Fan Feng¹, Yuan Yao², and Jie Liu^{1,2,*}

¹*Department of Computational Medicine & Bioinformatics, University of Michigan, Ann Arbor, MI, USA*

²*Department of Computer Science & Engineering, University of Michigan, Ann Arbor, MI, USA*

**corresponding author: Jie Liu (drjieliu@umich.edu)*

Abstract

Human epigenome and transcription activities have been characterized by a number of sequence-based deep learning approaches which only utilize the DNA sequences. However, transcription factors interact with each other, and their collaborative regulatory activities go beyond the linear DNA sequence. Therefore leveraging the informative 3D chromatin organization to investigate the collaborations among transcription factors is critical. We developed ECHO, a graph-based neural network, to predict chromatin features and characterize the collaboration among them by incorporating 3D chromatin organization from 200-bp high-resolution Micro-C contact maps. ECHO predicts 2,583 chromatin features with significantly higher average AUROC and AUPR than the best sequence-based model. We observed that chromatin contacts of different distances affected different types of chromatin features' prediction in diverse ways, suggesting complex and divergent collaborative regulatory mechanisms. Moreover, ECHO was interpretable via gradient-based attribution methods. The attributions on chromatin contacts identify important contacts relevant to chromatin features. The attributions on DNA sequences identify TF binding motifs and TF collaborative binding. Furthermore, combining the attributions on contacts and sequences reveals important sequence patterns in the neighborhood which are relevant to target sequence's chromatin feature prediction. The attribution results that reveal TF collaboration activities are provided on a website <https://echo.dcm.bmed.umich.edu/echo/>.

Introduction

The human body consists of hundreds of different cell types in spite of the identical genotype [1], and the precise gene expression, cell identities and plasticity are dictated by transcriptional regulatory mechanisms [2]. In this process, transcription factors (TFs) bind DNA regulatory sequences by recognizing their binding motifs and form a complex system that regulates chromatin and transcription [3]. Motivated by this, most of current computational models characterize TF binding and other chromatin features only from the DNA sequences. However, growing evidence indicates that transcription factors interact with each other [4, 5] and may form condensates in the 3D chromatin organization [6]. Therefore, it is critical for a computational model to properly characterize collaborative transcription regulation in the light of 3D chromatin organization.

According to whether 3D chromatin organization is utilized, we categorize the deep learning based computational works for predicting chromatin features into *sequence-based* and *graph-based* models. Well-known sequence-based models, such as DeepSEA [7], DanQ [8], DeepBind [9], Basset [10], Basenji [11], and SATORI [12], predict chromatin features only from DNA sequences and ignore the informative chromatin structures. To the best of our knowledge, the only graph-based chromatin feature prediction model ChromeGCN [13] uses a gated graph convolution network to leverage the neighborhood information from a 1kb resolution Hi-C contact map, but it does not fully characterize cooperation among chromatin features.

In this paper, we proposed ECHO (Epigenomic feature analyzer with 3D CHromosome Organization), a graph-based neural network to predict the chromatin features and identify collaboration among them by including 3D chromatin organization. In ECHO, nucleosome-resolution Micro-C contact maps were represented as graphs, in which nodes were non-overlapping 1000 base pair (bp) long genomic segments and weighted edges were chromatin contacts between these segments. Inspired by recent work [14, 15], we transformed the graph structure data to grid structure, which was then operated by 1D convolutions to leverage the neighborhood information. ECHO accurately predicted chromatin features with an average AUROC 0.921 and an average AUPR 0.378 in the prediction of 2,583 chromatin features, significantly higher than the best *sequence-based* model with an average AUROC 0.885 and an average AUPR 0.318. In addition, by evaluating the model using contacts of different distances, we identified patterns about how chromatin contacts of different distances affected chromatin feature prediction.

Moreover, the contributions of Micro-C contacts and DNA sequences to the investigated chromatin features were characterized by applying an attribution method to ECHO. For TFs with known motifs, the corresponding high attribution score regions on DNA sequences match their binding motifs. For TFs without known motifs, previous sequence-based methods may fail to extract prediction patterns from the binding DNA sequences. From the comparison of ECHO and sequence-based model's results, we found that these TFs' prediction accuracy improved more than that of TFs with known motifs. Therefore, ECHO leveraged chromatin structures and extracted information from the neighborhood to assist prediction. As we attributed the TFs to the neighborhood, the high attribution score regions in the neighbor sequences also matched other TFs' binding motifs, which indicated that ECHO recovered TF collaborative binding activities. Furthermore, important sequence patterns were revealed from high attribution score regions in the neighborhood, and some were consistent with existing biological knowledge.

Results

A graph neural network that leverages 3D chromatin organization and predicts chromatin features

We proposed ECHO, a graph neural network model to predict various chromatin features, including transcription factors (TFs) binding activities, histone modifications (HMs), and DNase I hypersensitive sites (DHSs), and characterize their collaboration in the light of high-resolution 3D chromatin organization. ECHO took inputs including one-hot representations of 1000-bp DNA sequences from the reference genome and chromatin contacts from Micro-C contact maps, and output a vector of predicted chromatin features. The architecture of ECHO consisted of sequence layers, graph layers, and one prediction layer (Figure 1a). The sequence layers extracted sequence features. The graph layers aggregated neighborhood information to extract features from the neighborhood. The prediction layer made predictions from previous graph layers. Unlike previous sequence-based models such as DeepSEA [7] and DanQ [8], ECHO leveraged neighborhood information to assist the prediction of chromatin features. Unlike ChromeGCN [13] which used graph convolution networks (GCNs) and treated neighbors 'equally', we first transformed the graph structure of chromatin contact data to a grid structure [14, 15], and then we performed convolution to learn features from sequential and spatial chromatin structures. Moreover, different from ChromeGCN which required full batch training, ECHO used mini-batch training which made it applicable to large graphs.

ECHO Predicted chromatin features more accurately than baseline methods

We first compared the prediction performance of ECHO and three sequence-based models, including DeepCNN (the deep neural network with six convolutional layers used in ExPecto framework [16]), DeepSEA [7], and DanQ [8]. AUROC and AUPR scores were calculated for individual chromatin features. Overall, ECHO predicted chromatin features accurately with an averaged AUROC of 0.921, which is significantly higher than DeepCNN (AUROC 0.885, p -value $6.73E-97$), DeepSEA (AUROC 0.881, p -value

1.89E−124), and DanQ (AUROC 0.881, p -value 1.57E−122) (Figure 1b). Similarly, ECHO yielded an averaged AUPR of 0.378, which is significantly higher than DeepCNN (AUROC 0.318, p -value 7.20E−25), DeepSEA (AUROC 0.312, p -value 1.96E−25), and DanQ (AUROC 0.316, p -value 1.41E−29) (Figure 1b). When we separated model performance in terms of the three categories of chromatin features, namely TF, HM, and DHS, we observed that the ECHO’s dramatically improved the prediction for HM, moderately improved the prediction for TF, and mildly improved the prediction for DHS (Figures 1c and S1).

The current results from ECHO used DeepCNN as the pre-train model to pre-train the sequence layers, and we were curious whether the improvement was consistent if we changed DeepCNN to DeepSEA and DanQ. Experiment results showed that when applying the same ECHO framework to different pre-train sequence-based models such as DeepSEA and DanQ, the prediction was improved consistently, ECHO pre-trained by DeepSEA had an average AUROC of 0.918 and AUPR of 0.373, and ECHO pre-trained by DanQ yielded an average AUROC of 0.919 and AUPR of 0.386 (Table S2).

In addition, we compared ECHO with a graph-based model ChromeGCN [13]. ChromeGCN’s full batch training is computationally prohibitive on our evaluation datasets (over 2.9 million sequence segments and 2,583 chromatin features). Therefore, we used the small dataset processed by ChromeGCN to compare our ECHO and ChromeGCN. The small dataset was for GM12878 only, which included 103 chromatin features (90 TFs, 11 HMs, and 2 DHSs) and a 1kb resolution Hi-C contact map. For both ECHO and ChromeGCN, each 2000-bp input DNA sequence and its reverse complement sequence were first embedded using one-hot encoding (use vectors of 0.25 for sequence gaps and unannotated regions). DeepCNN was selected as the pre-train model. In ECHO, 30 spatial and 10 sequential neighbors were sampled for each input sequence. In ChromeGCN, a combination of constant neighborhood and Hi-C contact maps was used. In both ChromeGCN and ECHO, hidden representations of each original sequence with its reverse complement were extracted when the pre-train model achieved a minimum loss on validation sets, which were then input to the graph layers. The outputs for the original sequence and its reverse complement were averaged in the calculation of the loss. The hyperparameters were all kept the same for ECHO and ChromeGCN. ECHO predicted GM12878’s chromatin features accurately with a mean AUROC of 0.924, a mean AUPR of 0.429 and a mean recall at 50% FDR of 0.399, which outperformed ChromeGCN (AUROC 0.916, AUPR 0.406 and recall at 50% FDR 0.372) (Table S3 and Figure S2).

The influences of Micro-C contact distances on chromatin feature prediction

We next investigate whether the improved chromatin feature prediction depends on the chromatin contact distances. To compare the influences of long-range and short-range Micro-C contacts on chromatin feature prediction, we separated the Micro-C contacts into four non-overlapping groups by contact distance within 0–1 kb, 1–5 kb, 5–20 kb, and over 20 kb. The four groups have 14, 30, 26, and 13 million chromatin contacts, respectively. For each group, we ran ECHO only using the specified range of Micro-C contacts without sampling sequential neighbors. The effects of each contact distance group were evaluated by the same model performance criteria, including AUROC and AUPR. By comparing the performance of each group with the baselines and original ECHO with all Micro-C contacts, we found both the long-range and short-range Micro-C contacts contributed to chromatin feature prediction. Overall, short-range Micro-C contacts produced a more significant on model improvement, which was indicated by the more significant p -values from the paired t -tests with baseline DeepCNN (p -values for AUROC: 0–1 kb (1.95E−29), 1–5 kb (1.91E−42), 5–20 kb (1.62E−13), 20– kb (3.25E−4); p -values for AUPR: 0–1 kb (1.03E−10), 1–5 kb (1.12E−10), 5–20 kb (3.81E−4), 20– kb (3.24E−2) Figure 2a).

We further investigated the influences of Micro-C contact distances on the prediction of TFs and HMs separately. Surprisingly, three patterns were revealed (Figure 2b, 2c, 2d, and 2e). First, for predicting TFs related to chromatin structure maintenance such as CTCF, RAD21, and SMC3, ECHO using short-range Micro-C contacts with contact distances within 0–1 kb and ECHO using long-range Micro-C contacts with contact distances over 20 kb improved the model the most (Figure 2b). The high AUPR and AUROC

scores for over 20 kb Micro-C contacts indicated that ECHO leveraged long-range interactions between these structural proteins. Second, for chromatin features mostly binding to promoters and enhancers and related to gene activation such as H3K4me3, H3K4me1, H3K27ac, H3K9ac, and POLR2A, ECHO achieved the highest AUROC and AUPR scores with 0–5 kb short-range Micro-C contacts, while the performance deteriorated as contact distance increased (Figures 2c and S4).

The last pattern was that the Micro-C contacts with contact distance within 1–20 kb (especially 1–5 kb) improved the prediction of gene elongation marks (including H3K36me3, H3K79me2, and H4K20me1, Figure 2d) and repressive marks (including H3K9me3 and H3K27me3, Figure 2e) the most. Gene elongation marks usually spanned longer than gene activation marks, resulting in correlation over longer distances. DNA methylation was found to be better correlated between distant loci in the B compartment where the repressive marks were enriched [17], and these correlation patterns were also revealed by ECHO. Overall, these patterns validated our assumption that Micro-C contacts of different contact distances could be related to chromatin features in different ways, and they indicated that ECHO identified these complex collaborative regulatory mechanisms, which might contribute to its great improvement on HM and TF prediction.

Chromatin contact resolution is critical in ECHO's prediction of chromatin features

To investigate whether the resolution of chromatin contacts is critical to improve the prediction performance, we compared the performance of ECHO using the high-resolution Micro-C contact maps and the relatively low-resolution Hi-C contact maps. In the previous settings, high-resolution 200-bp Micro-C contact maps were used to provide spatial neighborhood information for ECHO. Different from Hi-C which used a restriction enzyme [18], Micro-C used MNase for chromatin fragmentation [19], providing higher resolution than Hi-C. Since Hi-C could not reach the resolution higher than 1 kb, we up-sampled the 1 kb resolution Hi-C contact maps to 200-bp resolution as follows. To determine the up-sampled contact value between two 200-bp bins, we first found where these two bins located in the original 1 kb Hi-C contact maps, then the up-sampled contact value was set to be the contact value between the located 1 kb bins in the Hi-C contact maps. Consistent with the normalization of Micro-C contact map, both the HFF and hESC Hi-C contact maps were normalized and merged by taking the maximum contact value at each position to generate a weighted adjacency matrix. Because the adjacency matrix from the up-sampled Hi-C contact maps was much denser than the Micro-C contact maps, we filtered out low contact values with some threshold. In this experiment, we chose two thresholds and sampled 50 spatial neighbors and 10 sequential neighbors for each sequence similarly, resulting in 72M and 110M Hi-C contacts, respectively. ECHO using Micro-C contact maps predicted chromatin features with a mean AUROC of 0.921, which was significantly higher than ECHO with 101M Hi-C contacts (AUROC 0.915, p -value $6.57\text{E}-5$), ECHO with 72M Hi-C contacts (AUROC 0.914, p -value $3.67\text{E}-7$) and DeepCNN (AUROC 0.885, p -value $6.73\text{E}-97$). Moreover, ECHO with Micro-C contact maps yielded a mean AUPR of 0.378, which outperformed ECHO with 101M Hi-C contacts (AUPR 0.366, p -value $5.52\text{E}-2$), ECHO with 72M Hi-C contacts (AUPR 0.364, p -value $1.90\text{E}-2$) and DeepSEA (AUPR 0.318, p -value $7.20\text{E}-25$). As the results demonstrated, ECHO with either Micro-C or Hi-C contact maps achieved better performance than baseline DeepCNN. However, ECHO with Micro-C contact maps using fewer contacts outperformed ECHO with Hi-C contact maps, which indicated that high-resolution Micro-C contact maps provided more precise chromatin contact information for predicting chromatin features (Figure 3).

Identifying important chromatin contacts contributing to predicting chromatin features

For characterizing the contributions of Micro-C contacts and DNA sequences towards predicting chromatin features, we applied an attribution method [20] to our ECHO framework, which calculated the attribution scores of the inputs as $\text{gradient} \times \text{input}$ (Figure 4a and Methods). Unlike sequence-based models which only attributed the chromatin features to the sequence itself, ECHO also attributed chromatin feature prediction to the sequences of its neighbors implied by chromatin contacts, which characterizes the contributions

of neighbor sequences to the chromatin features.

Moreover, the attribution scores computed on the contact matrices identified the important contacts relevant to the targeted chromatin feature (Figure S5). To validate the contact attribution, we took CTCF for example. As an architectural protein bridging genome topology and functions [21], CTCF interactions play an important role in DNA looping and transcriptional regulation. Suppose that each contact connected one central sequence and one of its neighbor sequences, we first identified all chromatin contacts whose central sequences were bound to CTCF, and calculated their attribution scores related to CTCF prediction. Then the contacts with attribution scores greater than 0.1 were divided into four groups with contact distance between 5–10 kb, 10–30 kb, 30–100 kb, and over 100 kb. In each group, we calculated the proportion of these contacts whose other anchor was also bound by CTCF, namely the percentage of CTCF interaction (Figure 4b and 4c).

As a baseline, we identified all Micro-C contacts whose central sequence bound by CTCF without filtering the attribution scores, and obtained the percentages of potential CTCF interactions in the same way. As the results demonstrated, within short-range interaction, the percentages of CTCF interactions were significantly lower than those within long-range interaction. However, as the contact distance increased, long-range CTCF interactions were more essential to CTCF binding. When predicting the binding sites of CTCF, if ECHO assigned a high attribution score to a >30 kb Micro-C contact, then CTCF also bound the other anchor of the contact with a very high probability (Figure 4c). Moreover, this pattern still existed when we changed the threshold of attributions scores to 0.3, 0.5, and 0.7, which indicated that CTCF interactions had more effects on CTCF binding as contact distance increased. The distance range matched well with the common CTCF-mediated chromatin looping, indicating that ECHO not only captured the relationships between loops and CTCF interactions but also enhanced the prediction of CTCF binding sites by identifying its potential interacting domains.

Identifying TF binding motifs and collaborative binding mechanisms among TFs

As mentioned before, the attribution scores contributing to the chromatin features on the central sequence were calculated on both the central sequence and its neighbor sequences (Figure 4a). We found that the highly attributed regions on the central sequences matched known binding motifs from JASPAR [22] (examples given in Figure S6a). Furthermore, some high attribution score regions in the neighbor sequences also matched known motifs from JASPAR, which identified TF collaborative binding (examples given in Figure S6c). In addition, TF binding motifs can be identified from attribution scores of corresponding TF binding sequences by using TF-MoDISco [23] (Figure S6b).

The TF collaborative binding mechanisms were further investigated by combining attributions on DNA sequences and Micro-C contacts. For example, a contact between two candidate cis-regulatory elements (cCREs) at chr 2: 28,810,975–28,811,136 and chr 2: 28,821,422–28,821,694 received an attribution score of 1.0 for CTCF labels. By attributing the TFs binding on the central sequence to both the cCREs, we found that the high attribution score regions on the two cCREs matched with the CTCF motif from JASPAR (Figure 5). This indicated that ECHO learned both the CTCF motifs and the CTCF interaction patterns, and the CTCF on neighbor sequence's side cCRE contributed to the binding of CTCF on the central sequence's side cCRE.

We also found a candidate enhancer-promoter interaction between chr 2: 113,627,160–113,627,340 and chr 2: 113,602,125 – 113,602,462 receiving an attribution score 0.925 related with CTCF binding (Figure 5b). By calculating the attribution scores for TFs binding on the central sequence, we identified an ELF1 motif on the central sequence side's cCRE and three TF binding motifs on the neighbor sequence side's cCRE, which contributed to TFs binding on the central sequence. Surprisingly, we identified the binding motif of ZBTB3 which frequently appears in the proximity to CTCF [24], from the neighbor sequence's attribution scores. However, ZBTB3 was not selected in our TF profiles, i.e., ZBTB3 was not labelled on any sequence. As a comparison, we computed attribution scores by using DeepCNN who could only

attribute each sequence's binding TFs to itself, so the ZBTB3 binding motif was not detected as expected. Furthermore, the NRF1 motifs identified by ECHO and DeepCNN were not in the same location, but the attribution scores assigned by ECHO matched the NRF1 motif better than DeepCNN (see Figure S6b and the high attribution score regions highlighted in rectangles for ECHO and highlighted in ovals for DeepCNN).

Discovering sequence patterns in the neighborhood contributing to chromatin features on the central sequences

In this experiment, we investigated whether specific sequence patterns in the neighborhood existed and contributed to the related chromatin features on the central sequences. From the previous results, ECHO outperformed sequence-based models on multiple individual chromatin feature prediction (Figure 1c). Sequence-based models performed worse, particularly on TFs without known motifs, since they failed to extract patterns from the central sequences. However, ECHO, by leveraging neighborhood information, improved the prediction accuracy more significantly on TFs without known motifs than TFs with known motifs (Figures 6c and S7). Thus, we believed that ECHO's improvement in chromatin feature prediction largely came from leveraging neighbor sequences.

This motivated us to capture specific patterns contained in the neighbor sequences and helpful to TF binding. We first calculated attribution scores of chromatin contacts for investigated TFs. Then we selected the highly attributed contacts whose central sequence sides were these TFs' binding sites, and these binding events could be successfully predicted by ECHO. Next, two hundred neighbor sequences from the other anchor of these chromatin contacts were sampled. Next, common sequence patterns were identified from the attribution scores with respect to TFs binding on the central sequences by TF-MoDISco [23].

Examples were given to validate our assumptions (Figure 6a). We first found that ECHO predicted SMAD5 more precisely on all the collected cell lines. Then a FOS:JUN related pattern and a CAGGTG pattern were identified in its neighborhood. Corresponding evidences were found that SMADs had a low affinity for DNA, so they relied heavily on co-factors for the transcriptional regulation [25]. For ZEB1 whose binding motif contains CAGGTG pattern, it bound to coactivator p300 and was highly correlated with SMAD binding by promoting the formation of a p300-Smad transcriptional complex [26]. Additional evidence was found that SMAD interacted with AP-1 [27] whose binding site was identified as 5'-TGA G/C TCA-3', and FOS:JUN is a subunit of AP-1.

We further explore whether these significant patterns existed around HMs. We identified sequence patterns regarding H3K4me3 on cell line GM12878. For all the central sequences predicted to be H3K4me3 in GM12878, we identified four hundred neighbor sequences from important contacts, and their attribution scores were studied by TF-MoDISco. In the end, important patterns were identified and validated by current biological knowledge (Figure 6b). One of the pattern, ATGGCGG, matched parts of CTCF and YY1 motifs, and both YY1-binding and CTCF-YY1 co-binding regions were known to be enriched by H3K4me3 [28]. Another pattern matched NFYA motif from JASPAR, which agreed with the fact that the presence of H3K4me3 depended on NF-Y [29].

Discussion

Our model predicts chromatin features from DNA sequences in the light of high-resolution 3D chromatin organization, which differs from previous sequence-based deep learning models which only utilize DNA sequences. As a result, ECHO significantly outperforms sequence-based models on chromatin feature prediction in terms of AUROC and AUPR scores. The improvement is consistent in terms of the specific sequence-based model ECHO builds on and the specific chromatin feature to be predicted. ECHO also outperforms the only graph-based chromatin feature prediction model in the literature, ChromeGCN [30], which demonstrates that convolution over sampled neighborhood strategy in ECHO is important.

Whereas ECHO and ChromeGCN [13] explicitly leverage chromatin contacts, DNA interactions are

implicitly captured by SATORI [12] and Basenji [11]. SATORI captures TF-TF interactions by combining CNN with self-attention mechanisms. SATORI first finds interactions between positions on DNA sequences with attention matrices, then position interactions are converted to CNN filter interactions and finally translated to TF interactions. However, SATORI captures interactions within the input DNA sequence and cannot capture long-range interactions. Basenji adopts dilated convolution to increase receptive field size, which enables to deal with 131 kb long DNA sequences and capture long-range DNA interactions. While ChromeGCN converts the entire DNA sequence into the graph structure by leveraging Hi-C contact maps, it captures DNA interactions within the entire genome. Moreover, ChromeGCN uses a gated graph convolution network (GCN), which can determine whether to use the aggregated neighborhood information. However, ChromeGCN treats neighbors ‘equally’, but for a DNA sequence segment, some chromatin contacts are useful to gene regulation, while others can be noise. Nonetheless, ECHO further transforms the graph structure data from higher-resolution contact maps to grid structure and applies 1D convolutions to learn the neighborhood aggregation patterns.

Our work also systematically characterizes the contributions of the spatial and sequential genomic neighborhood towards predicting the chromatin features of a central sequence. Different from previous sequence-based prediction models which only calculate attribution scores on each central sequence for its own chromatin features, and also different from ChromeGCN which only computes the attribution scores on the Hi-C contacts excluding DNA sequences, ECHO attributes the chromatin feature prediction to both chromatin contacts and DNA sequences (including central sequences and their neighbors). Moreover, ECHO, which mainly consists of convolutional layers, is more interpretable compared to ChromeGCN. Some popular interpretation methods, such as gradient-based attribution methods and visualizing first convolutional layers, may not perform well on ChromeGCN for the following reasons. First, the full batch training in ChromeGCN is computationally prohibitive to compute the gradients end-to-end. Other typical gradient-based attribution methods, such as DeepLIFT [20] and layer-wise relevance propagation (LRP) [31], may not perform well, since original DeepLIFT cannot handle gating functions in ChromeGCN and LRP can fail with Sigmoid activation function [30]. In addition, the motifs learned from the first convolutional layer are the same as motifs learned from sequence-based models, since pre-training sequence-based models freezes the first convolutional layer in ChromeGCN.

Another contribution in our work is that ECHO with mini-batch training enables to deal with large-scale graphs and datasets. The data used by ECHO includes 200-bp high-resolution Micro-C contact maps with more than 77 million contacts, and more than 2.9 million input DNA sequences, and 2,583 chromatin features.

One of the limitations in our work is that some valuable chromatin contacts may be overlooked in ECHO’s neighborhood sampling procedures. For each central sequence, ECHO samples a number of neighbor sequences according to the contact strength. When one central sequence has multiple neighbor sequences, some important Micro-C contacts with weaker contact strength may be left out. Another limitation is that we only merge two Micro-C contact maps (hESC and HFF) to provide chromatin contact information, but the chromatin features are from multiple cell lines. Although many chromatin structures are preserved across different cell types, we suspect some of the chromatin contacts related to transcription regulation are cell type-specific. Therefore, we anticipate ECHO performs better when additional Micro-C data sets become available in the future.

Methods

Different from previous chromatin feature prediction models, including DeepSEA [7], DanQ [8], DeepBind [9], and Basenji [11], which ignore the important 3D chromatin organization, our model ECHO extracts both sequential features along the DNA sequence and spatial features from $\mathcal{G}$, which explicitly describes high-resolution 3D chromatin organization. In the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{A})$, each node $v \in \mathcal{V}$ represents a non-

overlapping DNA segment of fixed length (i.e., 200 bp), and $\mathcal{V}$ covers the entire DNA sequence. Each edge $e \in \mathcal{E}$ represents a chromatin contact between two DNA segments, and the strength of the chromatin contacts is captured by $\mathcal{A}$. Because some of the 3D chromatin contacts are cell-type specific, we merge several existing Micro-C contact matrices from different cell types by taking the maximum contact value at each entry of the adjacency matrix $\mathcal{A}$.

Model Description

ECHO takes a one-hot encoding representation of one DNA sequence segment $x^{(i)}$ as input, and computes a vector of chromatin features $\hat{y}^{(i)}$ with a number of neural networks layers, including sequence layers, graph layers, and one prediction layer. Since each layer in a deep learning model is a function mapping the layer inputs to the layer outputs, we use functions f , g , and p to represent sequence layers, graph layers, and prediction layer, respectively. The architectural and algorithmic details of our model are given in Figure 1a and Algorithm 1, and the notations used in our work are summarized in Table S1. Next, we describe each of the three layers in detail.

The sequence layers f are designed to extract sequence feature information

$$\phi^{(i)} = f(x^{(i)}),$$

where $i \in \{1, \dots, N\}$ is the index of 1000-bp DNA sequences (including 400-bp flanking regions both upstream and downstream a 200-bp central sequence) ordered by their locations on the reference genome, and N indicates the number of input sequences. $\phi^{(i)} \in \mathbb{R}^K$ represents the hidden representation of sequence i from the sequence layers f . To reduce training parameters and accelerate training, we adopt a pre-train approach which is also used in ChromeGCN [13]. Existing sequence-based models, e.g. DanQ [8] and DeepSEA [7] are first pre-trained to learn features from sequences, and all the pre-trained layers except the last fully connected layer in the sequence-based models are frozen as our sequence layers f .

Since the convolution kernels operate on receptive fields of fixed size and order, original convolutions cannot be performed on a graph [14]. Therefore, between the sequence layers and graph layers, functions SAMPLE and STACK are introduced to transform the graph structure data into grid structure data. Specifically, we first sample a fixed number of sequences from the neighbor set and stack the hidden representations of these sequences to form a feature matrix ϕ which allows convolution operations to extract sequential and spatial features, namely

$$\begin{aligned}\Xi_s^{(i)} &= \text{STACK} \left(\text{SAMPLE} \left(\phi, \mathcal{N}_s^{(i)} \right) \right), \\ \Xi_c^{(i)} &= \text{STACK} \left(\text{SAMPLE} \left(\phi, \mathcal{N}_c^{(i)} \right) \right).\end{aligned}$$

Let $\mathcal{N}_s^{(i)}$ represent a sequential neighbor set containing the neighbors along the chromatin fiber. Let $\mathcal{N}_c^{(i)}$ indicate a spatial neighbor set from the 3D chromatin structure, i.e., a union of sequences which contact with central sequence i , and here only the first-order neighbors with direct contacts are considered. The neighbor set is defined as the union of these two sets, namely

$$\mathcal{N}^{(i)} = \mathcal{N}_s^{(i)} \cup \mathcal{N}_c^{(i)}.$$

Suppose that k_s sequential neighbors from both downstream and upstream of the central sequence i are sampled, then we stack the hidden representations of these selected sequences according to their location

orders on the genome to generate a feature matrix

$$\Xi_s^{(i)} = \begin{bmatrix} (\phi^{(i-k_s)})^T \\ (\phi^{(i-k_s+1)})^T \\ \vdots \\ (\phi^{(i+k_s)})^T \end{bmatrix} \in R^{(1+2k_s) \times K}.$$

Next the spatial neighbor set is defined as

$$\mathcal{N}_c^{(i)} = \{t | \mathcal{A}[i, t] > \tau\}$$

where $\mathcal{A}[i, t]$ is the normalized contact value between sequences i and t , and $\tau > 0$ is a threshold to filter out extremely small contact values and noisy contacts. Then we order the sequences from the neighbor set $\mathcal{N}_c^{(i)}$ by the contact values with the central sequence i , then k_c sequences with top contact values are sampled from the set $\mathcal{N}_c^{(i)}$. Similarly, a feature matrix $\Xi_c^{(i)} \in R^{k_c \times K}$ is generated by stacking the hidden representations of sampled sequences with an order of the contact values. If one central sequence does not have as many neighbors as expected, dummy sequences $M' \in R^K$ with features of all zeros are added to feature matrices to ensure they are of the same size.

The graph layers g perform convolution on the feature matrices $\Xi_c^{(i)}$ and $\Xi_s^{(i)}$ of central sequence i as well as its neighbors to aggregate information from the neighborhood,

$$h_s^{(i)} = g_s(\Xi_s^{(i)}), \quad h_c^{(i)} = g_c(\Xi_c^{(i)}),$$

where $h_s^{(i)}$ and $h_c^{(i)}$ are updated hidden representation extracted by graph layers g_s and g_c . The feature matrix $\Xi_c^{(i)} \in R^{k_c \times K}$ enables 1D convolution by taking k_c as the number of channels and K as the feature size, so as the feature matrix $\Xi_s^{(i)}$. For the two types of feature matrices, their corresponding graph layers g_c and g_s are applied. The structures of graph layers g_c and g_s are the same except for the first layer which receives inputs in different sizes (the numbers of sampled sequential and spatial neighbors are different). Both g_c and g_s consist of convolution layers and a global average pooling [32] in the last layer. Each final feature map reflects one type of sequential and spatial neighborhood information aggregation pattern.

The prediction layer p predicts chromatin features from the updated hidden representations of each input sequence,

$$\hat{y}^{(i)} = p(h_s^{(i)} \parallel h_c^{(i)}),$$

where $\parallel$ indicates a concatenation operation. The hidden features output by the graph layers $h_s^{(i)}$ and $h_c^{(i)}$ are concatenated as the embedding of the central sequence i , which is fed into the prediction layer p with one fully connected layer.

Although higher-order neighborhood information is not investigated in our work, it can also be learned by mixing powers of adjacency matrix [33]. For example, to learn the neighborhood information within an n -th order, additional spatial neighbor lists $[\mathcal{N}_c^{(i)}]^j$ need to be sampled from the j -th power of adjacency matrix $\mathcal{A}^j$ to generate the feature matrix $[\Xi]_c^j$. Then a new embedding $\parallel_{j=2, \dots, n} g_c^j([\Xi]_c^j)$ is concatenated to the first-order node embedding where g_c^j are the graph layers corresponding to j -th order neighborhood

information extraction.

Algorithm 1: workflow of ECHO

Input: Graph($\mathcal{V}, \mathcal{E}, \mathcal{A}$). Input features $x^{(i)} \forall i = 1, \dots, |\mathcal{V}|$. Sequence layers f . Graph layers g_c and g_s . Prediction layers p . Number of sampled sequence and first-order neighbors k_s and k_c . First-order neighbor list $\mathcal{N}_c^{(i)} \forall i = 1, \dots, |\mathcal{V}|$.

Output: Predicted chromatin features $\hat{y}^{(i)} \forall i = 1, \dots, |\mathcal{V}|$.

$\phi^{(i)} \leftarrow f(x^{(i)}) \quad \forall i = 1, \dots, |\mathcal{V}|$

for $i = 1, \dots, |\mathcal{V}|$ **do**

Dummy node $M' \leftarrow \mathbf{0}$

$\Xi_s^{(i)} \leftarrow [\phi^{(i-k_s)}, \dots, \phi^{(i+k_s)}]^T$

$\{i^{[1]}, i^{[2]}, \dots, i^{[|\mathcal{N}_c^{(i)}|]}\} \leftarrow \text{sort}(\mathcal{N}_c^{(i)})$

if $|\mathcal{N}_c^{(i)}| < k_c$ **then**

$\Xi_c^{(i)} \leftarrow [\phi^{(i^{[1]})}, \dots, \phi^{(i^{[|\mathcal{N}_c^{(i)}|]})}, \dots, M']^T$

else

$\Xi_c^{(i)} \leftarrow [\phi^{(i^{[1]})}, \dots, \phi^{(i^{[k_c]})}]^T$

end

end

$h_s^{(i)} \leftarrow g_s(\Xi_s^{(i)}), \quad h_c^{(i)} \leftarrow g_c(\Xi_c^{(i)}) \quad \forall i = 1, \dots, |\mathcal{V}|$

$\hat{y}^{(i)} \leftarrow p(h_s^{(i)} \| h_c^{(i)}) \quad \forall i = 1, \dots, |\mathcal{V}|$

Model Training

For training our models, we first downloaded the human reference genome GRCh38 and removed sequence gaps and unannotated regions. The rest of the genome was segmented into 200-bp bins. Next we collected 2,583 chromatin feature profiles including 882 TFs, 1,510 HMs and 191 DHS profiles from the Encyclopedia of DNA elements (ENCODE) [34] and the International Human Epigenome Consortium (IHEC) [35]. The bins were labelled in a way that if more than half of a bin was in the peak region, then the corresponding entry in its chromatin feature vector was set to be 1, and 0 otherwise. We had over 2.9 million 200-bp segments bound with at least one TF, resulting in 585,137,600-bp DNA sequence (20.4% of the human reference genome). We added a 400-bp flanking region to both upstream and downstream of the 200-bp sequence to generate a 1000-bp sequence. As input, each sequence was represented by a 4×1000 one-hot encoding matrix. These input sequences were split to a training set, a validation set, and a testing set without overlapping. Here we chose chromosomes 2, 8, and 21 as the testing set, chromosomes 3 and 12 as our validation set, and the remaining chromosomes were used to train our model.

The default pre-train model in ECHO was DeepCNN if not specified. The length of each hidden feature vector $\phi^{(i)}$ extracted by sequence layers f was 2600. The weighted adjacency matrix $\mathcal{A}$ was generated by merging Micro-C contact matrices of two cell lines, HFF and hESC as follows. We first calculated the total read counts in both the contact maps for each chromosome and the ratio between the two maps. Each contact value in the contact matrix with a larger total read count was multiplied with the ratio to ensure that the two contact matrices had the same total read count. Then we merged the normalized contact matrices by taking the maximum contact value at each entry in the matrices. In addition, we used a threshold to filter out noisy contacts in the adjacency matrix, i.e., all the entries with contact values < 2 were removed. Next we sampled 50 spatial neighbors and 10 sequential neighbors for each sequence, resulting in a total number of 77M chromatin contacts used.

ECHO and all the baseline models were trained on a NVIDIA Tesla V100 GPU with a batch size 64 and optimized by stochastic gradient descent with a momentum of 0.9 and a learning rate of 0.5. For the loss function, we chose a mean binary cross entropy loss

$$\text{BCELoss} = -\frac{1}{N} \frac{1}{L} \sum_{i=1}^N \sum_{l=1}^L y_s^{(i)} \log(\hat{y}_l^{(i)}) + (1 - y_l^{(i)}) \log(1 - \hat{y}_l^{(i)}),$$

where N indicates the number of samples, L is the length of target chromatin feature vector, $y_l^{(i)}$ and $\hat{y}_l^{(i)}$ represent the l -th element in target and predicted chromatin feature vector for sequence i , respectively.

Attribution Methods

Even though our ECHO model predicts chromatin features accurately, researchers still hope to identify DNA segments and chromatin contacts contribute to the prediction of a specific chromatin feature. Visualizing the filters from the first convolutional layer is useful to discover DNA binding motifs, but the patterns learned by individual filters can be redundant, and one motif pattern can be the result of cooperation among multiple filters [23]. In our work, we adopt a gradient-based attribution method. The inputs to our model include both DNA sequences and their interactions, and we use $\text{gradient} \times \text{input}$ [20] to calculate attribution scores on the inputs (DNA sequences and Micro-C contacts). The score on each base pair indicates which base pair of a DNA sequence and which chromatin contacts contribute to the prediction of a given chromatin feature. The attribution method is described in Figure 4a and explained thoroughly in the following sections.

Attribution on chromatin contacts

For each sequence i , generating its feature matrix $\Xi^{(i)}$ is a multiplication between a binary sampling matrix and the hidden representation matrix $\phi \in R^{N \times k}$, $\Xi_s^{(i)} = P_s^{(i)} \phi$ and $\Xi_c^{(i)} = P_c^{(i)} \phi$, where $P_s^{(i)} \in \{0, 1\}^{k_s \times N}$ and $P_c^{(i)} \in \{0, 1\}^{k_c \times N}$ are two binary sampling matrices used to sample sequential and spatial neighbor sequences, respectively. Both the two matrices have exactly one 1 in each row, indicating which sequence is sampled. Then we use $\text{gradient} \times \text{input}$ to calculate the attribution scores of the sampling matrices for chromatin feature l

$$[S_c^{(i)}]_l = P_c^{(i)} \odot \frac{\partial \hat{y}_l^{(i)}}{\partial P_c^{(i)}}, \quad [S_s^{(i)}]_l = P_s^{(i)} \odot \frac{\partial \hat{y}_l^{(i)}}{\partial P_s^{(i)}},$$

where $\odot$ represents the Hadamard product. $[S_c^{(i)}]_l$ and $[S_s^{(i)}]_l$ are the attribution scores of the two binary sampling matrices for chromatin feature l . If attribution scores need to be calculated for a set of chromatin features $\mathcal{L}$ (e.g. same type of chromatin features from different cell lines), we have

$$[S_c^{(i)}]_{\mathcal{L}} = \sum_{l \in \mathcal{L}} [S_c^{(i)}]_l, \quad [S_s^{(i)}]_{\mathcal{L}} = \sum_{l \in \mathcal{L}} [S_s^{(i)}]_l.$$

Then $[S_c^{(i)}]_l$ and $[S_s^{(i)}]_l$ are compressed to two interaction importance vectors $[V_c^{(i)}]_l \in R^{1 \times N}$ and $[V_s^{(i)}]_l \in R^{1 \times N}$ by taking the one non-zero value at each row. Each non-zero element in the vectors indicates the importance of sampling the corresponding neighbor sequence, which shows the contact importance with central sequence i . Since one neighbor sequence may be sampled as the sequential and spatial neighbors at the same time, we take the maximum of the two vectors to generate one single vector which is taken as the i -th row of a sparse interaction importance matrix $M_l \in R^{N \times N}$. An element $M_l[i, j]$ indicates the importance of a contact between sequence j and the central sequence i for chromatin feature l , while if sequence j is not sampled as i 's neighbor sequence, then the element is set 0. In the end, we take the absolute value and normalize each row of the interaction importance matrix as

$$\widehat{M}_l[i, j] = \frac{|M_l[i, j]|}{\max_{t \in \{1, \dots, N\}} (|M_l[i, t]|)} \in [0, 1].$$

Attribution on DNA sequences

In addition to calculating attribution scores of the central sequences like sequence-based models, ECHO computes an attribution score $[S^{(j)}]_l^{(i)}$ for the neighbor sequence $j \in \mathcal{N}^{(i)}$ for chromatin feature l on the

central sequence i (Figure 4a)

$$[S^{(j)}]_l^{(i)} = x^{(j)} \odot \left(\frac{\partial y_l^{(i)}}{\partial \Xi^{(i)}} \cdot \frac{\partial \Xi^{(i)}}{\partial x^{(j)}} \right) \quad \forall j \in \mathcal{N}^{(i)} \cup \{i\},$$

where $\mathcal{N}^{(i)}$ indicates the neighbor set of the central sequence i . Similarly, to calculate the attribution scores for a set of chromatin features $\mathcal{L}$, we have

$$[S^{(j)}]_{\mathcal{L}}^{(i)} = \sum_{l \in \mathcal{L}} [S^{(j)}]_l^{(i)}.$$

The high attribution score regions in the central sequence and its neighbors significantly contribute to the prediction of the central sequence’s chromatin features.

Combining attribution on DNA sequences and chromatin contacts

The attribution methods on DNA sequences and Micro-C contacts are combined to further interpret ECHO, e.g., discovering important sequence patterns in the neighborhood. To identify such sequence patterns for specific chromatin features, we have two requirements. The first requirement is that the contacts which connect central sequences and the neighbor sequence need to have high attribution scores for the investigated chromatin features. Then the second one is that the sequence patterns are generated from high attribution score regions from those neighbor sequences. Therefore, we first identify highly attributed chromatin contacts according to the chromatin features investigated, and select the contacts if the investigated chromatin features are successfully predicted on the central sequence sides. Next, each pair of central and neighbor sequences connected by the contact is further attributed regarding the contributions to the chromatin features on the central sequence. To explore the TF collaborative binding mechanism, the high attribution score regions on both the central and neighbor sequences are compared with known TF motifs from JASPAR [22]. To discover important sequence patterns in the neighborhood contributing to the chromatin features on the central sequences, we collect a number of neighbor sequences satisfying the requirements discussed above and compute their attribution scores. Then sequence patterns are generated from these sequences with their attribution scores by TF-MoDISco [23].

Data availability

The chromatin contacts are obtained from 200-bp Micro-C contact maps of two cell lines, hESC and HFF. Our 2,583 chromatin feature profiles are collected from the Encyclopedia of DNA elements (ENCODE) [34] and the International Human Epigenome Consortium (IHEC) [35]. The download urls of the chromatin feature profiles and the source code of ECHO are available in the GitHub repository: <https://github.com/liu-bioinfo-lab/echo>. The attribution results that reveal TF collaboration activities are provided on a website <https://echo.dcm.med.umich.edu/echo/> so that other researchers can further explore them.

Funding

The work was supported by National Human Genome Research Institute award R35HG011279.

References

- [1] Xiaoping Han, Ziming Zhou, Lijiang Fei, Huiyu Sun, Renying Wang, Yao Chen, Haide Chen, Jingjing Wang, Huanna Tang, Wenhao Ge, et al. Construction of a human cell landscape at single-cell level. *Nature*, 581(7808):303–309, 2020.

- [2] Ty C Voss and Gordon L Hager. Dynamic regulation of transcriptional states by chromatin and transcription factors. *Nature Reviews Genetics*, 15(2):69–81, 2014.
- [3] Samuel A Lambert, Arttu Jolma, Laura F Campitelli, Pratyush K Das, Yimeng Yin, Mihai Albu, Xiaoting Chen, Jussi Taipale, Timothy R Hughes, and Matthew T Weirauch. The human transcription factors. *Cell*, 172(4):650–665, 2018.
- [4] Gerald Stampfel, Tomáš Kazmar, Olga Frank, Sebastian Wienerroither, Franziska Reiter, and Alexander Stark. Transcriptional regulators form diverse groups with context-dependent regulatory functions. *Nature*, 528(7580):147–151, 2015.
- [5] Arttu Jolma, Yimeng Yin, Kazuhiro R Nitta, Kashyap Dave, Alexander Popov, Minna Taipale, Martin Enge, Teemu Kivioja, Ekaterina Morgunova, and Jussi Taipale. DNA-dependent formation of transcription factor pairs alters their binding specificity. *Nature*, 527(7578):384–388, 2015.
- [6] Krishna Shrinivas, Benjamin R Sabari, Eliot L Coffey, Isaac A Klein, Ann Boija, Alicia V Zamudio, Jurian Schuijers, Nancy M Hannett, Phillip A Sharp, Richard A Young, et al. Enhancer features that drive formation of transcriptional condensates. *Molecular cell*, 75(3):549–561, 2019.
- [7] Jian Zhou and Olga G Troyanskaya. Predicting effects of noncoding variants with deep learning–based sequence model. *Nature methods*, 12(10):931–934, 2015.
- [8] Daniel Quang and Xiaohui Xie. DanQ: a hybrid convolutional and recurrent deep neural network for quantifying the function of DNA sequences. *Nucleic acids research*, 44(11):e107–e107, 2016.
- [9] Babak Alipanahi, Andrew DeLong, Matthew T Weirauch, and Brendan J Frey. Predicting the sequence specificities of DNA-and RNA-binding proteins by deep learning. *Nature biotechnology*, 33(8):831–838, 2015.
- [10] David R Kelley, Jasper Snoek, and John L Rinn. Basset: learning the regulatory code of the accessible genome with deep convolutional neural networks. *Genome research*, 26(7):990–999, 2016.
- [11] David R Kelley, Yakir A Reshef, Maxwell Bileschi, David Belanger, Cory Y McLean, and Jasper Snoek. Sequential regulatory activity prediction across chromosomes with convolutional neural networks. *Genome research*, 28(5):739–750, 2018.
- [12] Fahad Ullah and Asa Ben-Hur. A self-attention model for inferring cooperativity between regulatory features. *Nucleic acids research*, 05 2021.
- [13] Jack Lanchantin and Yanjun Qi. Graph convolutional networks for epigenetic state prediction using both sequence and 3D genome data. *BioRxiv*, page 840173, 2019.
- [14] Hongyang Gao, Zhengyang Wang, and Shuiwang Ji. Large-scale learnable graph convolutional networks. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1416–1424. ACM, 2018.
- [15] Mathias Niepert, Mohamed Ahmed, and Konstantin Kutzkov. Learning convolutional neural networks for graphs. In *International conference on machine learning*, pages 2014–2023. PMLR, 2016.
- [16] Jian Zhou, Chandra L Theesfeld, Kevin Yao, Kathleen M Chen, Aaron K Wong, and Olga G Troyanskaya. Deep learning sequence-based ab initio prediction of variant effects on expression and disease risk. *Nature genetics*, 50(8):1171–1179, 2018.

- [17] J. Fortin and K. D. Hansen. Reconstructing A/B compartments as revealed by Hi-C using long-range correlations in epigenetic data. *Genome Biology*, 16, 2015.
- [18] Erez Lieberman-Aiden, Nynke L Van Berkum, Louise Williams, Maxim Imakaev, Tobias Ragoczy, Agnes Telling, Ido Amit, Bryan R Lajoie, Peter J Sabo, Michael O Dorschner, et al. Comprehensive mapping of long-range interactions reveals folding principles of the human genome. *science*, 326(5950):289–293, 2009.
- [19] Tsung-Han S Hsieh, Assaf Weiner, Bryan Lajoie, Job Dekker, Nir Friedman, and Oliver J Rando. Mapping nucleosome resolution chromosome folding in yeast by micro-C. *Cell*, 162(1):108–119, 2015.
- [20] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. *arXiv preprint arXiv:1704.02685*, 2017.
- [21] C.-T. Ong and V. G. Corces. CTCF: an architectural protein bridging genome topology and function. *Nature Reviews Genetics*, 15(4):236–246, 2014.
- [22] Oriol Fornes, Jaime A Castro-Mondragon, Aziz Khan, Robin van der Lee, Xi Zhang, Phillip A Richmond, Bhavi P Modi, Solenne Correard, Marius Gheorghe, Damir Baranašić, Walter Santana-Garcia, Ge Tan, Jeanne Chèneby, Benoit Ballester, François Parcy, Albin Sandelin, Boris Lenhard, Wyeth W Wasserman, and Anthony Mathelier. JASPAR 2020: update of the open-access database of transcription factor binding profiles. *Nucleic acids research*, 48(D1):D87–D92, 11 2019.
- [23] Avanti Shrikumar, Katherine Tian, Žiga Avsec, Anna Shcherbina, Abhimanyu Banerjee, Mahfuza Sharmin, Surag Nair, and Anshul Kundaje. Technical note on transcription factor motif discovery from importance scores (TF-MoDISco) version 0.5. 6.5. *arXiv preprint arXiv:1811.00416*, 2018.
- [24] Alan P Boyle, Lingyun Song, Bum-Kyu Lee, Darin London, Damian Keefe, Ewan Birney, Vishwanath R Iyer, Gregory E Crawford, and Terrence S Furey. High-resolution genome-wide in vivo footprinting of diverse transcription factors in human cells. *Genome research*, 21(3):456–464, 2011.
- [25] D.S.J. Miller and C.S. Hill. TGF- β superfamily signaling. In Ralph A. Bradshaw and Philip D. Stahl, editors, *Encyclopedia of Cell Biology*, pages 37–50. Academic Press, Waltham, 2016.
- [26] Antonio A Postigo, Jennifer L Depp, Jennifer J Taylor, and Kristen L Kroll. Regulation of Smad signaling through a differential recruitment of coactivators and corepressors by ZEB proteins. *The EMBO journal*, 22(10):2453–2462, 2003.
- [27] Franck Verrecchia, Laurence Vindevoghel, Robert J Lechleider, Jouni Uitto, Anita B Roberts, and Alain Mauviel. Smad3/AP-1 interactions control transcriptional responses to TGF- β in a promoter-specific manner. *Oncogene*, 20(26):3332–3340, 2001.
- [28] Petra C Schwalie, Michelle C Ward, Carolyn E Cain, Andre J Faure, Yoav Gilad, Duncan T Odom, and Paul Flicek. Co-binding by YY1 identifies the transcriptionally active, highly conserved set of CTCF-bound regions in primate genomes. *Genome biology*, 14(12):1–15, 2013.
- [29] Andrea Fossati, Diletta Dolfini, Giacomo Donati, and Roberto Mantovani. NF-Y recruits Ash2L to impart H3K4 trimethylation on CCAAT promoters. *PloS one*, 6(3):1–9, 03 2011.
- [30] Marco Ancona, Enea Ceolini, Cengiz Öztireli, and Markus Gross. Towards better understanding of gradient-based attribution methods for deep neural networks. *arXiv preprint arXiv:1711.06104*, 2017.

- [31] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS one*, 10(7):e0130140, 2015.
- [32] Min Lin, Qiang Chen, and Shuicheng Yan. Network in network, 2014.
- [33] Sami Abu-El-Haija, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. Mixhop: Higher-order graph convolutional architectures via sparsified neighborhood mixing. *arXiv preprint arXiv:1905.00067*, 2019.
- [34] Carrie A Davis, Benjamin C Hitz, Cricket A Sloan, Esther T Chan, Jean M Davidson, Idan Gabdank, Jason A Hilton, Kriti Jain, Ulugbek K Baymuradov, Aditi K Narayanan, Kathrina C Onate, Keenan Graham, Stuart R Miyasato, Timothy R Dreszer, J Seth Strattan, Otto Jolanki, Forrest Y Tanaka, and J Michael Cherry. The Encyclopedia of DNA elements (ENCODE): data portal update. *Nucleic acids research*, 46(D1):D794–D801, 11 2017.
- [35] David Bujold, David Anderson de Lima Morais, Carol Gauthier, Catherine Côté, Maxime Caron, Tony Kwan, Kuang Chung Chen, Jonathan Laperle, Alexei Nordell Markovits, Tomi Pastinen, Bryan Caron, Alain Veilleux, Pierre Étienne Jacques, and Guillaume Bourque. The international human epigenome consortium data portal. *Cell Systems*, 3(5):496 – 499.e2, 2016.
- [36] W James Kent, Charles W Sugnet, Terrence S Furey, Krishna M Roskin, Tom H Pringle, Alan M Zahler, and David Haussler. The human genome browser at UCSC. *Genome research*, 12(6):996–1006, 2002.

Figure 1: Architectural details of ECHO and its empirical performance in chromatin feature prediction tasks.

(a) Model architecture. The model inputs are one-hot representations of DNA sequences. The inputs x are first fed into sequence layers f to extract hidden representations ϕ . For central sequence i , we first sample a fixed number of sequential neighbor sequences from its sequential neighbor set $\mathcal{N}_c^{(i)}$ and a fixed number of spatial neighbor sequences from the spatial neighbor set $\mathcal{N}_s^{(i)}$. The hidden representations of sampled sequential neighbors are stacked into a feature matrix $\Xi_c^{(i)}$ which is input to graph layers g_c to learn an updated hidden representation $h_c^{(i)}$. The feature matrix $\Xi_s^{(i)}$ stacked by spatial neighbors is fed into g_s to learn an updated hidden representation $h_s^{(i)}$. A feature embedding concatenated by the two hidden representations will be fed into the final prediction layer p to compute a chromatin feature vector $y^{(i)}$. **(b)** The AUROC and AUPR scores for individual chromatin feature are provided according to ECHO and three baselines including DeepCNN, DeepSEA and DanQ. The given p -values from paired t -tests indicate that ECHO achieves a significant improvement. **(c)** Scatter plots of each chromatin feature's AUROC score (top panels) and AUPR scores (bottom panels) from ECHO and DeepCNN. ECHO outperforms DeepCNN on TF, HM and DHS profile prediction for both AUROC scores and AUPR scores.

Figure 2: How much chromatin contacts helps to predict a chromatin feature is contact-distance specific.

(a) Comparing the performance of ECHO using Micro-C contacts from specified contact distance ranges. The x-axis shows four groups with different contact distance ranges and the total number of Micro-C contacts used. The p -values calculated by paired t -tests are given. (b-e) show the effects of Micro-C contact distances on the prediction of chromatin structure maintenance related TFs, gene activation related chromatin features, gene elongation marks, and gene repressive marks, respectively.

Figure 3: Comparing ECHO's performance between when we used 200-bp Micro-C contact maps and when we used up-sampled 200-bp Hi-C contact maps.

The p -values from paired t -tests indicate that Micro-C contact maps provide more precise high resolution chromatin contact information for chromatin feature prediction. The x-axis shows the contact map types and the total numbers of used contacts resulting from two contact value cut-off thresholds. The median AUROC and AUPR scores are provided above the median lines.

Figure 4: Details of attribution methods and validations of attributions on Micro-C contacts.

(a) Workflow of attribution methods. If sequences u and v are the spatial neighbors of the central sequence i , sequences $i-1$ and $i+1$ are i 's sequential neighbors, then in the forward propagation, sequence layers f are applied to extract hidden representation $\phi^{(i)}$ for each sequence. Next, two binary sampling matrices, $P_c^{(i)}$ and $P_s^{(i)}$, are input to multiply with ϕ to form two feature matrices, $\Xi_c^{(i)}$ and $\Xi_s^{(i)}$. The two feature matrices are fed into graph layers (g_c and g_s) and one prediction layer p to compute a predicted chromatin feature vector $\hat{y}^{(i)}$. Next we take the gradient of the specific label l on i with respect to the inputs (two sampling matrices, the central sequence and neighbor sequences) independently, the attribution scores of each input which indicate the contributions to the central sequence i 's label l , are calculated by *gradient $\times$ input* [20]. The attributions on the two sampling matrices, $[S_c^{(i)}]_l$ and $[S_s^{(i)}]_l$, are first compressed into two vectors, $[V_c^{(i)}]_l$ and $[V_s^{(i)}]_l$, and then taken the maximum to become the i -th row in the interaction importance matrix M_l . In addition, the attribution scores $[S^{(i-1)}]_l^{(i)}$ of the sequence $i-1$ for label l on sequence i are computed. **(b)** Logarithm plots to show the total numbers of Micro-C contacts whose central sequence side bound by CTCF. Micro-C contacts are separated into four groups according to the contact distances: 5–10 kb, 10–30 kb, 30–100 kb, and 100+ kb. In each group, the contacts with CTCF related attribution scores greater than 0.1, 0.3, 0.5, and 0.7 are selected, and group ‘all’ is a baseline where contacts are not filtered by attribution scores. **(c)** Bar-plots to show the percentage of CTCF interactions, which is calculated as the number of contacts whose two sides were bound by CTCF over the total number of contacts whose central sequence was bound by CTCF.

Figure 5: Visualization of attribution scores on chromatin contacts, central sequences, and neighbor sequences. The hESC Micro-C contact map is visualized in the UCSC genome browser [36]. The green arrow represents a Micro-C contact pointing to the central sequence side. The attribution scores of the Micro-C contacts for CTCF are provided. The candidate cis-regulatory elements (cCREs) within both anchors of the contact are shown as small colored rectangles, which are attributed for TFs binding on the central sequence. High attribution score regions are plotted in the black rectangles. The binding motifs from JASPAR are marked with *. As a comparison, the attribution scores of cCREs on the neighbor sequence sides contributing to TFs binding on the neighbor sequences are computed using the baseline DeepCNN, which are shown in the large blue blocks. The corresponding high attribution score regions are plotted in ovals. (a) The high attribution score regions on both the two cCREs match known CTCF motifs from JASPAR. (b) The regions within the red rectangles match binding motifs of TFs not included in our collected data. The attribution scores of the neighbor sequence's cCREs computed by ECHO identify ZBTB3 which was not identified by DeepCNN. The attribution scores from ECHO detected NRF1's motif, which is more proximate than the one identified by DeepCNN.

Figure 6: Model performance improvement compared to DeepCNN and visualization of sequence patterns in the neighborhood of investigated chromatin features.

(a) ECHO predicts RBFOX2 and SMAD5 more accurately on all the collected cell lines than the baseline DeepCNN. The patterns generated by TF-MoDISco [23] using the selected neighbor sequences' attribution scores for the corresponding TFs binding on central sequences, which are compared with known binding motifs from JASPAR [22]. (b) ECHO outperforms DeepCNN on GM12878 H3K4me3 prediction and the important patterns in the neighborhood are visualized. 400 neighbor sequences which have high attribution score contacts with successfully predicted GM12878 H3K4me3 sites are selected. The attribution scores of the neighbor sequences are input to TF-MoDISco to identify frequent patterns. Top four patterns are given, and some of them match several TFs' binding motifs from JASPAR. (c) 100 TFs with the cell lines which achieve greatest improvement on AUROC compared to DeepCNN.