UniversalEPI: harnessing attention mechanisms to decode chromatin interactions in rare and unexplored cell types

Aayush Grover[1,2], Lin Zhang[3], Till Muser[3], Simeon Häfliger[1], Minjia Wang[1], Josephine Yates[1], Eliezer M. Van Allen[4], Fabian J. Theis[5,6], Ignacio L. Ibarra[5*], Ekaterina Krymova[3*], Valentina Boeva[1,2,7*]

[1]Department of Computer Science, ETH Zurich, Zurich, Switzerland;
[2]Swiss Institute of Bioinformatics, Lausanne, Switzerland;
[3]Swiss Data Science Center, EPF Lausanne and ETH Zurich, Zurich, Switzerland;
[4]Broad Institute of MIT and Harvard, Harvard Medical School, Dana-Farber Cancer Institute, Boston, USA;
[5]Institute of Computational Biology, Helmholtz Center Munich, Munich, Germany;
[6]School of Computation, Information and Technology, Technical University of Munich, Munich, Germany;
[7]Université Paris Cité, Institut Cochin, INSERM U1016, Paris, France.

*Corresponding should be addressed: ignacio.ibarra@helmholtz-muenchen.de; ekaterina.krymova@sdsc.ethz.ch; valentina.boeva@inf.ethz.ch;

Abstract
Chromatin interactions between cis-regulatory elements (CREs) and specific gene promoters offer insights into the mechanisms of gene activation and repression. This knowledge is crucial for evaluating the effects of non-coding variants and chromatin alterations in cell differentiation or disease. Experimental techniques Hi-C and single-cell Hi-C provide information on CRE-promoter interactions, but the high costs and labor requirements limit their widespread use. We introduce UniversalEPI, an attention-based neural network designed to predict CRE-promoter interactions up to 2 Mb apart across unseen cell types without model retraining and using only DNA sequence and chromatin accessibility data as input. In UniversalEPI, chromatin accessibility profiles with masked non-accessible regions and DNA sequence data get passed through several convolution layers to identify binding motifs of three key transcription factors ubiquitously involved in promoter activation and CRE-promoter looping; a transformer block then takes the information from the convolutional levels passed through stochastic gating to predict chromatin interactions. While being a relatively lightweight model because of its architecture, UniversalEPI outperforms the state-of-the-art methods in inferring CRE-promoter interactions in unseen cell types across human cancerous and non-cancerous datasets. UniversalEPI can thus also be used to incorporate pseudo-bulk chromatin accessibility data from single-cell ATAC-seq experiments to pinpoint cell subtype-specific interactions. We illustrate this ability of UniversalEPI by tracking the dynamics of CRE-promoter interactions during macrophage reprogramming and the transition between cancer transcriptional states in esophageal adenocarcinoma. UniversalEPI thus advances in silico 3D chromatin modeling, supporting studies of chromatin dynamics in organism development and impacts of non-coding variants in cell differentiation and disease.

Keywords: Chromatin, Cis-regulatory elements, Enhancer-promoter interactions, deep learning, transformer

1 Introduction

In complex organisms like humans, different cell types have different transcriptional programs leading to varying cellular functions which are important for the organism's survival. These differences are linked to specific chromatin interactions between gene promoters and cis-regulatory elements (CRE), often referred to as enhancers, which regulate gene expression [1, 2]. Disruptions in chromatin interactions are associated with various genetic diseases, including cancer, making their study crucial for understanding disease initiation and progression. Additionally, chromatin interactions provide insights into the functions of non-coding genomic regions, which are pivotal in studying cell differentiation during organism development and disease mechanisms. Chromatin interactions can be profiled experimentally via Hi-C, an unbiased but costly high-throughput chromosome conformation capture (3C) technique. Due to the high time and resource demands of Hi-C, interest has increasingly shifted to in silico methods for inferring these interactions.

Several recent studies have proposed deep learning-based architectures to predict chromatin architecture. This includes methods predicting chromatin interaction in a specific cell type for hold-out genomics regions from the DNA sequence alone [3–5]. These methods can be used to assess the effects of genomic variants on chromatin structure in a cell type of interest. While these methods generate realistic Hi-C maps, by design, they do not generalize to cell types unseen during the model training. Methods that can predict cell-type-specific chromatin interactions across cell types require additional epigenetic features as inputs, such as accessible chromatin regions and histone marks [6–8]. Nevertheless, these methods may not truly allow for a generalization across cell types, as by construction, they capture the information related to the regulatory mechanisms by transcription factors (TFs) expressed in the cell types used for training. Moreover, these methods require diverse input features which are not always readily available. Some other methods focus only on predicting interactions between a given enhancer and a promoter using epigenetic features [9, 10]. However, such methods cannot be used to understand the mechanisms of non-coding mutations as these

methods cannot capture the effect of mutations in the insulators, which can affect an enhancer-promoter interaction.

The recently published method C.Origami provides a state-of-the-art solution to predicting cell-type-specific 3D-chromatin interactions using bulk data as input [11]. Specifically, C.Origami, based on convolutional neural networks (CNNs) combined with transformers, accurately predicts chromatin interactions across cell types from DNA sequence, profiles of accessible chromatin, and CTCF binding affinities. While C.Origami was designed to work with bulk inputs, the ChromaFold method was recently developed to predict chromatin interactions from single-cell chromatin accessibility profiles leveraging co-accessibility across metacells [12]. However, despite the progress made so far, there is still room for significant improvements including increasing model transferability to cell types unseen during training by limiting the model's visual scope to ubiquitously functional genomic elements and creating a lightweight model that can work with both single-cell and bulk chromatin accessibility data.

In this work, we propose UniversalEPI, a lightweight deep-learning model consisting of CNN layers and transformer blocks. It overcomes the shortcomings of earlier methods by (*i*) using, for training, DNA binding information of ubiquitously expressed TFs CTCF, YY1, and SP1 that are involved in determining enhancer-promoter architecture across cell types [13–15], (*ii*) using DNA sequence and chromatin accessibility features as inputs to make the model cell-type-specific without heavily depending on large datasets, and (*iii*) training solely on DNA accessible regions to reduce noise and model size while still capturing all functionally relevant interactions in large genomic domains (up to 2Mb). Furthermore, we incorporate uncertainty estimation by integrating both aleatoric and epistemic uncertainty using stochastic training loss [16] and deep ensemble [17], enhancing the predictive reliability of our model. Unlike earlier methods that provide only point estimates, our approach also quantifies prediction confidence, capturing both inherent data noise and model uncertainty. UniversalEPI is evaluated based on its ability to predict cell-type-specific Hi-C interactions between accessible genomic regions in four human cell lines. As expected, UniversalEPI demonstrates superior performance to methods that rely

solely on DNA sequence; it performs comparably to or outperforms the state-of-the-art C.Origami method even when the latter uses experimental CTCF affinities in addition to bulk chromatin accessibility for inference. As further validation, we show the applicability of our method in assessing chromatin dynamics during cell reprogramming on the macrophage activation data generated by Reed *et al* [18]. Finally, we benchmark the ability of UniversalEPI to predict Hi-C from pseudo-bulks of single-cell chromatin accessibility profiles and highlight the differences in chromatin dynamics across different transcriptional states in esophageal adenocarcinoma.

2 Results

2.1 UniversalEPI: a transformer-based model to predict genomic interactions from DNA sequence and chromatin accessibility

We developed UniversalEPI, a two-stage deep learning model that accurately predicts chromatin interactions across cell types. The two sequentially trained neural networks of UniversalEPI use DNA sequences and signals from assays for transposase-accessible chromatin using sequencing (ATAC-seq) as input and chromatin immunoprecipitation followed by sequencing (ChIP-seq) and Hi-C data as target (Figure 1a). To infer the Hi-C signal in a test cell type, UniversalEPI requires DNA sequence and cell-type-specific ATAC-seq profiles as input (Figure 1b).

The first neural network of UniversalEPI was trained to predict the genome-wide binding occupancy of TFs ubiquitously involved in promoter activation and chromatin looping: maximal intensities of SP1, CTCF, and YY1 ChIP-seq [13–15]. This was achieved by training a one-dimensional 5-layer CNN producing a 4-dimensional output on the concatenation of one-hot encoded 1Kb DNA sequences and its corresponding ATACseq p-value signals (Methods). Since the orientation of CTCF plays an important role in determining the stability of topologically associated domains (TADs), we split the CTCF ChIP-seq based on orientation (forward or reverse) and predicted the corresponding maximal ChIP-seq intensities as separate targets. The pre-trained convolutional layers from the first stage thus generated

sequence embeddings, indicating the presence or absence of the three TFs in each accessible region in the genome; this design of sequence embeddings was chosen to prevent capturing DNA motifs of other, potentially cell-type-specific TFs and enable the model to generalize predictions across different cell types at the second stage.

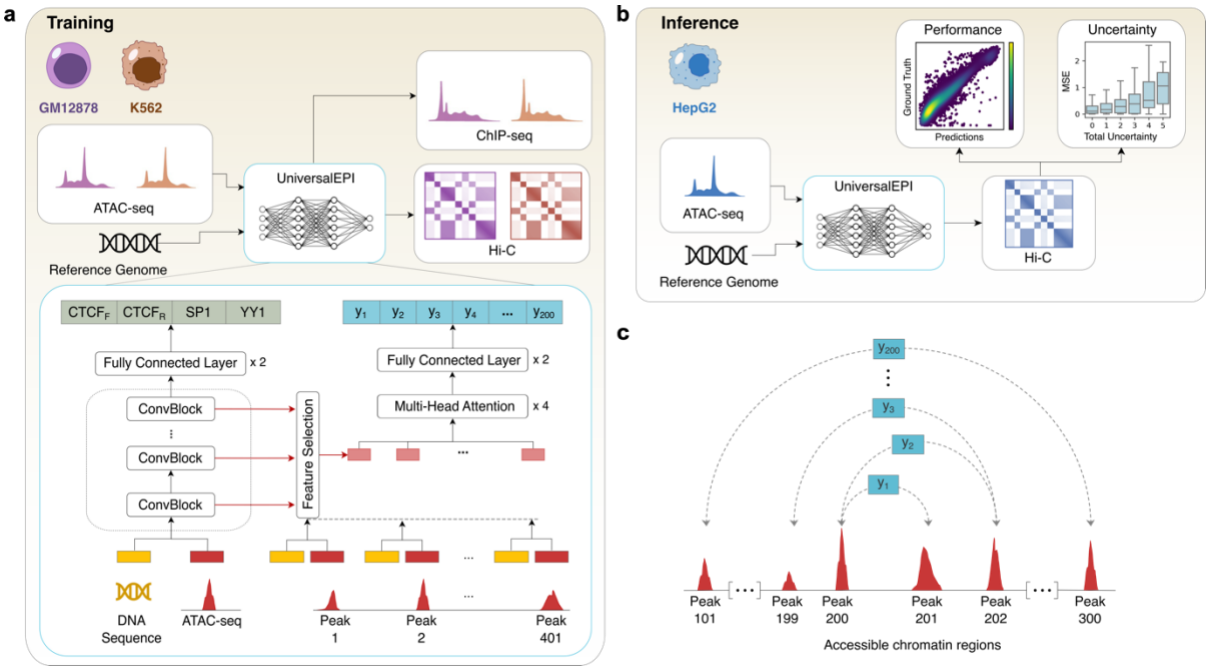

Figure 1. Outline of the UniversalEPI model architecture. a, UniversalEPI uses one-hot encoded DNA sequence and ATAC-seq p-value tracks as input to predict cell-type-specific transcription factor binding and chromatin interactions. It can be trained on data from multiple cell lines. The first stage of UniversalEPI consists of a convolutional architecture that predicts transcription factor binding affinity from DNA sequence and ATAC-seq. The second stage automatically selects the features from the pre-trained convolutional layers from the first stage using a stochastic gating mechanism. The multi-head attention block then uses these features to predict cell-type-specific Hi-C interactions between accessible regions. **b,** The trained UniversalEPI can then be used to predict Hi-C interactions for an unseen cell type. The model also outputs uncertainty associated with each prediction, which increases with the mean squared error (MSE) of the model. **c,** UniversalEPI predicts pairwise Hi-C interactions (y_i) between the central 200 accessible regions in a zigzag manner.

In the second stage, we trained the final model to predict Hi-C values between accessible regions based on the DNA sequence embeddings calculated in the first stage (Figure 1c). We employed a transformer-based encoder because of its ability to capture long-range interactions and create large-context embeddings of genomic elements (Methods). This allowed us to efficiently learn the inter-dependency

between the input regions and, consequently, predict the Hi-C interactions between regulatory elements with high accuracy in a cell-type-specific manner. Specifically, the model input in the second stage corresponded to 401 consecutive accessible regions (or peaks) spanning in total about 4Mb. Each region was 1Kb long, centered around the position of the maximum ATAC-seq signal. Skipping the information about the DNA sequence in between accessible regions allowed for a rather lightweight architecture of UniversalEPI (2.5M parameters) while maintaining a large receptive field and making accurate predictions for interactions between regulatory elements located up to 2Mb apart from each other. We provided the total predictive uncertainty as a combination of aleatoric (data) and epistemic (model) uncertainty. The former was estimated by minimizing the negative log-likelihood loss to incorporate the predicted variance, while the latter was quantified based on the deep ensemble (Methods).

To ensure the stability of the proposed architecture, we trained both stages of UniversalEPI in a cross-validation setting. Model 1 was trained on GM12878 and K562 cell lines and tested on unseen chromosomes of cell lines IMR90 (Hi-C), A549 (ChIPseq), and HepG2 (Hi-C and ChIP-seq). Model 2 was trained on IMR90 and HepG2 cell lines and tested on unseen chromosomes of GM12878 and K562 (Hi-C and ChIP-seq). Approximately 15% of the genome was set aside for validation (chromosomes 5, 12, 13, and 21) and approximately 15% for testing (chromosomes 2, 6, and 19). The first stage of UniversalEPI predicted the transcription factor binding on unseen chromosomes and unseen cell types for CTCF, YY1, and SP1 with an average Pearson's correlation of 0.78, 0.64, and 0.51, respectively (Figure S1). We further validated that the first stage of our model learned biologically relevant information by analyzing the gradients using the DeepLIFT attributions scores [19]. The first stage of UniversalEPI captured the binding motif of all three TFs on the test set (Figure S1).

2.2 UniversalEPI learns information about biologically-relevant features

Experimental Hi-C values, quantifying chromatin interactions between each pair of genomic regions, are inversely correlated with the distance between regions and strongly follow the power law resulting

from polymer physics [20]. Similarly, UniversalEPI's predictions followed the power law on the test chromosomes of the unseen cell lines, highlighting the model's ability to accurately capture the effect of distance on chromatin structure (Figure 2a). Further, to show that UniversalEPI was sensitive to the variability in the ATAC-seq signal, we computed the distance-stratified correlation between the predicted and ground-truth Hi-C signal between pairs of accessible regions (Methods). Using as inputs cell-line-specific ATAC-seq profiles, UniversalEPI obtained equally high values of distance-stratified correlation for all four cell lines (Figure 2b).

At last, we assessed the sensitivity of the UniversalEPI model to changes in the DNA sequence through in silico inversions of CTCF binding sites located at the TAD boundaries. TADs influence which cis-regulatory elements interact with gene promoters and are typically bordered by the CTCF protein bound in the convergent orientation [21]. Disruptions in TADs have been linked to diseases like cancer, making the information about chromatin interactions within and across TADs essential for exploring disease mechanisms [22–24]. In silico inversions of one of the endpoints of a TAD boundary containing a CTCF site in a convergent orientation were expected to reduce the strength of chromatin interactions between the original TAD boundary (Figure 2c). Indeed, when we provided the original and modified inputs to UniversalEPI, we predominantly observed negative values of the log-fold change across the four test cell lines (Figure 2d), indicating a decrease in the predicted interaction between the original TAD boundaries after the inversion.

We further illustrated the effect of in silico TAD boundary inversion on the predictions of UniversalEPI with a randomly selected TAD in IMR90 (chr6:17,280,000-17,585,000). UniversalEPI predicted a strong reduction in interactions for the TAD boundary (highlighted) as well as between the regions within the original TAD (Figure 2e). Moreover, we observed a significant increase in predicted interactions with a reverse CTCF site located nearly 200Kb away, outside of the original TAD boundary, and a new interaction formed across the original TAD with the promoter of the *NUP153* gene. This result

confirmed that UniversalEPI was capable of capturing the effects of CTCF orientation on the stability of TADs.

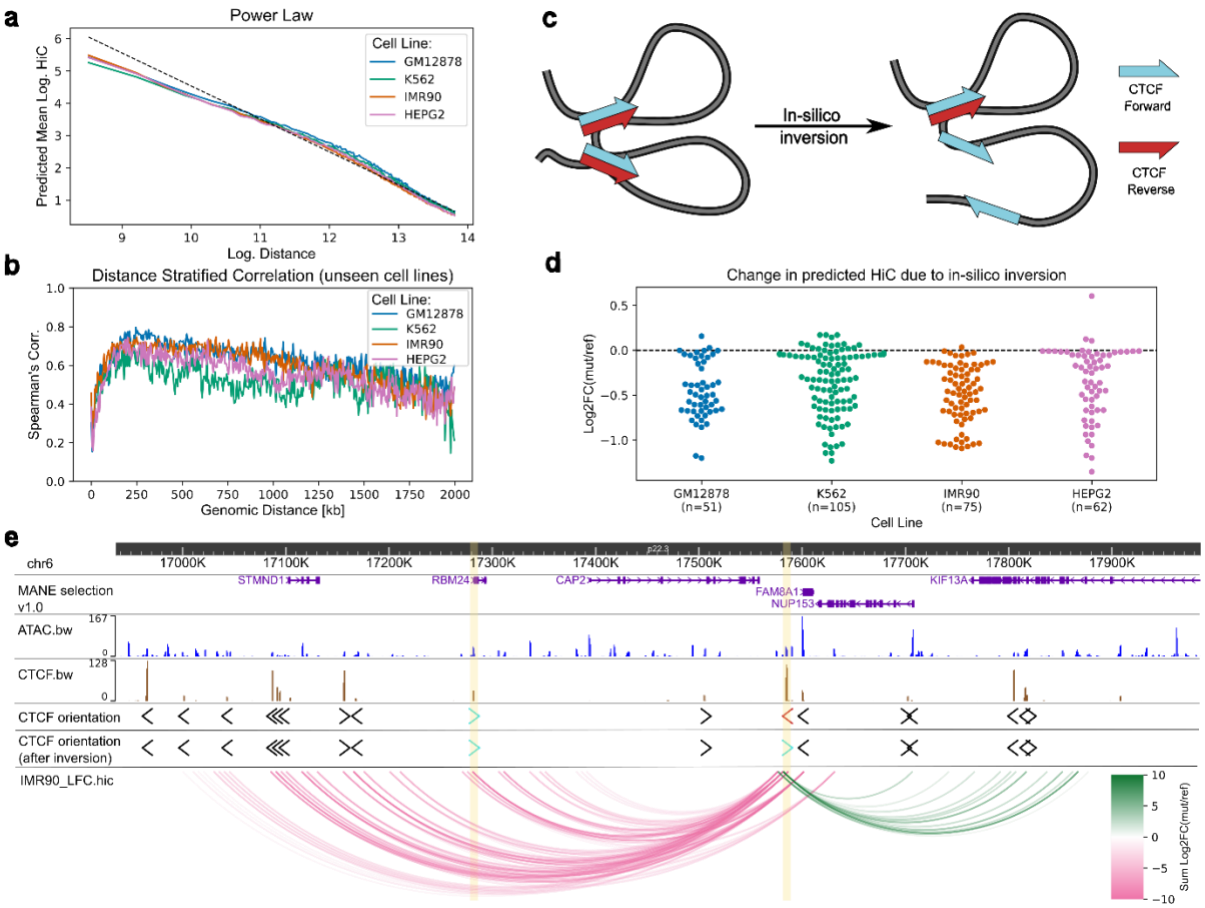

Figure 2. UniversalEPI learns information from biologically-relevant variables. a, The predicted Hi-C signal by UniversalEPI follows the power law resulting from polymer physics. The average predicted signal is shown for test chromosomes of test cell lines (*i.e.*, cell lines unseen during training): GM12878, K562, IMR90, and HepG2. The theoretical relationship between the log-distance and log-Hi-C values is depicted as the dotted line. **b,** Distance-stratified Spearman's correlation for all the unseen cell lines displaying UniversalEPI's sensitivity to the input ATAC-seq. **c,** Graphical illustration depicting the effect of in-silico inversion of a CTCF site on genome organization. **d,** Log fold change between the UniversalEPI predictions before and after the in-silico inversion. The experiment is done for all TADs in test chromosomes in unseen cell lines: GM12878, K562, IMR90, and HepG2. **e,** An example of UniversalEPI predictions for a TAD in IMR90 in chromosome 6 (unseen cell line and unseen chromosome). The original TAD is highlighted in yellow. The change in predicted Hi-C is calculated using $\log_2$ fold change (logFC) between before and after the in-silico inversion. In the presence of multiple ATAC-seq peaks within the same 5Kb bin, the logFC is summed among all pairs of interactions contained in the two bins.

2.3 Benchmark of the UniversalEPI method on the unseen cell types

Since predicting chromatin interactions across cell types based on cell-type-specific inputs is the main goal of this work, we compared UniversalEPI against the existing methods in terms of their ability to predict chromatin interactions in cell types not used for model training. We also introduced three baselines that captured the effect of distance between interacting elements and the variability between the train and test Hi-C interaction matrices (Methods). We calculated method performances for all interactions within 1Mb regions, which was the maximum common receptive field across all existing methods. We observed that the performance of UniversalEPI was comparable to the original version of the C.Origami method, which used CTCF ChIP-seq as an additional input (Figure 3a and Figure S2). When trained without CTCF as input, C.Origami showed significantly worse performance than UniversalEPI (Figure 3a). UniversalEPI also outperformed the sequence-only method Akita as this, by design, was not expected to generalize across cell types. When stratified by cell lines, UniversalEPI consistently obtained a Spearman's correlation between predicted and ground-truth Hi-C values above 90% (Figure 3b). On IMR90 and K562, UniversalEPI achieved comparable results to C.Origami, and it substantially outperformed C.Origami on HepG2 and GM12878 cell lines.

Since distance between the interacting elements is a major driving factor in determining the strength of interaction, we also compared the existing methods after removing the distance effect. This was done by calculating the distance-stratified correlation (Figure 3c), which assessed the ability of tested methods to capture biological features across several distance ranges. We observed that UniversalEPI captured the biological information better than all the existing methods, including C.Origami. The possible reason of that can be that C.Origami, being an end-to-end model, which, in contrast to UniversalEPI, does not limit the model's visual scope to specific TF motifs, learned cell-type-specific information that is not transferable to all cell types.

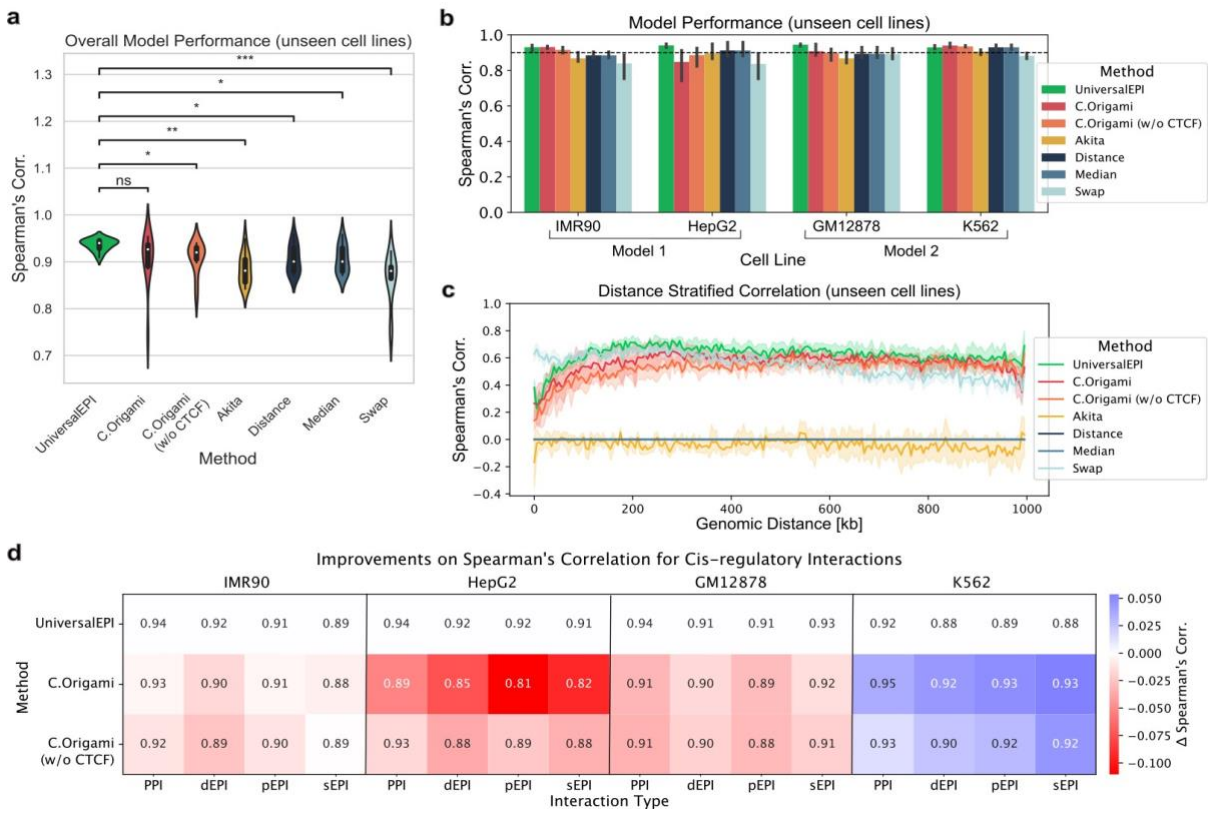

Figure 3. Benchmark of the UniversalEPI method on the unseen cell types. a, UniversalEPI predicts chromatin interactions in unseen cell types with comparable performance to C.Origami, which additionally uses CTCF ChIP-seq as input. UniversalEPI outperforms C.Origami when the latter uses the same inputs and Akita, which makes predictions solely based on DNA sequence. Each point corresponds to Spearman's correlation coefficient between predictions and ground-truth HiC for an unseen chromosome in an unseen cell line, GM12878, K562, IMR90, and HepG2. A Wilcoxon signed rank test is used to compare the methods; ***: $P < 0.001$, **: $P < 0.01$, *: $P < 0.05$. **b,** UniversalEPI is the only method that achieves Spearman's correlation greater than 0.9 (depicted by a dotted line) for all unseen cell types. For each method, Model 1 is trained on GM12878 and K562 cells whereas Model 2 is trained on IMR90 and HepG2 cells. **c,** Distance-stratified correlation shows that UniversalEPI captures the biological information better than other methods. Each line demonstrates the mean distance-stratified Spearman's correlation across four cell lines; the shaded area marks the 95% confidence interval. **d,** UniversalEPI outperforms C.Origami and its variant without CTCF ChIP-seq in predicting different types of cis-regulatory interactions in all but the K562 cell line. The change in Spearman's correlation is calculated with respect to UniversalEPI. Positive values suggest that C.Origami or its variant have higher scores than UniversalEPI. PPI: Promoter-Promoter Interaction, dEPI: distal Enhancer-Promoter Interactions (further than 2Kb), pEPI: proximal Enhancer-Promoter Interactions (within 2Kb), and sEPI: Super Enhancer-Promoter Interactions.

Finally, we compared the top three methods – UniversalEPI and C.Origami (with and without CTCF ChIP-seq used as input), on their ability to predict enhancer-promoter interactions (EPIs) and promoter-promoter interactions (PPIs). To study the ability of the models to capture distance-dependent interactions, we split EPIs into two categories – proximal EPI (pEPI, 200bp–2Kb) and distal EPI (dEPI, >2Kb). The annotations for pEPIs, dEPIs, and PPIs were obtained from ENCODE [25]. We also compared the three methods based on their ability to predict interactions between promoters and super-enhancers, corresponding to clusters of active enhancers often regulating cell identity genes [26] and being defined for each cell line in the dbSuper database [27] (sEPIs). UniversalEPI outperformed C.Origami and its counterpart without CTCF binding in the input on all but the K562 cell line (Figure 3e), thereby better capturing important cis-regulatory interactions. Overall, we concluded that DNA sequence and ATAC-seq were sufficient to accurately predict cell-type-specific chromatin interactions.

2.4 UniversalEPI captures chromatin dynamics during cell reprogramming

To validate the applicability of UniversalEPI for predicting changes in regulatory interactions upon cell reprogramming, we analyzed the model's predictions upon reprogramming of macrophages, for which ATAC-seq and Hi-C data were collected at eight different time points following the macrophage activation with a mixture of Lipopolysaccharide (LPS) and Interferon gamma (IFN_{γ}) [18] (Figure 4a). In a zero-shot setting, our model trained on GM12878 and K562 cell lines demonstrated high prediction accuracy, with Pearson's and Spearman's correlation with ground-truth Hi-C values exceeding 90% (Figure 4b).

Reed *et al.* identified specific chromatin loops that exhibited differential interactions upon macrophage activation—lost, gained, or remaining static (Figure 4c). For these selected chromatin loops, UniversalEPI effectively identified treatment-induced changes in chromatin architecture. This was measured by the maximum-confidence log2-fold change between the Hi-C interactions predicted at 24h and 0h timepoints which also takes into account the uncertainty of prediction (Figure 4d,

Methods). Moreover, UniversalEPI's predictive power allowed it to distinguish gained and lost loops from static loops with a performance level that is at least five times more accurate than random classification (Figure 4e).

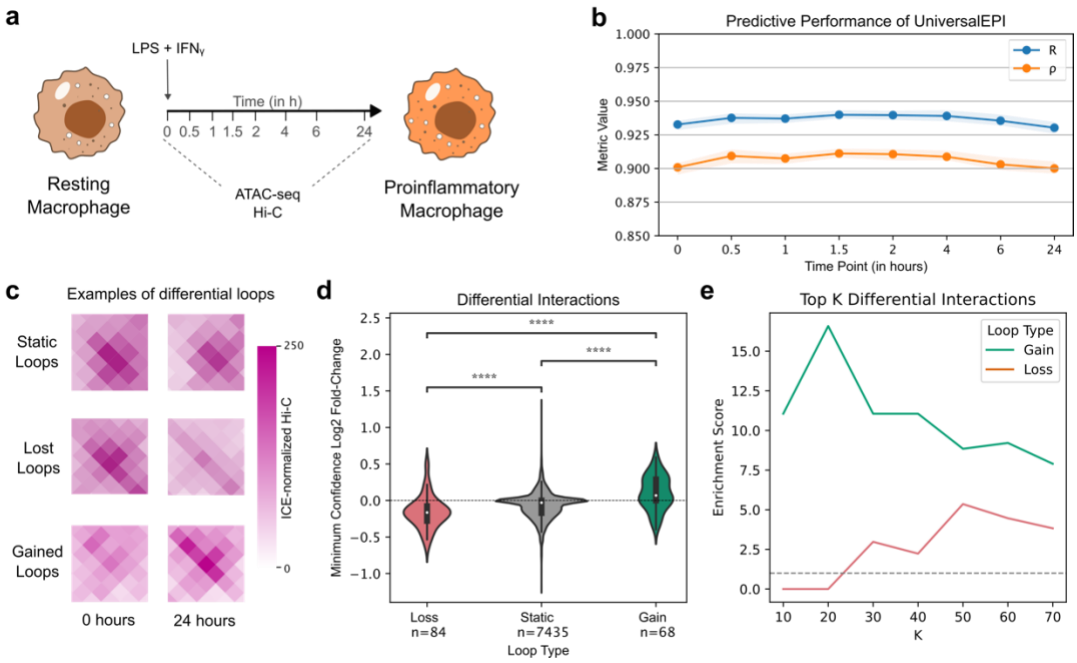

Figure 4. UniversalEPI captures chromatin dynamics during cell reprogramming. a, Experimental layout defined in Reed *et al*. [18]. Macrophages derived from the THP-1 monocytic cell line are activated using LPS + IFN_γ treatment. ATAC-seq and Hi-C are measured at 8 time points. **b,** Performance of UniversalEPI (trained on GM12878 and K562 cells) using ATAC-seq for each of the 8 time points, quantified by Pearson's (*R*) and Spearman's (*ρ*) correlation. The dot represents the mean and the shaded area represents the range across 10 ensembles. **c,** Examples of static, lost, and gained loops in the chromatin between 0 and 24 hours identified by Reed *et al*. **d,** The maximum-confidence log2-fold change between Hi-C values at the 24-hour time point over the 0th time point measured by UniversalEPI predictions for the differential loops. The significance is measured using a Mann-Whitney-Wilcoxon test with Bonferroni correction. ****: $P \leq 10^{-4}$. **e,** Gain and loss enrichment score over random for selecting a gained or lost loop respectively in the K most differential interactions as predicted by UniversalEPI. The blue dotted line represents the enrichment score of labelling a random interaction as gained or lost.

2.5 UniversalEPI unravels chromatin structure in complex tissues using pseudo-bulk single-cell ATAC-seq data

To show the applicability of our computational method to single-cell studies, we used UniversalEPI to predict the chromatin structure using as inputs pseudo-bulks of single-cell ATAC-seq (scATAC-seq) data (Figure 5a). We compared UniversalEPI with ChromaFold [12], the only other method that predicts bulk Hi-C from scATACseq and DNA sequence as inputs. scATAC-seq and Hi-C data from GM12878, K562, HepG2, and IMR90 cell lines were used in this analysis. Healthy and cancer cell lines GM12878 and HepG2 cells were employed for training both models. Keeping the study design similar to the one proposed by Gao *et al.* [12], we used Hi-C at 10Kb resolution with HiC-DC+ z-scores as targets [28].

The two methods were compared based on their ability to predict Hi-C values on unseen chromosomes in the unseen cell lines IMR90 and K562. UniversalEPI outperformed ChromaFold on both unseen cell lines (Figure 5b). UniversalEPI also obtained a higher distance-stratified correlation, especially for short-to-medium-ranged interactions (50-600Kb) (Figure 5c), highlighting the ability of our model to maintain its generalization capacity while extending to pseudo-bulk scATAC-seq data.

Further, we used the multiome data (scATAC-seq and scRNA-seq) from 8 esophageal adenocarcinomas (EAC) [29] to assess potential differences in chromatin structure between two cancer transcriptional states: differentiated and undifferentiated malignant cells, expressing cNMF4 and cNMF5 programs respectively described in Yates *et al.* [29]. Pseudo-bulk ATAC-seq profiles were obtained from the merging scATAC-seq data of all cells across patients expressing a given program. These pseudo-bulk ATAC-seq profiles were then used to predict Hi-C interactions between accessible regions using the UniversalEPI model pre-trained on GM12878 and K562 cell lines (Figure 5d).

Based on predicted Hi-C data, we defined promoter activity for each gene and used it to juxtapose with measured gene expression. Promoter activity was derived as a simple function of ATAC-seq promoter signal, ATAC-seq enhancer signals, and the predicted Hi-C interaction strengths between the

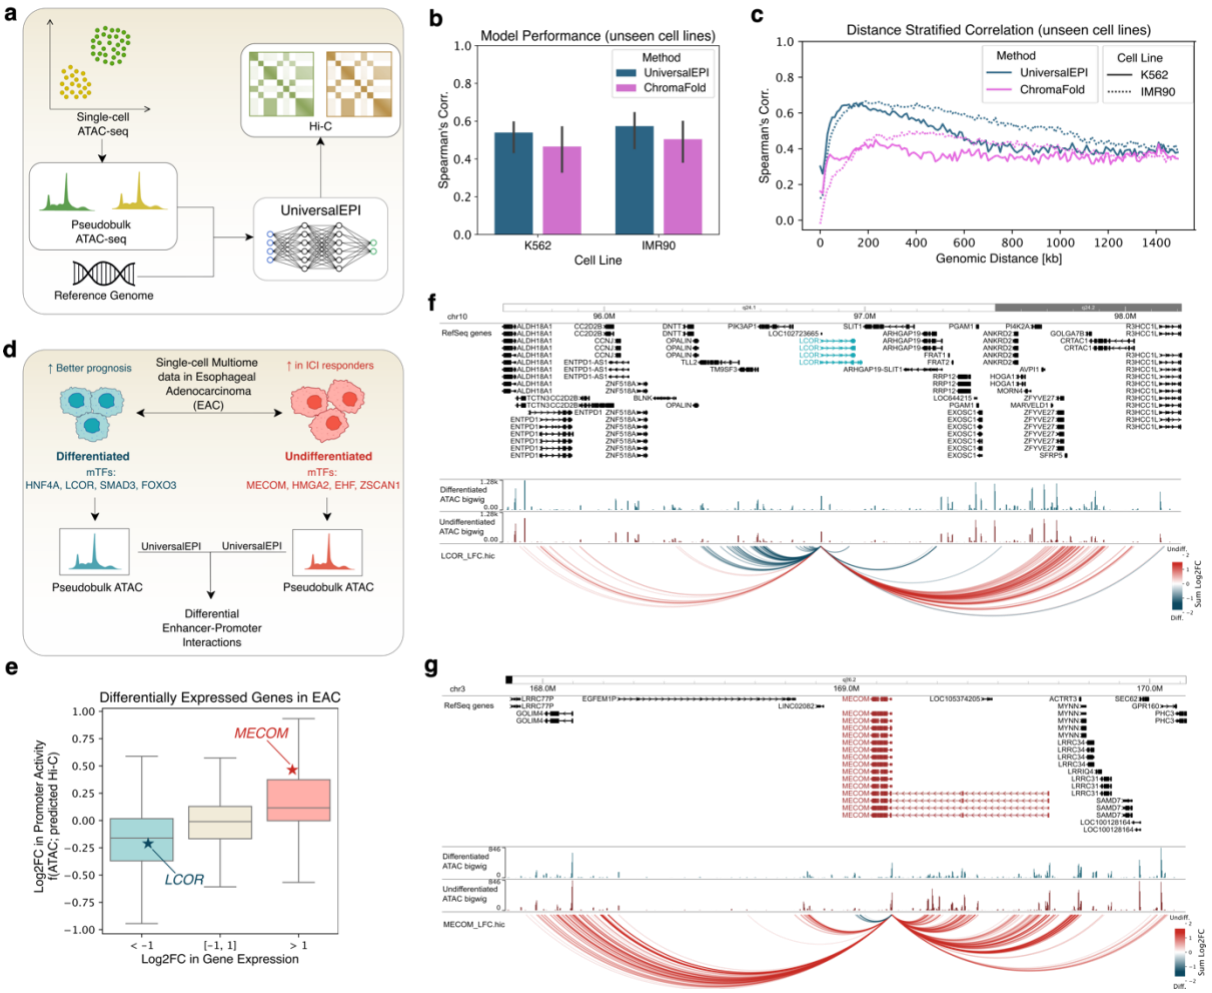

Figure 5. UniversalEPI unravels chromatin structure in complex tissues using pseudo-bulk single-cell ATAC-seq data. a, UniversalEPI can predict cell-type-specific Hi-C from different cell types identified using scATAC-seq data. **b,** Comparison between UniversalEPI and ChromaFold on their ability to predict bulk Hi-C from scATAC-seq. Spearman's correlation is calculated on unseen chromosomes (chr5, chr18, chr20, chr21) of unseen cell lines. **c,** Distance-stratified Spearman's correlation based on the unseen chromosomes and cell lines. **d,** The multiome data (scRNA-seq and scATAC-seq) of differentiated and undifferentiated malignant cells are obtained from esophageal adenocarcinoma (EAC) tumors [29]. Master transcription factors (mTFs) for each program are identified in Yates *et al.* The scATAC-seq data from the candidate cells expressing each program are merged from multiple patient tumors to obtain program-specific pseudobulk ATAC-seq profiles. These are then used as inputs to the pretrained UniversalEPI model to obtain Hi-C interactions for each program. **e,** Log2-fold change between bulkified gene expression profiles in undifferentiated and differentiated cells is compared against the log2-fold change in estimated promoter activity. Promoter activity acts as a representative for gene expression and is defined as a function of ATAC-seq and the predicted Hi-C. The blue star highlights the LCOR gene, which is a mTF for differentiated cells, whereas the red star highlights the *MECOM* gene, which is a mTF for undifferentiated cells. The whiskers of the box plot show the range with ouliers removed. **f-g,** The change in predicted Hi-C

interactions between accessible regions, measured by the sum of log2-fold change between undifferentiated and differentiated cells, are shown for the *LCOR* and *MECOM* gene promoters.

promoter and enhancers (Methods). The derived promoter activity not only showed a high correlation with gene expression but also a positive correlation (Pearson's R = 0.224) with differential gene expression between undifferentiated and differentiated cells in EAC (Figure 5e, Figure S11). Yates *et al.* also identified master transcription factors (mTFs) for each program, including *LCOR* and *MECOM* in differentiated and undifferentiated EAC cells, respectively [29]. These TFs play a pivotal role in defining and maintaining malignant cell identity. Specifically, *LCOR* drives tumor cell differentiation and reduces tumor growth while low *LCOR* expression is associated with suppression of antigen processing and presentation, enabling resistance to immune checkpoint blockade (ICB) therapy in breast cancer [30, 31]. Conversely, overexpression of *MECOM* has been linked with activating oncogenes [32], promoting stem-like properties, and reducing apoptosis in different cancers [33, 34]. Our predicted Hi-C-based promoter activity was in agreement with these findings showing higher activity of *LCOR* in differentiated cells relative to the undifferentiated cells and vice versa for *MECOM* (Figure 5e).

Furthermore, we studied the change in chromatin structure for these two mTFs between the two programs. Stronger interactions were observed between the *LCOR* gene promoter and proximal enhancers (within 500Kb) in differentiated cells whereas distal enhancers (approximately 1Mb) regulated this gene in undifferentiated cells (Figure 5f). Similarly, *the MECOM gene promoter formed strong interactions with enhancers in its neighborhood,* thereby potentially activating the gene in undifferentiated cells. Only a small set of proximal enhancers showed stronger interaction with the gene promoter in differentiated cells, highlighting the different looping mechanisms that regulate the expression of *MECOM* (Figure 5g).

3 Discussion

In this work, we presented UniversalEPI, a novel attention-based model that accurately predicts chromatin interactions on unseen cell types without retraining. One major drawback of current approaches, which is requiring several data modalities as inputs, was solved by UniversalEPI operating on DNA sequence and chromatin accessibility data. Furthermore, since our model learns only from accessible regions of the genome, it demonstrates notable scalability (Figure 1). In contrast to other techniques, such as C.Origami, which was originally trained on a single cell line [11], this enables our model to be trained on several cell lines concurrently in a fair amount of time.

Using CNN layers and transformer blocks, UniversalEPI effectively captures the essential transcription factor binding motifs and spatial dependencies underlying chromatin interactions. The model's high prediction accuracy demonstrated by the strong distance-stratified correlation between the predicted and ground-truth Hi-C values confirms its sensitivity to biologically relevant information, which is critical for accurate modeling of chromatin architecture (Figure 2). This enables UniversalEPI to consistently achieve high correlations between predictions and experimental Hi-C values and even outperform existing models, including C.Origami, for unseen cell types (Figure 3).

By applying UniversalEPI to predict treatment-induced chromatin changes in macrophages stimulated with LPS and INF_{γ}, we were able to further validate the model performance in a zero-shot setting (Figure 4). The model's ability to detect dynamic changes in chromatin organization, which can be crucial in examining how different cell types respond to drugs and environmental stimuli, was demonstrated by its accurate identification of differential chromatin loops. UniversalEPI also displayed high performance when predicting bulk Hi-C using pseudo-bulk single-cell ATAC-seq data (Figure 5b-c). Furthermore, different looping patterns were identified by UniversalEPI for master transcription factors during malignant cell differentiation in esophageal adenocarcinoma (Figure 5e-g). This feature

increases the method's utility for providing insights into the heterogeneity in chromatin organization observed in complex tissues or across cell states.

Despite showing robust performance across cell types, UniversalEPI was only trained on a limited number of cell lines. Therefore, additional training may be required for its use in rare and highly specialized cell types with unique regulatory mechanisms. However, this issue can be mitigated in the near future with the increasing production of ChIP-seq and Hi-C data in new cell types. Additionally, the first stage of UniversalEPI could be further improved by applying techniques like adversarial training [35–37], preventing the model from capturing cell-type-specific information and, consequently, enhancing its ability to generalize across cell types.

In conclusion, UniversalEPI achieves state-of-the-art performance in predicting chromatin interactions in unseen cell types from only DNA sequence and chromatin accessibility profiles. This brings a major advancement in chromatin interaction predictions that can be widely applied to study complex regulatory landscapes of the genome. Its predictive power would allow UniversalEPI to be used for studying changes in chromatin organization in the presence of non-coding mutations and structural variations and thereby, deciphering the regulatory mechanisms of genetic diseases. Finally, the computational architecture of UniversalEPI sheds light on the chromatin folding mechanisms between promoters and cis-regulatory regions: It demonstrates top performance while discarding all information about non-accessible chromatin and cell-type-specific TF binding motifs.

Methods

Hi-C data processing

The raw hg38-aligned Hi-C files for four cell lines (GM12878, K562, IMR90, and HepG2) were obtained directly from the 4D Nucleome Data Portal [38, 39] under the accession codes listed in Supplementary Table 1. We then extracted the 5Kb resolution contact matrices. To remove any systematic biases, we applied ICE normalization [40] using HiCExplorer v2.2.1 [41] with filter thresholds of −1.1 and 4.5.

Since we are only interested in intra-chromosomal contacts, the ICE normalization was also done independently on each chromosome of each cell line.

The Hi-C contact matrices must be normalized across cell lines for the model predictions to be comparable to the training data. We achieved this by applying a distance-stratified robust z-score normalization using the healthy GM12878 cell line as a reference. GM12878 was selected as a reference cell line due to the high sequencing depth of its Hi-C matrix and hence, better data quality (approximately 759M interactions as compared to approximately 275M interactions in the other cell lines). For each chromosome of the GM12878 cell line, we stored the median and median absolute deviation (MAD) of all measured Hi-C interactions corresponding to regions that are *d* bins apart where $d \in [0,800]$ (each bin being 5Kb of the genome). In some chromosomes, there were too few long-range interactions for accurate calculations of median and MAD. Hence, we fitted a B-spline, using SciPy's splrep function [42], setting the smoothing condition *s* to 5. This spline was then used to smooth our median and MAD curves across the genomic distance. We denote the smoothed median and MAD as $m_d^{GM12878}$ and $mad_d^{GM12878}$ respectively. For a new cell type, e.g., the K562 cell line, the normalized Hi-C score ($\overline{y}$) between regulatory elements in bins *i* and *j* is then calculated as follows:

$$\bar{y}_{i,j}^{\mathrm{K562}} = \left(\frac{y_{i,j}^{\mathrm{K562}} - \mathrm{m}_{|j-i|}^{\mathrm{K562}}}{\mathrm{mad}_{|j-i|}^{\mathrm{K562}}} \right) \mathrm{mad}_{|j-i|}^{\mathrm{GM12878}} + \mathrm{m}_{|j-i|}^{\mathrm{GM12878}}$$

where *y* is the ICE-normalized Hi-C value. Figure S3 depicts the effect of this normalization on all cell lines.

Fundenberg *et al* [4] applied Gaussian smoothing with unit variance to fill in the missing interactions and reduce the noise of the Hi-C experiment. However, this can lead to a significant reduction in strong Hi-C contacts in the presence of missing values at shorter distances. To mitigate this issue, we applied Gaussian smoothing with distance-dependent variance, which increased with distance (Figure S7). Particularly, we used four kernels with variances of 0.5, 0.65, 0.8, and 1 for interactions between

regions that are closer than 25Kb, between 25Kb and 100Kb, between 100Kb and 250Kb, and greater than 250Kb respectively. Finally, we applied logarithmic transformation to the interaction scores. Several methods like Akita [4] and ChromaFold [12] also remove the effect of distance from the model by applying z-scores. While this approach highlights long-range interactions, we aim to capture the true biological interactions and hence, retain the effect of distance in our Hi-C matrices.

ATAC-seq data processing

In this work, we used the GRCh38/hg38 human reference genome. ATAC-seq data from five human cell lines was used in this study consisting of healthy (GM12878 and IMR90) and cancer (K562, A549, and HepG2) cells. The carefully processed signal p-value bigwig profiles and the pseudoreplicated peak files were downloaded directly from the ENCODE portal (http://www.encodeproject.org/) under the accession codes listed in Supplementary Table 1. For multiple ATAC-seq peaks within the 500bp genomic region, we retain the peak with the highest read count and remove the remaining peaks. This resulted in approximately 175K peaks for each cell line.

To ensure consistency between the ATAC-seq bigwig profiles of different cell types, we normalized the bigwig signal using GM12878 as the reference cell line. First, we identified a conserved set of CTCF sites by selecting the common peaks from CTCF ChIP-seq tracks of all five cell lines. For each identified peak, we then obtained the signal values from bigwigs of each cell line using deepTools multiBigwigSummary v3.5.3 [43]. Finally, we applied the Trimmed Mean of M-values (TMM) normalization of EdgeR [44] using the extracted ATAC-seq signal values on these conserved CTCF sites to obtain the scaling factor for each cell line. GM12878 cell line was used as a reference while TMM normalization was applied. The normalized bigwig was then constructed by multiplying the scaling factor with the original signal track. The effect of normalization for all the cell lines can be observed in Figure S4.

ChIP-seq data processing

The ”narrow peak” files of each TF from the list of target TFs (CTCF, YY1, and SP1) for all five cell lines (GM12878, K562, IMR90, HepG2, and A549) were obtained from the ENCODE portal. The accession codes can be found in Supplementary Table 1. To account for the role of CTCF orientation in the formation and stability of TADs, we split CTCF ChIP-seq peaks based on their orientation. Specifically, we utilized the consensus motif MA0139.1 from the JASPAR database [45] to derive a position-specific scoring matrix, which was subsequently convolved over the forward and reverse strands of the reference genome. To determine the optimal cutoff for the resulting convolved signal, we employed a false negative rate of 1%, based on empirical estimates [46] of the intergenic nucleotide frequency background distribution. Signals larger than the cutoff indicated the presence of the consensus motif on the respective strand. Since the resolution of the ChIP-seq signal was lower than the motif length, a peak was labeled as *forward* and *backward* facing if both signals exceeded the cutoff value (18% of all peaks). Overall, 49% of the CTCF peaks were assigned to a motif on the forward strand and 48% to a motif on the reverse strand. For 20% of the peaks, no corresponding strand could be determined. Since these peaks were also highly correlated with low enrichment, they were dropped from the dataset during training.

Model architecture

The architecture of UniversalEPI only accounts for the information coming from accessible chromatin regions (ATAC-seq peaks); all information about genomic regions between accessible regions is provided to the model via the value of the distance between ATAC-seq peaks. The model, therefore, predicts the entries in the interaction matrices only for the pairwise interactions between genomic regions containing ATAC-seq peaks. We propose a two-stage deep learning model to learn the mapping $F: \mathbb{X} \mapsto \mathbb{Y}$

where $\mathbb{X}$ is the space of genomic information about ATAC-seq peaks and $\mathbb{Y}$ is the space of interactions. Namely, for each $X \in \mathbb{X}$, $X = \{x_1, x_2, .., x_{401}\}$ with $x_i \in \mathbb{R}^{1000\times5}$ being stacked one-hot encoded DNA sequences and ATAC-seq signal of one ATAC-seq peak, the mapping should produce $Y = \{y^{(i,j)} \mid y^{(i,j)} \in \mathbb{R}, i \leq j, i,j = 101,...,300\} \in \mathbb{Y}$ with elements representing the interaction between the pair of peak locations i and j in the central 200 ATAC-seq peaks. Note that the entry $y^{(i,j)}$, which represents the interaction between the pair of peaks locations i and j, depends not only on the data at those locations but also on the information across other peaks in their neighborhood.

The first stage learns genomic features for the peak regions, relevant to predicting generalized chromatin features across cell types, whereas the second stage predicts HiC values by modeling the interactions between the learned features at peak locations. A schematic overview of the proposed method is depicted in Figure S5. Below, we describe the two stages in detail.

Stage 1: Representation Learning

In the first stage, we train a representation network $f_\vartheta(x)$, where $x \in \mathbb{R}^{1000\times5}$, to predict binding affinity of the target TFs $y_{\mathrm{TF}} \in \mathbb{R}^{N_{\mathrm{TF}}}$ using the data across different cell lines, where N_{TF} is the number of the target TFs. This approach ensures that the learned features capture information that is both invariant across different cell lines and relevant for predicting the target Hi-C. Inspired by DeepC [3], we adopt a convolution-based model for this purpose. The model consists of five convolutional layers with {30,60,60,90,90} channels, and kernel sizes of {11,11,11,5,5} respectively. After each convolutional layer, max pooling with widths of {4,5,5,4,2} is applied to aggregate the learned features, which is followed by Leaky Rectified Linear Unit (LeakyReLU) [47] with a slope of 0.2 to introduce non-linearity. A fully connected layer takes the output of the convolutional layers and maps to the target TFs. To prevent over-fitting, a 20% dropout is applied during training. The model is optimized by minimizing the empirical version of the mean squared error (MSE) defined as

$$\min_{\theta} \mathbb{E}_{(x,y_{\mathrm{TF}})\sim P_{\mathrm{TF}}}[||f_\theta(x) - y_{\mathrm{TF}}||^2], \qquad (1)$$

where paired x and y_{TF} are assumed to be sampled from a joint data distribution P_{TF}.

The use of convolutional layers enables the network to effectively extract hierarchical features from the input data. After applying the non-linear activations, the outputs of convolutional blocks serve as intermediate representations and are used as the input for the second stage. We use a linear projection head for each feature layer to map it to the space of the dimension C, which is set to 180. For one peak x, we obtain a set of N_L projected features $h(x) = \{h_1(x), h_2(x), ..., h_{NL}(x)\}$, $h_i(x) \in \mathbb{R}^C$, where $N_L = 5$ is the total number of convolutional blocks. Thus, for the subsequent 401 ATAC-seq peaks, we get the stacked projected features of the dimension $401 \times C \times N_L$, which we denote by $h(X)$. The feature selection module is then employed to select from N_L features, as described in details below.

Stage 2: Hi-C Prediction

• *Automatic feature selection* is enabled by adopting the stochastic gating mechanism proposed by Yamada *et al*. [48] to automatically select deep features learned during the first stage for predicting Hi-C interactions. Unlike traditional deterministic approaches, which rely on predefined criteria to select features, stochastic gating allows the model to explore a broader range of potential feature subsets and identify those crucial for the training task. Stochastic gating approximates the ℓ_0 sparsity constraint by continuous relaxation of Bernoulli distribution.

We jointly optimize the feature selection function s_μ and the target network g_ϕ, which predicts the Hi-C values. The total objective can be written as

$$\min_{\phi,\nu} \mathbb{E}_{(X,Y)\sim P_{\mathrm{HiC}}} \mathbb{E}_Z \left[L_T \left(g_\phi(s_\nu(h(X), Z)), Y \right) + \lambda L_R(Z) \right], \qquad (2)$$

where L_T and L_R denote the target loss function and feature sparsity regularization, respectively, parameter λ controls the regularization strength, Z is a random sparsity-defining vector of size N_L, X and Y are sampled from a joint data distribution P_{HiC}. The components of Z provide the gating mechanism: For k-th feature $z_k = \max(0, \min(1, v_k + \epsilon_k))$, where v_k is learned and ϵ_k is sampled from

$\mathcal{N}(0,\sigma^2)$ with a predefined σ during training. Following Yamada *et al*. [48], we set σ = 0.5. Herein we propose to use the stochastic gating $s_\nu(h(X),Z) = \{s_\nu(h(x_1),Z), s_\nu(h(x_2),Z), ..., s_\nu(h(x_{401}),Z)\}$ at the feature level instead of the node level (the latter was initially proposed by Yamada *et al.* [48]) and aggregate the selected features using average pooling to reduce the model complexity. The gating mechanism is applied for each ATAC-seq peak x_i: $s_\nu(h(x_i), Z) = \sum_{k=1}^{N_L} h_k(x_i) \cdot z_k$. The sparsity regularization is the sum of the probabilities that the gates are active, which is equal to $\mathbb{E}_Z L_R(Z) = \sum_{k=1}^{N_L} P(z_k > 0) = \sum_{k=1}^{N_L} \Phi\left(\frac{\nu_k}{\sigma}\right)$, where Φ is the Gaussian cumulative distribution function. During the inference stage, we set $\hat{z}_k = \max(0,\min(1,\nu_k))$.

- *Hi-C prediction network* g_ϕ comprises a transformer-based encoder g_ϕ^{ENC} that extracts information about the local and global context of the input sequence followed by a task-specific lightweight multilayer perceptron (MLP) decoder g_ϕ^{DEC}. The encoder consists of four multi-head attention blocks as proposed by Vaswani *et al* [49]. We used four attention heads, a dropout rate of 0.1, and set the hidden dimension d_{model} equal to 32. To more accurately encode the genomic distance between input tokens, which is crucial for predicting interaction level, here we propose a genomic-distance-aware positional encoding. We compute the relative genomic distance between each token and a reference token and encode the distance using the sine-cosine positional encoding scheme [49]. Each token corresponds to a region centered around a single ATAC-seq peak. We designate the position of the highest ATAC-seq peak value as the token position. The middle token is chosen as the reference. We encode a maximum distance of 3 Mbp with a resolution of 500 bp. The encoder extracts a representation for each token, capturing its relation with the other tokens in the input sequence. The output of g_ϕ^{ENC} is defined as $E_{\phi,\nu} = \{e_1,e_2,...,e_{401}\}$ with the i-th token representation $e_i \in \mathbb{R}^{d_{\text{model}}}$, *i.e.* $E_{\phi,\nu} = g_\phi^{\text{ENC}}(s_\nu(h(X),Z))$. To compute the Hi-C interaction value between two peaks locations i and j, the encoder output of the i-th and j-th token are concatenated and fed into the decoder predicting Hi-C value, *i.e.*, $\hat{y}_{\phi,\nu}^{(i,j)} = g_\phi^{\text{DEC}}(e_i, e_j)$ and e_j depend on ϕ,ν. We use two MLP layers with ReLU activation for the decoder. During training, sequences are randomly flipped to ensure that the decoder remains

invariant to the order of the tokens. The target loss L_T between the true interactions Y and the estimated interactions $\hat{Y}_{\phi,\nu}$ in (Equation 2) is defined as

$$L_T(\hat{Y}_{\phi,\nu}, Y) = \frac{1}{N}\sum_{i,j}(\hat{y}_{\phi,\nu}^{(i,j)} - y^{(i,j)})^2 \quad (3)$$

where *i,j* are encoded following the DeepC approach [3] in a vertical, zigzag pole over the center of the sequence window (Figure 1c). *N* is the total number of interaction pairs. During training the empirical version of the loss (Equation 2) is optimized in both ν and ϕ.

- *Auxiliary information* can be optionally used for Hi-C prediction. Since the Hi-C values between two bins are also dependent on the mappability of the independent bins [50], we also include the 36bp-mappability track, obtained from the UCSC Genome Browser [51, 52], as input to the Stage 2 of UniversalEPI. First, we use a linear head to embed auxiliary information. We set the embedding dimension the same as the hidden dimension of the transformer. The embedded vector is then concatenated with the input embeddings, and resulting in a new input to the attention layers. This method allows the use of any auxiliary information, such as mappability and ATAC-seq.

Uncertainty estimation

We incorporate both *aleatoric* and *epistemic* uncertainty estimation into UniversalEPI, providing information on the reliability of model predictions. To capture the *aleatoric* uncertainty, we assume that each observed log Hi-C value is a sample from a Gaussian distribution. Instead of the point estimate $\hat{Y}_{\phi,\nu}$ comprised of $\hat{y}_{\phi,\nu}^{(i,j)}$ provided by the entire Hi-C prediction component $g_\phi \circ s_\nu$, we model the parameter of the distribution: the mean $M_{\phi,\nu} = \{\mu_{\phi,\nu}^{(i,j)} \mid \mu_{\phi,\nu}^{(i,j)} \in \mathbb{R}\}$ and the variance $\Sigma_{\phi,\nu} = \{[\sigma_{\phi,\nu}^{(i,j)}]^2 \mid [\sigma_{\phi,\nu}^{(i,j)}]^2 \in \mathbb{R}_{\geq 0}\}$ between two peaks locations *i* and *j*, which are structured the same way as Y . We apply an exponential activation function to the variance output to guarantee non-negativity. Instead of the MSE loss L_T defined in Equation 3, herein we minimize in ϕ and μ the negative log-

likelihood loss weighted by the β-exponentiated variance estimates [16], referred to as β−NLL loss, which is defined as

$$L_{\beta\text{-NLL}}(M_{\phi,\nu}, \Sigma_{\phi,\nu}, Y) = \frac{1}{2N} \sum_{i,j} \text{sg}([\sigma_{\phi,\nu}^{(i,j)}]^2)^{\beta} \left(\log \left([\sigma_{\phi,\nu}^{(i,j)}]^2 \right) + \frac{(y_{\phi,\nu}^{(i,j)} - \mu_{\phi,\nu}^{(i,j)})^2}{[\sigma_{\phi,\nu}^{(i,j)}]^2} \right) \quad (4)$$

with the stop-gradient operator sg, and the parameter β controlling the trade-off between the regression accuracy and log-likelihood estimation. We set β = 0.5 following [16], which has been empirically found to provide the best trade-off. To estimate *epistemic* uncertainty, we utilize the deep ensemble method [17] by training K models with different random initializations, resulting in an ensemble $\{(\phi_k, \nu_k), k = 1,..,K\}$. The total predictive uncertainty $[\sigma_T^{(i,j)}]^2$ is a sum of aleatoric $[\sigma_a^{(i,j)}]^2$ and epistemic $[\sigma_e^{(i,j)}]^2$uncertainties. It is quantified by combining these components as follows:

$$[\sigma_T^{(i,j)}]^2 = [\sigma_a^{(i,j)}]^2 + [\sigma_e^{(i,j)}]^2 = \frac{1}{K} \sum_k [\sigma_{\phi_k,\nu_k}^{(i,j)}]^2 + \frac{1}{K} \sum_k (\mu_{\phi_k,\nu_k}^{(i,j)} - \mu_T^{(i,j)})^2 \quad (5)$$

with the empirical estimate of the mean $\mu_T^{(i,j)} = \frac{1}{K} \sum_k \mu_{\phi_k,\nu_k}^{(i,j)}$. We use an ensemble size of K = 10 for the reported results.

Training details

Our method is implemented with the PyTorch framework [53]. The training of the proposed method involves a two-stage model, where each stage was trained separately with different training setups. The TF prediction network was trained with Adam [54] optimizer using a learning rate of 10^{-4}, a weight decay of 10^{-4}, with a batch size of 1024 over 100 epochs. The best model was identified using Pearson's correlation score on the validation dataset. For each active binding site, we select a 1Kb segment of DNA sequence around the peak center and the corresponding ATAC-seq as input. Transcription factors CTCF, YY1, and SP1 were empirically chosen as target TFs based on their ability to assist in chromatin organization prediction (Figure S7). ZNF143 was initially considered another target TF candidate [55–

57]. However, a recent study [58] found the commonly used antibody of ZNF143 cross-reacting with CTCF, challenging reported associations between ZNF143 binding and chromatin looping. This, along with the fact that including ZNF143 did not improve Hi-C prediction accuracy (Figure S7), led us to exclude this TF from further consideration in our work. The first stage model was trained on the GM12878 and K562 cell lines. We randomly split the chromosomes into training, validation and test set. Chromosomes 5, 12, 13, and 21 were used for validation, 2, 6, and 19 for testing, and the remaining chromosomes for training. After training, the backbone of the TF prediction network was used as a feature extractor and its model parameters were frozen.

In the second stage, the Hi-C prediction network was trained using the AdamW [59] optimizer with a learning rate of 10^{-3}. We set the feature sparsity regularization λ to 0.01, which was tuned to select the minimum number of feature layers necessary. Since distance between elements is highly predictive of the Hi-C values, we employed a position embedding upscaling by a factor of $\sqrt{d_{model}}$. We trained two sets of models, one set uses GM12878 and K562 cell lines for training, and IMR90 and HepG2 cell lines for testing whereas the other uses HepG2 and IMR90 cells for training and GM12878 and K562 cells were heldout for testing. The second stage model was trained using the same chromosome split as the first stage. We found that the Hi-C data at 5Kb resolution also contains arbitrary bins with very few reads mapped to them (Figure S8). To ensure that these noisy samples do not affect our analysis, we removed these and their one-hop neighboring bins from all datasets. To increase robustness, we augmented the accessible regions in the training dataset with 10% of inaccessible regions of length 1Kb each that were chosen randomly from the genome ensuring that the model can also predict accurately for various inaccessible regions, which can arise due to mutations. The combined regions were then used to form the 401 consecutive peaks which is the input of UniversalEPI. We trained the network for 20 epochs and choose the model for evaluation with the highest Spearman's correlation on the validation data. Training the Hi-C prediction network took 12 hours on 0.2 of an A100 GPU with a batch size of 32, and 24 hours on a single RTX2080Ti with a batch size of 16.

Evaluation details

UniversalEPI was evaluated on unseen chromosomes of unseen cell types. To ensure that reliable targets were used for evaluation, we first flagged the bins that overlap with unmappable [1] and blacklisted[2] regions as reported by Kelley *et al* [60]. We then removed the interactions containing a flagged endpoint. Finally, we calculated Spearman's correlation and Pearson's correlation using the smoothed log Hi-C of the remaining interactions.

To understand the biological information captured by our model, we removed the strong influence of distance on Hi-C prediction by calculating distance-stratified correlation. Specifically, we calculated the correlation using all interactions that lie at a particular distance *d* away from each other, where *d* lies between 0 and 2Mb with a step size of 5Kb.

State-of-the-art methods

In this section, we describe the configurations of the state-of-the-art methods C.Origami [11] and Akita [4], against which we compare our approach. We did not compare against DeepC [3], as it has approximately 10 times more parameters than Akita and the two models are shown to perform comparably [4].

Various existing methods follow different Hi-C preprocessing techniques, predict Hi-C upto different distances and at different resolutions. For example, Akita predicts Hi-C at 2Kb resolution upto 1Mb after removal of effect of distance whereas C.Origami retains the effect of distances and predicts Hi-C at 10Kb resolution capturing interactions upto 2Mb away. Moreover, Akita employs 2D Gaussian smoothing which is not used by C.Origami.

To mitigate these differences, we re-trained the existing methods on our preprocessing of 5Kb Hi-C data and tuned their hyperparameters based on their performance on validation chromosomes of

[1] https://storage.googleapis.com/basenji barnyard2/umap k36 t10 l32 hg38.bed
[2] https://storage.googleapis.com/basenji barnyard2/hg38.blacklist.rep.bed

seen cell lines. We used the open-sourced codes as a starting point. For both methods, we did not modify the loss functions, optimizers, or learning rate schedulers. In case of C.Origami, we reduced the input size from 2Mb to 1Mb to compensate from increase in Hi-C resolution from 10Kb to 5Kb. This majorly led to changes in the hidden layer sizes and number of convolution layers in the encoder. To obtain the C.Origami model without the CTCF ChIP-seq as input, we simply excluded this track from the inputs without changing the configuration of the model. Finally, both versions of C.Origami models were trained for a maximum of 80 epochs. The top performing model was selected based on the validation loss and used for inference.

On the other hand, for Akita, we kept the input size of 1Mb intact. Since the Hi-C resolution is reduced from 2Kb to 5Kb and consequently, the target length is reduced from 512 to 201, we reduced the complexity of the convolution trunk. Specifically, we replaced the 11 convolution layers with 7 convolution layers. Moreover, the pooling size was increased from 2 to 5 for the final 4 convolution layers. This configuration led to the best performing Akita model. The Akita model was trained with two heads, corresponding to the two training cell lines. The model was trained for a maximum of 10000 epochs. The training was stopped early if the validation loss did not improve for 12 epochs. This resulted in training stopping after 35 epochs on average for the two models (one trained on GM12878 and K562 cells and the other on HepG2 and IMR90 cells). Adjusting these hyperparameters enabled each state-ofthe-art method to achieve performance levels comparable to those reported in their respective publications.

Baseline methods

We define 3 baseline methods, (1) Distance, (2) Median, and (3) Swap. These are computationally very simple as compared to the complex deep learning methods.

As the Hi-C interactions between two genomic segments heavily depend on the distance between the segments, we introduce the Distance baseline which only depends on the genomic distance between the open regions. In particular, we define the prediction between regions *i* and *j* as follows:

$$\hat{y}_{i,j} = -\log(d_{i,j})$$

where $d_{i,j}$ is the genomic distance (in 5Kb bins) between regions *i* and *j*.

Similarly, we define the Median baseline where prediction is the median of all the interactions between the open regions in the training data that are at a particular distance away. Mathematically, the prediction for interaction between regions *i* and *j*, that are $d_{i,j}$ bins away, is given by:

$$\hat{y}_{i,j} = \text{Median}([y_1, y_2, \ldots, y_n])$$

where $y_1 \ldots y_n$ are Hi-C interactions from the training data between all pairwise open regions that are $d_{i,j}$ bins apart.

Finally, we define the third baseline called Swap to capture the similarity in Hi-C across cell types. Here, we take the mean of Hi-C interactions from all the cell lines in the training data and use this as the predicted Hi-C matrix. If GM12878 and K562 are taken as training cell lines, then the predicted Hi-C interaction for the new unseen cell line would be given by:

$$\hat{y}_{i,j} = \frac{1}{2}(y_{i,j}^{GM12878} + y_{i,j}^{K562})$$

TAD annotations

TADs were identified using boundary calls from the 4DN data portal [38, 39]. We defined a consecutive pair of boundary calls (*i,j* where *i* < *j*) as a TAD if there was an overlapping ATAC-seq peak and a CTCF ChIP-seq peak within 10Kb of boundary calls *i* and *j*. Additionally, we required that the CTCF orientation is forward for *i* and reverse for *j*. The CTCF orientation was determined following the same approach

used to prepare target data for the first stage of UniversalEPI. Applying these criteria yielded approximately 200 TADs per cell line.

In-silico inversion of TAD boundary

For each identified TAD, we modified the ATAC-seq peak at boundary *j* by reversing the ATAC-seq signal, mappability profile, and by taking the reverse complement of the DNA sequence, resulting in a modified peak. We then extracted features from this modified peak using the first stage of UniversalEPI. These features, along with those from the 400 neighboring peaks, were used as input to the second stage of UniversalEPI.

Since interactions between two 5Kb bins can be influenced by multiple accessible chromatin regions within the two bins, it is challenging to isolate the effect of inverting a TAD endpoint in such cases. To address this, we focused our analysis on TADs with a unique ATAC-seq peak within 5Kb of each endpoint.

Macrophage activation data processing

We validated UniversalEPI on the data collected for activated macrophages derived from the THP-1 monocytic cell line. Since we require the ATAC-seq peak calls and the signal p-value bigwig track, we used the raw ATAC-seq paired-end fastq files provided by Reed *et al.* [18] and performed preprocessing similar to the ones done by the authors. Adaptors and low-quality reads were first trimmed using Trim Galore! v0.6.10. Reads were then aligned using BWA mem. Samtools v1.13 [61] was then used to sort the aligned reads. Duplicated reads were removed using PicardTools whereas mitochondrial reads were filtered using Samtools. The replicates for each timestamp were merged using Samtools. Finally, ENCODE's ATAC-seq pipeline [62] was used to obtain the peaks and signal p-value bigwig track. Deduplication was done as described before for the generated peak files. This resulted in approximately 200K peaks for each time point. For the generated bigwig tracks, edgeR's TMM normalization [44] was done with the GM12878 cell line as the reference (Figure S5).

The .hic files, for each of the 8 different time stamps, were directly obtained from Reed *et al* [18]. Nearly 370M interactions were observed on average for each of the time points. ICE normalization [40] was applied to each .hic file followed by the z-score normalization using the GM12878 cell line as the reference (as described above).

Differential loop prediction

For the validation of the ability of UniversalEPI to identify differential loops, we used the selected loops provided by Reed *et al*. The *gain.early* and *gain.late* were merged to form the set of gained loops. A similar approach was used for lost loops. To ensure a reliable set of differential loops, the sets of gained, static, and lost loops were further filtered based on fold change (*FC*) of ICE-normalized Hi-C between the 24 hours and 0 hour time points. Specifically, lost loops with $FC < 0.67$, gained loops with $FC > 1.5$, and static loops with $FC \in [0.67, 1.5]$ were retained.

The ATAC-seq peak sets at 0th and 24th hour time points were merged to get a set of regions for which Hi-C signal was predicted for the two time points using UniversalEPI. If multiple peaks existed within 500bp of each other, the peak with the maximum overall enrichment was retained resulting in a total of 235,690 peaks. UniversalEPI predictions were then independently calculated on this merged peak set using ATAC-seq bigwig signals for the two time points. Finally the maximum-confidence log2-fold change was calculated between the predictions at these two time points. Assuming UniversalEPI predicted an interaction for two time points t_1 and t_2 with means μ_1 and μ_2 and variances σ_1^2 and σ_2^2 respectively, the maximum-confidence log2-fold change between time points t_2 and t_1 was set to 0 if the intervals $[\mu_1-\sigma_1, \mu_1+\sigma_1]$ and $[\mu_2-\sigma_2, \mu_2+\sigma_2]$ overlapped. In case of no overlap, the maximum-confidence log2-fold change was given by $log_2(\frac{e^{\mu_2}}{e^{\mu_1}})$.

Comparison with ChromaFold

To compare UniversalEPI's performance against ChromaFold [12], we trained and evaluated both the models using the same input single-cell ATAC-seq (scATAC-seq) data. Since UniversalEPI is designed to

work with bulk ATAC-seq signal and ATAC-seq peaks, we first bulkified the scATAC-seq data. The bam files were directly downloaded from ENCODE (Supplementary Table 1) and converted to signal and peak files using the techniques explained in sections above. These bulkified ATAC-seq signals and peaks were used to train the UniversalEPI model and the evaluation of both the models.

Both the models were trained using data from GM12878 and HepG2 cell lines and evaluated on K562 and IMR90 cell lines. We followed the same processing as introduced in Gao *et al.* [12] *i.e.* we used ICE-normalized Hi-C at 10Kb resolution. Z-scores were then calculated using HiC-DC+ [28] and the resulting values were clipped between -16 and 16. Chromosomes 5, 18, 20, and 21 were used for testing whereas chromosomes 3 and 15 for validation, and the remaining chromosomes for training, as done in Gao *et al* [12]. This allowed us to directly use the preprocessing, model, and training scripts from the source code of ChromaFold.

The models were compared using interactions between accessible chromatin regions. The interactions were also filtered based on unmappable and blacklisted regions. UniversalEPI and ChromaFold were also evaluated based on distancestratified correlation, evaluating the performance of these methods in predicting interactions between genomic regions located upto 1.5Mb away from each other.

Esophageal adenocarcinoma data processing

The multiome (single-nuclei ATAC-seq and single-nuclei RNA-seq) profiles of 10 esophageal adenocarcinoma (EAC) patients were obtained from Yates *et al* [29]. Using the cell scores from Yates *et al.*, we selected the top 20% unique cells in each of the differentiated (cNMF4) and undifferentiated (cNMF5) programs as the representative cells.

The representative cell barcodes are extracted from the individual patients' bam file and the resulting patient-specific bam files are merged to obtain the pseudo-bulk ATAC-seq bam. For each program, ATAC-seq bam files were then used to obtain the signal p-values bigwig and narrowpeak files using ENCODE's ATAC-seq pipeline [62] as done before. The ATAC-seq peaks were merged for the two

programs to obtain a common set of peaks for downstream differential analysis. This was followed by deduplication (as done for datasets before) which resulted in a total of 280,368 peaks. The signal p-value bigwigs were normalized using edgeR's TMM normalization [44] was done with the GM12878 cell line as the reference.

The pseudo-bulk gene expression is obtained by summing the gene counts from all representative cells. Finally, the gene expression counts for each program were converted to counts-per-million (CPM) to mitigate the library-size bias.

Promoter activity

The promoter activity (A_P) was obtained using accessibility of the promoter (ATAC_P), accessibility of all the interacting enhancers (ATAC_E), and the interaction strength between the enhancer and promoter ($\text{Hi-C}_{E,P}$). The accessibility of each CRE was extracted as the maximum signal in the vicinity (±50bp) of the gene transcriptional start site or the point of maxima called by peak calling of MACS2 for the enhancers. Specifically, we defined the promoter activity as a linear combination of the promoter accessibility and the Hi-C weighted by the total accessibility of enhancer and promoter. Inspired by Loubiere *et al.* [63], the total accessibility of enhancer and promoter was defined using the multiplicative model. This is mathematically given by

$$A_P = \beta_0 + \beta_1 \log_2(\text{ATAC}_P + 1) + \beta_2 \sum_E (\text{ATAC}_P \cdot \text{ATAC}_E \cdot \text{Hi-C}_{E,P})$$

where β_0 is the intercept and β_1, and β_2 are the coefficients, all of which are learnt by a lasso regressor. To mitigate the effect of weak interactions, only those enhancers were considered that had a predicted log Hi-C interaction of greater than 1 with the promoter.

Funding

This project is partially funded by the Swiss Data Science Center (SDSC) collaborative projects grant (C22-09); Swiss Government Excellence Scholarship (ESKAS-Nr: 2021.0468) to A.G..

Competing interests

F.J.T. consults for Immunai Inc., Singularity Bio B.V., CytoReason Ltd and Omniscope Ltd, and has ownership interest in Dermagnostix GmbH and Cellarity. I.L.I. currently works at Bioptimus.

References

[1] Uyehara, C.M., Apostolou, E.: 3D enhancer-promoter interactions and multi-connected hubs: Organizational principles and functional roles. Cell Reports **42**(4) (2023)

[2] Pollex, T., Rabinowitz, A., Gambetta, M.C., Marco-Ferreres, R., Viales, R.R., Jankowski, A., Schaub, C., Furlong, E.E.: Enhancer–promoter interactions become more instructive in the transition from cell-fate specification to tissue differentiation. Nature Genetics **56**(4), 686–696 (2024)

[3] Schwessinger, R., Gosden, M., Downes, D., Brown, R.C., Oudelaar, A.M., Telenius, J., Teh, Y.W., Lunter, G., Hughes, J.R.: DeepC: predicting 3D genome folding using megabase-scale transfer learning. Nature Methods **17**(11), 1118–1124 (2020)

[4] Fudenberg, G., Kelley, D.R., Pollard, K.S.: Predicting 3D genome folding from DNA sequence with Akita. Nature Methods **17**(11), 1111–1117 (2020)

[5] Zhou, J.: Sequence-based modeling of three-dimensional genome architecture from kilobase to chromosome scale. Nature Genetics **54**(5), 725–734 (2022)

[6] Li, W., Wong, W.H., Jiang, R.: DeepTACT: predicting 3D chromatin contacts via bootstrapping deep learning. Nucleic Acids Research **47**(10), 60–60 (2019)

[7] Zhang, S., Chasman, D., Knaack, S., Roy, S.: In silico prediction of high-resolution Hi-C interaction matrices. Nature Communications **10**(1), 5449 (2019)

[8] Yang, R., Das, A., Gao, V.R., Karbalayghareh, A., Noble, W.S., Bilmes, J.A., Leslie, C.S.: Epiphany: predicting Hi-C contact maps from 1D epigenomic signals. Genome Biology **24**(1), 134 (2023)

[9] Whalen, S., Truty, R.M., Pollard, K.S.: Enhancer–promoter interactions are encoded by complex genomic signatures on looping chromatin. Nature Genetics **48**(5), 488–496 (2016)

[10] Chen, K., Zhao, H., Yang, Y.: Capturing large genomic contexts for accurately predicting enhancer-promoter interactions. Briefings in Bioinformatics **23**(2), 577 (2022)

[11] Tan, J., Shenker-Tauris, N., Rodriguez-Hernaez, J., Wang, E., Sakellaropoulos, T., Boccalatte, F., Thandapani, P., Skok, J., Aifantis, I., Fenyö, D., *et al.*: Cell-type-specific prediction of 3D chromatin organization enables high-throughput in silico genetic screening. Nature Biotechnology **41**(8), 1140–1150 (2023)

[12] Gao, V.R., Yang, R., Das, A., Luo, R., Luo, H., McNally, D.R., Karagiannidis, I., Rivas, M.A., Wang, Z.-M., Barisic, D., *et al.*: ChromaFold predicts the 3D contactmap from single-cell chromatin accessibility. Nature Communications **15**(1), 9432 (2024)

[13] Ong, C.-T., Corces, V.G.: CTCF: an architectural protein bridging genome topology and function. Nature Reviews Genetics **15**(4), 234–246 (2014)

[14] Makhlouf, M., Ouimette, J.-F., Oldfield, A., Navarro, P., Neuillet, D., Rougeulle,C.: A prominent and conserved role for YY1 in Xist transcriptional activation. Nature Communications **5**(1), 4878 (2014)

[15] O'Connor, L., Gilmour, J., Bonifer, C.: Focus: Epigenetics: The role of the ubiquitously expressed transcription factor Sp1 in tissue-specific transcriptional regulation and in disease. The Yale journal of biology and medicine **89**(4), 513 (2016)

[16] Seitzer, M., Tavakoli, A., Antic, D., Martius, G.: On the pitfalls of heteroscedastic uncertainty estimation with probabilistic neural networks. arXiv preprint arXiv:2203.09168 (2022)

[17] Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. Advances in neural information processing systems **30** (2017)

[18] Reed, K.S., Davis, E.S., Bond, M.L., Cabrera, A., Thulson, E., Quiroga, I.Y., Cassel, S., Woolery, K.T., Hilton, I., Won, H., et al.: Temporal analysis suggestsa reciprocal relationship between 3D chromatin structure and transcription. Cell reports **41**(5) (2022)

[19] Shrikumar, A., Greenside, P., Kundaje, A.: Learning important features through propagating activation differences. In: International Conference on Machine Learning, pp. 3145–3153 (2017). PMLR

[20] Lieberman-Aiden, E., Van Berkum, N.L., Williams, L., Imakaev, M., Ragoczy, T., Telling, A., Amit, I., Lajoie, B.R., Sabo, P.J., Dorschner, M.O., *et al.*: Com- prehensive mapping of long-range interactions reveals folding principles of the human genome. Science **326**(5950), 289–293 (2009)

[21] Wit, E., Vos, E.S., Holwerda, S.J., Valdes-Quezada, C., Verstegen, M.J., Teu- nissen, H., Splinter, E., Wijchers, P.J., Krijger, P.H., Laat, W.: CTCF binding polarity determines chromatin looping. Molecular Cell **60**(4), 676–684 (2015)

[22] Valton, A.-L., Dekker, J.: TAD disruption as oncogenic driver. Current opinion in genetics & development **36**, 34–40 (2016)

[23] Hnisz, D., Weintraub, A.S., Day, D.S., Valton, A.-L., Bak, R.O., Li, C.H., Goldmann, J., Lajoie, B.R., Fan, Z.P., Sigova, A.A., *et al.*: Activation of proto-oncogenes by disruption of chromosome neighborhoods. Science **351**(6280), 1454–1458 (2016)

[24] Kloetgen, A., Thandapani, P., Ntziachristos, P., Ghebrechristos, Y., Nomikou,

S.,Lazaris, C., Chen, X., Hu, H., Bakogianni, S., Wang, J., *et al.*: Three-dimensional chromatin landscapes in T cell acute lymphoblastic leukemia. Nature Genetics **52**(4), 388–400 (2020)

[25] Moore, J., Purcaro, M., Pratt, H., Epstein, C., Shoresh, N., Adrian, J., Kawli, T., Davis, C., Dobin, A., Kaul, R., et al.: & Weng, Z.(2020). Expanded encyclopaediasof DNA elements in the human and mouse genomes. Nature **583**(7818), 699–710

[26] Tang, F., Yang, Z., Tan, Y., Li, Y.: Super-enhancer function and its application in cancer targeted therapy. NPJ precision oncology **4**(1), 2 (2020)

[27] Khan, A., Zhang, X.: dbSUPER: a database of super-enhancers in mouse and human genome. Nucleic Acids Research **44**(D1), 164–171 (2016)

[28] Sahin, M., Wong, W., Zhan, Y., Van Deynze, K., Koche, R., Leslie, C.S.: HiC-DC+ enables systematic 3D interaction calls and differential analysis for Hi-C and HiChIP. Nature Communications **12**(1), 3366 (2021)

[29] Yates, J., Mathey-Andrews, C., Park, J., Garza, A., Gagné, A., Hoffman, S., Bi, K., Titchen, B., Hennessey, C., Remland, J., et al.: Cell states and neighborhoodsin distinct clinical stages of primary and metastatic esophageal adenocarcinoma. bioRxiv (2024)

[30] Celià-Terrassa, T., Liu, D.D., Choudhury, A., Hang, X., Wei Y., Zamalloa, J., Alfaro-Aco, R., Chakrabarti, R., Jiang, Y.-Z., Koh, B.I., *et al.*: Normal and cancerous mammary stem cells evade interferon-induced constraint through the miR-199a–LCOR axis. Nature Cell Biology **19**(6), 711–723 (2017)

[31] Pérez-Núñez, I., Rozalen, C., Palomeque, J.Á., Sangrador, I., Dalmau, M., Com- erma, L., Hernandez-Prat, A., Casadevall, D., Menendez, S., Liu, D.D., *et al.*: LCOR mediates interferon-independent tumor immunogenicity and responsive- ness to immune-checkpoint blockade in triple-negative breast cancer. Nature Cancer **3**(3), 355–370 (2022)

[32] Bleu, M., Mermet-Meillon, F., Apfel, V., Barys, L., Holzer, L., Bachmann Salvy, M., Lopes, R., Amorim Monteiro Barbosa, I., Delmas, C., Hinniger, A., *et al.*: PAX8 and MECOM are interaction partners driving ovarian cancer. Nature Communications **12**(1), 2442 (2021)

[33] Ma, Y., Kang, B., Li, S., Xie, G., Bi, J., Li, F., An, G., Liu, B., Li, J., Shen, Y.,*et al.*: CRISPR-mediated MECOM depletion retards tumor growth by reducing cancer stem cell properties in lung squamous cell carcinoma. Molecular Therapy **30**(11), 3341–3357 (2022)

[34] Lou, M., Zou, L., Zhang, L., Lu, Y., Chen, J., Zong, B.: MECOM and the PRDM gene family in uterine endometrial cancer: bioinformatics and experi- mental insights into pathogenesis and therapeutic potentials. Molecular Medicine**30**(1), 190 (2024)

[35] Madry, A.: Towards deep learning models resistant to adversarial attacks. arXiv preprint arXiv:1706.06083 (2017)

[36] Shafahi, A., Najibi, M., Ghiasi, M.A., Xu, Z., Dickerson, J., Studer, C., Davis, L.S., Taylor, G., Goldstein, T.: Adversarial training for free! Advances in Neural Information Processing Systems **32** (2019)

[37] Bai, T., Luo, J., Zhao, J., Wen, B., Wang, Q.: Recent advances in adversarial training for adversarial robustness. arXiv preprint arXiv:2102.01356 (2021)

[38] Dekker, J., Belmont, A.S., Guttman, M., Leshyk, V.O., Lis, J.T., Lomvardas, S., Mirny, L.A., O'shea, C.C., Park, P.J., Ren, B., *et al.*: The 4D nucleome project. Nature **549**(7671), 219–226 (2017)

[39] Reiff, S.B., Schroeder, A.J., Kırlı, K., Cosolo, A., Bakker, C., Mercado, L., Lee, S., Veit, A.D., Balashov, A.K., Vitzthum, C., *et al.*: The 4D Nucleome DataPortal as a resource for searching and visualizing curated nucleomics data. NatureCommunications **13**(1), 2365 (2022)

[40] Imakaev, M., Fudenberg, G., McCord, R.P., Naumova, N., Goloborodko, A., Lajoie, B.R., Dekker, J., Mirny, L.A.: Iterative correction of Hi-C data reveals hallmarks of chromosome organization. Nature Methods **9**(10), 999–1003 (2012)

[41] Wolff, J., Rabbani, L., Gilsbach, R., Richard, G., Manke, T., Backofen, R., Grüning, B.A.: Galaxy HiCExplorer 3: a web server for reproducible Hi-C, cap- ture Hi-C and single-cell Hi-C data analysis, quality control and visualization. Nucleic Acids Research **48**(W1), 177–184 (2020)

[42] Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cour- napeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., *et al.*: SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods **17**(3), 261–272 (2020)

[43] Ramírez, F., Ryan, D.P., Grüning, B., Bhardwaj, V., Kilpert, F., Richter, A.S., Heyne, S., Dündar, F., Manke, T.: deepTools2: a next generation web server for deep-sequencing data analysis. Nucleic Acids Research **44**(Web Server issue), 160 (2016)

[44] Robinson, M.D., Oshlack, A.: A scaling normalization method for differential expression analysis of RNA-seq data. Genome Biology **11**, 1–9 (2010)

[45] Rauluseviciute, I., Riudavets-Puig, R., Blanc-Mathieu, R., Castro-Mondragon, J.A., Ferenc, K., Kumar, V., Lemma, R.B., Lucas, J., Chèneby, J., Barana- sic, D., *et al.*: JASPAR 2024: 20th anniversary of the open-access database of transcription factor binding profiles. Nucleic Acids Research **52**(D1), 174–182 (2024)

[46] Swindell, W.R., Johnston, A., Sun, L., Xing, X., Fisher, G.J., Bulyk, M.L., Elder, J.T., Gudjonsson, J.E.: Meta-profiles of gene expression during aging: limited sim- ilarities between mouse and human and an unexpectedly decreased inflammatory signature. PloS One 7(3), 33204 (2012)

[47] Xu, B., Wang, N., Chen, T., Li, M.: Empirical evaluation of rectified activations in convolutional network. arXiv preprint arXiv:1505.00853 (2015)

[48] Yamada, Y., Lindenbaum, O., Negahban, S., Kluger, Y.: Feature selection using stochastic gates. In: International Conference on Machine Learning, pp. 10648– 10659 (2020). PMLR

[49] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in Neural Information Processing Systems **30** (2017)

[50] Carty, M., Zamparo, L., Sahin, M., González, A., Pelossof, R., Elemento, O., Leslie, C.S.: An integrated model for detecting significant chromatin interactions from high-resolution Hi-C data. Nature Communications **8**(1), 15454 (2017)

[51] Nassar, L.R., Barber, G.P., Benet-Pagès, A., Casper, J., Clawson, H., Diekhans, M., Fischer, C., Gonzalez, J.N., Hinrichs, A.S., Lee, B.T., *et al.*: The UCSC genome browser database: 2023 update. Nucleic Acids Research **51**(D1), 1188– 1195 (2023)

[52] Karimzadeh, M., Ernst, C., Kundaje, A., Hoffman, M.M.: Umap and Bismap: quantifying genome and methylome mappability. Nucleic Acids Research **46**(20), 120–120 (2018)

[53] Ansel, J., Yang, E., He, H., Gimelshein, N., Jain, A., Voznesensky, M., Bao, B., Bell, P., Berard, D., Burovski, E., Chauhan, G., Chourdia, A., Constable, W., Des-maison, A., DeVito, Z., Ellison, E., Feng, W., Gong, J., Gschwind, M., Hirsh, B., Huang, S., Kalambarkar, K., Kirsch, L., Lazos, M., Lezcano, M., Liang, Y., Liang,J., Lu, Y., Luk, C., Maher, B., Pan, Y., Puhrsch, C., Reso, M., Saroufim, M., Siraichi, M.Y., Suk, H., Suo, M., Tillet, P., Wang, E., Wang, X., Wen, W., Zhang,S., Zhao, X., Zhou, K., Zou, R., Mathews, A., Chanan, G., Wu, P., Chintala, S.: PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Trans- formation and Graph Compilation. In: 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS '24) (2024). https://doi.org/10.1145/3620665.3640366 ACM. https://pytorch.org/assets/pytorch2-2.pdf

[54] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprintarXiv:1412.6980 (2014)

[55] Rao, S.S., Huntley, M.H., Durand, N.C., Stamenova, E.K., Bochkov, I.D., Robin-son, J.T., Sanborn, A.L., Machol, I., Omer, A.D., Lander, E.S., *et al.*: A 3D map of the human genome at kilobase resolution reveals principles of chromatin looping. Cell **159**(7), 1665–1680 (2014)

[56] Bailey, S.D., Zhang, X., Desai, K., Aid, M., Corradin, O., Cowper-Sal· lari, R., Akhtar-Zaidi, B., Scacheri, P.C., Haibe-Kains, B., Lupien, M.: ZNF143 provides sequence specificity to secure chromatin interactions at gene promoters. Nature communications **6**(1), 6186 (2015)

[57] Zhou, Q., Yu, M., Tirado-Magallanes, R., Li, B., Kong, L., Guo, M., Tan, Z.H., Lee, S., Chai, L., Numata, A., *et al.*: ZNF143 mediates CTCF-bound promoter–enhancer loops required for murine hematopoietic stem and progenitorcell function. Nature Communications **12**(1), 43 (2021)

[58] Magnitov, M.D., Maresca, M., Alonso Saiz, N., Teunissen, H., Braccioli, L., Wit, E.: ZNF143 is a transcriptional regulator of nuclear-encoded mitochondrial genes that acts independently of looping and CTCF. bioRxiv, 2024–03 (2024)

[59] Loshchilov, I.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)

[60] Kelley, D.R., Reshef, Y.A., Bileschi, M., Belanger, D., McLean, C.Y., Snoek, J.: Sequential regulatory activity prediction across chromosomes with convolutional neural networks. Genome Research **28**(5), 739–750 (2018)

[61] Li, H.: The sequence alignment/map (SAM) format and SAMtools 1000 genome project data processing subgroup. Bioinformatics **25**, 1 (2009)

[62] Lee, J., Christoforo, G., Foo, C., Probert, C., Kundaje, A., Boley, N., Kim, D., Dacre, M., et al.: kundajelab/atac dnase pipelines: 0.3. 3. Zenodo (2016)

[63] Loubiere, V., Almeida, B.P., Pagani, M., Stark, A.: Developmental and house-keeping transcriptional programs display distinct modes of enhancer-enhancer cooperativity in Drosophila. Nature Communications **15**(1), 8584 (2024)