



## Early feature extraction drives model performance in high-resolution chromatin accessibility prediction

Aayush Grover, Till Muser, Liine Kasak, et al.

*Genome Res.* published online January 12, 2026

Access the most recent version at doi:[10.1101/gr.281042.125](https://doi.org/10.1101/gr.281042.125)

---

<b>P&lt;P</b>	Published online January 12, 2026 in advance of the print journal.
<b>Accepted Manuscript</b>	Peer-reviewed and accepted for publication but not copyedited or typeset; accepted manuscript is likely to differ from the final, published version.
<b>Open Access</b>	Freely available online through the <i>Genome Research</i> Open Access option.
<b>Creative Commons License</b>	This manuscript is Open Access. This article, published in <i>Genome Research</i> , is available under a Creative Commons License (Attribution-NonCommercial 4.0 International license), as described at <a href="http://creativecommons.org/licenses/by-nc/4.0/">http://creativecommons.org/licenses/by-nc/4.0/</a> .
<b>Email Alerting Service</b>	Receive free email alerts when new articles cite this article - sign up in the box at the top right corner of the article or <a href="#">click here</a> .

---



---

To subscribe to *Genome Research* go to:  
<https://genome.cshlp.org/subscriptions>

---

Published by Cold Spring Harbor Laboratory Press

# Early feature extraction drives model performance in high-resolution chromatin accessibility prediction

Aayush Grover<sup>1,2</sup>, Till Muser<sup>3</sup>, Liine Kasak<sup>1</sup>, Lin Zhang<sup>3</sup>,  
Ekaterina Krymova<sup>3\*</sup>, Valentina Boeva<sup>1,2,4\*</sup>

<sup>1</sup>Department of Computer Science, ETH Zurich, 8092 Zurich, Switzerland.

<sup>2</sup>Swiss Institute of Bioinformatics, 1015 Lausanne, Switzerland.

<sup>3</sup>Swiss Data Science Center, EPF Lausanne and ETH Zurich, 8092 Zurich, Switzerland.

<sup>4</sup>Université Paris Cité, Institut Cochin, INSERM U1016, 75006 Paris, France.

\*Corresponding author(s). E-mail(s): [ekaterina.krymova@sdsc.ethz.ch](mailto:ekaterina.krymova@sdsc.ethz.ch);  
[valentina.boeva@inf.ethz.ch](mailto:valentina.boeva@inf.ethz.ch);

## Abstract

Fine-grained prediction of chromatin accessibility from DNA sequence is a foundational step in modeling gene expression changes resulting from sequence variants. Yet, few methods operate at the resolution necessary to capture subtle effects of single-nucleotide changes. Furthermore, it remains unclear which architectural components, such as residual connections, normalization strategies, or attention mechanisms, drive performance in these high-resolution predictions. To address these knowledge gaps, we systematically evaluate classic architectural choices and introduce ConvNeXt V2 blocks, originally developed for computer vision, as high-resolution feature extractors in deep learning models for genomic data. Integrated into diverse architectures such as CNNs, LSTMs, dilated CNNs, and transformers, ConvNeXt V2 blocks consistently improve performance, leading to similar prediction accuracy across these different model types. This reveals that early feature extraction, rather than downstream architecture, is the primary determinant of prediction accuracy. A comprehensive evaluation of these models on ATAC-seq signal prediction at 4 bp resolution in a cell-type-specific manner identifies the ConvNeXt-based dilated CNN as the most robust performer, better preserving the signal's shape. Our codebase and benchmarks provide practical tools for high-resolution chromatin modeling.

**Keywords:** Deep Learning, Chromatin Accessibility, Variant Effects, Cancer

# 1 Introduction

2 Predicting chromatin accessibility from DNA sequence is one of the central challenges  
3 in regulatory genomics, with applications including the interpretation of non-coding  
4 genetic variants and the decoding of regulatory element architecture (Zhou and Troy-  
5 anskaya, 2015; Kelley et al., 2018; Avsec, Agarwal, et al., 2021; Pampari et al., 2024).  
6 Accurate, high-resolution maps of chromatin accessibility are essential for downstream  
7 tasks such as identifying regulatory elements, modeling 3D chromatin organization,  
8 and predicting gene expression (Thibodeau et al., 2021; Karbalayghareh et al., 2022;  
9 Tan et al., 2023; Z. Zhang et al., 2023; Grover et al., 2024). While experimental  
10 assays like FAIRE-seq (Giresi et al., 2007), DNase-seq (Song and Crawford, 2010), and  
11 ATAC-seq (Buenrostro et al., 2015) provide high-resolution measurements of chro-  
12 matin accessibility in specific cellular contexts, they are resource-intensive and often  
13 unavailable for specific genetic variants. In the absence of such experimental assays,  
14 neural network-based predictions offer a reliable and efficient alternative for esti-  
15 mating changes in chromatin accessibility induced by the genetic variants in a cell  
16 type-specific manner.

17 Numerous deep learning approaches have been proposed to predict chromatin  
18 accessibility from DNA sequence. Earlier methods (Zhou and Troyanskaya, 2015; Kel-  
19 ley et al., 2016; Quang and Xie, 2016; Min et al., 2017; Guo et al., 2020) treated this as  
20 a binary classification task, labeling regions as either open or closed. Some of the later  
21 methods, such as Deopen (Liu et al., 2018) demonstrated the performance of their  
22 model in predicting the maximum intensity of chromatin accessibility corresponding  
23 to the input DNA segment. Subsequently, Basenji (Kelley et al., 2018) and Enformer  
24 (Avsec, Agarwal, et al., 2021) showed that it was possible to predict the intensity of  
25 the chromatin accessibility signal corresponding to the input DNA sequence with a rel-  
26 atively lower resolution (128bp) by binning the signal tracks. Although these methods  
27 are highly accurate, this coarse formulation (*i.e.*, predicting chromatin accessibility

28 profiles at 128bp or lower resolution) is still insufficient to study local changes in chro-  
29 matin accessibility caused by point mutations. Moreover, it hinders their applicability  
30 for downstream tasks, including prediction of transcription factor binding (Cazares  
31 et al., 2023), cell-type-specific chromatin structure (Tan et al., 2023; R. Yang et al.,  
32 2023; Z. Zhang et al., 2023; Gao et al., 2024; Grover et al., 2024), and gene expression  
33 (Karbalayghareh et al., 2022; Z. Zhang et al., 2023). Therefore, there is a need for  
34 the development of deep learning methods that can accurately predict high-resolution  
35 chromatin accessibility signals.

36 Recently introduced ChromBPNet (Pampari et al., 2024) predicts chromatin acces-  
37 sibility at base-pair resolution using dilated convolutional networks, representing an  
38 advance in resolution capabilities. However, it is not extensively compared against  
39 other model architectures, like recurrent neural networks and transformers, leaving  
40 questions about optimal architectural choices unanswered. GOPHER (Toneyan et al.,  
41 2022) was introduced as an extensive benchmark that compared various deep learning  
42 methods, such as Basenji (Kelley et al., 2018) and BPNet (Avsec, Weilert, et al., 2021),  
43 in their ability to predict chromatin accessibility for different experimental setups.  
44 Both of these methods use dilated convolutional neural networks (dCNNs) to capture  
45 long-range information from the input DNA sequence. These methods were evaluated  
46 against a convolutional neural network (CNN) baseline, highlighting the predictive  
47 prowess of dCNN-based methods over CNN-based methods. However, recurrent neu-  
48 ral network-based methods and transformer-based methods are not explored in this  
49 benchmark, creating a gap in understanding which architectural paradigms perform  
50 best for high-resolution chromatin accessibility prediction.

51 In this work, we present a systematic evaluation of deep learning models for  
52 predicting ATAC-seq signals at 4 bp resolution from DNA sequence. We begin by  
53 extensively tuning three commonly used architectures – CNNs, recurrent neural net-  
54 works (specifically, long short-term memory, LSTM), and dCNNs, for high-resolution

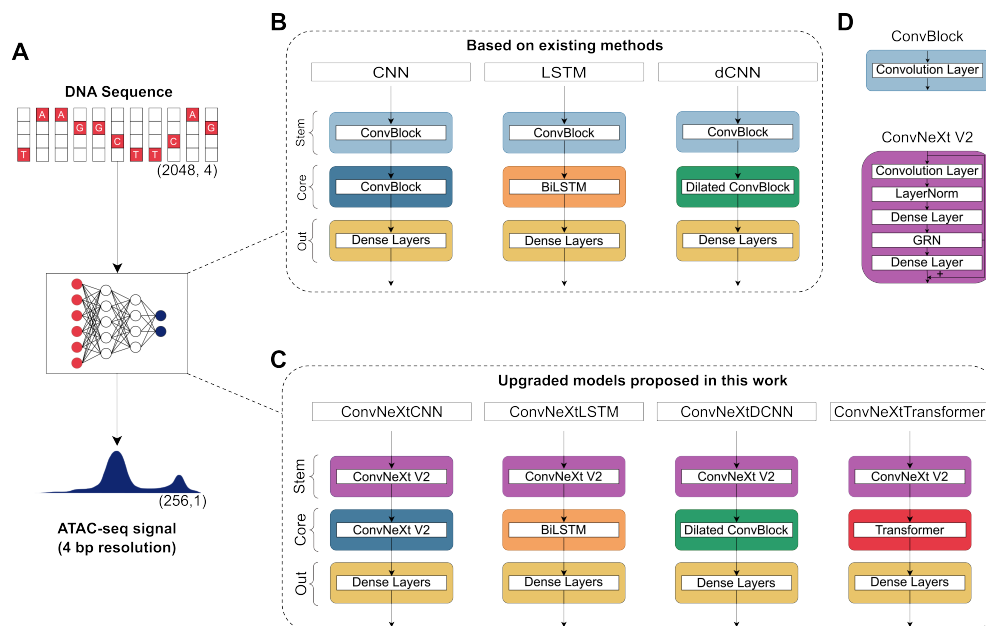
55 chromatin accessibility prediction. Building on this, we incorporate ConvNeXt V2  
56 (Woo et al., 2023) blocks, originally developed for vision tasks, as feature extractors  
57 across all architectures and introduce a novel transformer-based model tailored for  
58 this task. We evaluate each model’s performance on accessible regions and genome-  
59 wide, assess robustness to input shifts, examine interpretability, and test sensitivity to  
60 single-nucleotide variants (SNVs) using cancer patient data. This benchmark reveals  
61 how architectural choices impact predictive accuracy, robustness, and variant inter-  
62 pretation, providing a practical foundation for future development of deep learning  
63 methods in regulatory genomics.

## 64 **Results**

### 65 **Designing a deep-learning model benchmark for** 66 **high-resolution ATAC-seq prediction**

67 We evaluated four categories of deep learning architectures – convolutional neural  
68 networks (CNNs), recurrent neural networks (LSTMs), dilated CNNs (dCNNs), and  
69 transformers, based on their ability to predict ATAC-seq signal at high resolution.  
70 To ensure a fair comparison, we standardized the experimental setup across models:  
71 each received a 2048 bp one-hot encoded DNA sequence as input and was trained to  
72 predict the chromatin accessibility signal at 4 bp resolution across the central 1024  
73 bp region (Fig. 1). This resolution was chosen because transcription factor binding  
74 motifs typically span more than 4 bp (Stewart et al., 2012), making it sufficient to cap-  
75 ture relevant regulatory variation. The 2048 bp input window also provides adequate  
76 genomic context for the models to learn both local and distal sequence features.

77 We began by benchmarking architectures used in existing methods, broadly cate-  
78 gorized into CNNs (Zhou and Troyanskaya, 2015), LSTMs (Min et al., 2017; Liu et al.,  
79 2018), and dilated CNNs (Kelley et al., 2018; Pampari et al., 2024) (Fig. 1A). Since  
80 these models were originally developed for different prediction tasks and resolutions,



**Fig. 1 Overview of the benchmarking setup.** (A) Models are compared based on their ability to predict experimental ATAC-seq at 4 bp resolution from an input of 2048 bp long DNA sequences. (B) Models based on existing methods include CNN, LSTM, and dCNN-based architectures. Each model has a convolution-based stem to extract genomic features. (C) Models proposed in this work, including a transformer-based architecture. The new models use ConvNeXt V2 blocks to effectively extract features from the input DNA sequence. (D) ConvBlock uses a single convolutional layer whereas ConvNeXt V2 block additionally uses dense layers, layer normalization and Global Response Normalization (GRN).

81 we conducted an extensive hyperparameter search to adapt each to our standard-  
 82 ized 4 bp resolution setup. Next, to improve feature extraction, we replaced standard  
 83 convolutional layers in the stem of all three architectures with transformer-inspired  
 84 ConvNeXt V2 blocks, originally designed for vision tasks but well-suited for capturing  
 85 local and global genomic features (Woo et al., 2023) (Fig. 1B).

86 Previous studies have shown that transformer architectures can outperform  
 87 convolution-based models in capturing long-range dependencies in genomic sequences  
 88 (Avsec, Agarwal, et al., 2021; Ji et al., 2021; M. Yang and Ma, 2023; Dalla-Torre  
 89 et al., 2024). However, their application to high-resolution chromatin accessibility  
 90 prediction has not been explored. To address this limitation, we designed a novel

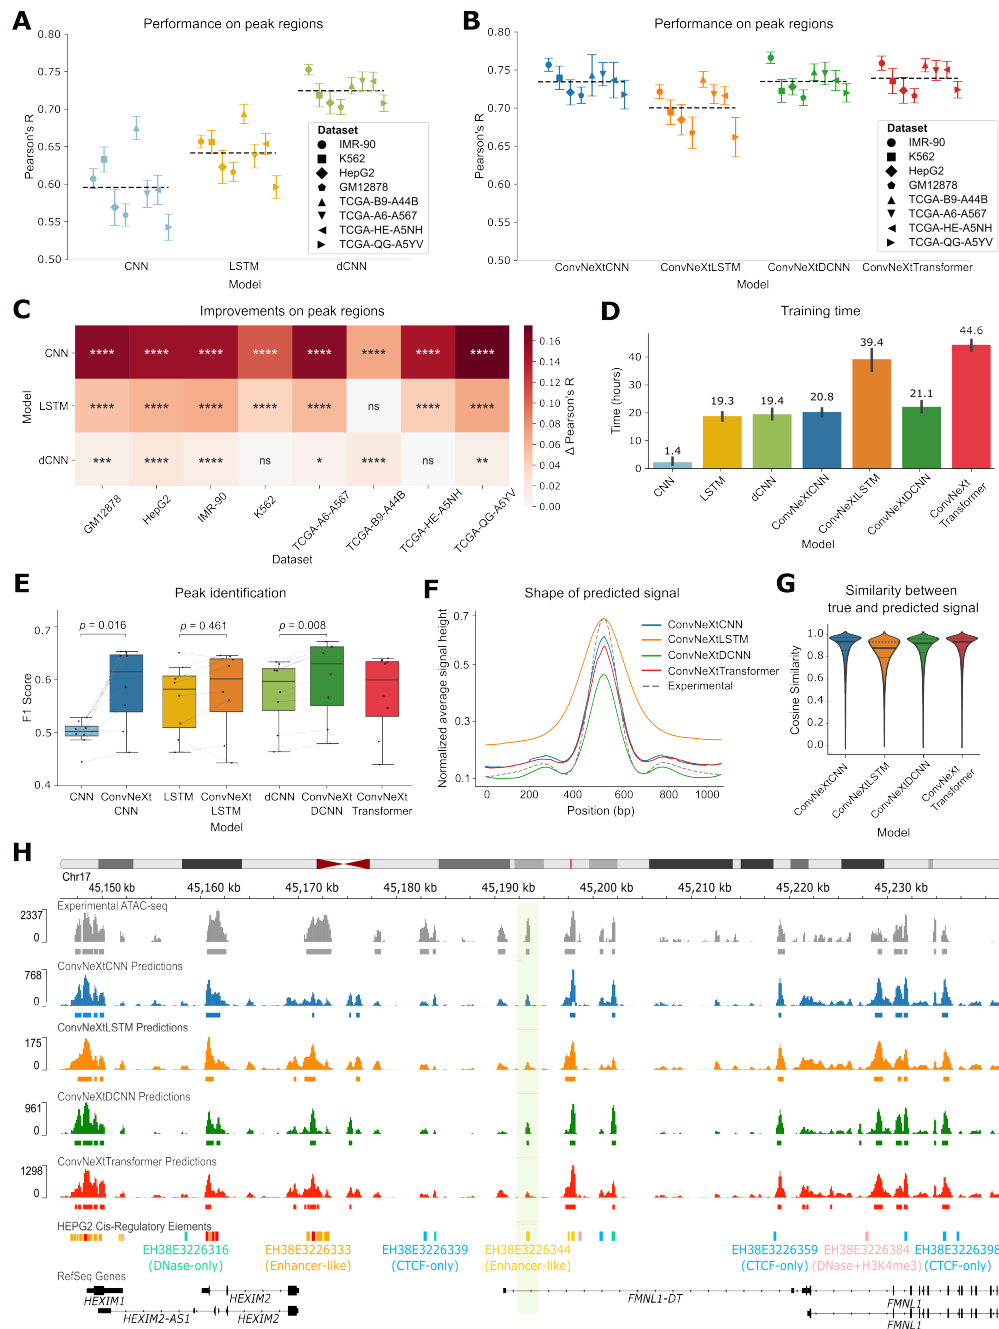
91 transformer-based model incorporating rotary positional embeddings (Su et al., 2024)  
92 (Fig. 1B).

93 The detailed architecture of each model can be found in Supplementary Fig. S1.  
94 Together, these models form a comprehensive benchmark across four deep learning  
95 architecture families.

## 96 **Improving model prediction accuracy with ConvNeXt-based** 97 **feature extraction**

98 Each model was extensively trained and tested on data from eight distinct datasets:  
99 four established human cell lines (GM12878, K562, IMR-90, and HepG2) and four  
100 cancer patient samples (two colon adenocarcinomas (COAD) and two kidney renal  
101 papillary cell carcinomas (KIRP)). We implemented a cell type-specific training  
102 approach to account for the distinct regulatory mechanisms that govern chromatin  
103 accessibility in different cellular contexts. Specifically, separate models were trained  
104 and evaluated independently for each cell line or cancer patient sample, without shar-  
105 ing parameters across cell types. To ensure statistical robustness, we implemented a  
106 5-fold cross-validation strategy for each model within each cell type or tumor sam-  
107 ple. The validation folds were stratified by chromosomes, allowing us to evaluate  
108 model performance on entirely unseen genomic regions while maintaining the unique  
109 characteristics of each cell type.

110 Within the test chromosomes, we evaluated model performance using two com-  
111plementary measurements: prediction accuracy on the reported ATAC-seq peaks and  
112 prediction accuracy on entire genomic regions. This dual assessment allowed us to  
113 benchmark both the models' ability to accurately characterize highly accessible chro-  
114matin regions (peak-focused evaluation) and their capacity to distinguish between  
115 accessible and inaccessible regions across the complete chromosomal landscape  
116 (genome-wide evaluation).



**Fig. 2 Model comparison across datasets.** (A) Pearson's correlation between true and predicted ATAC-seq signals in peak regions across eight distinct datasets for state-of-the-art models. (B) Performance comparison among the four models proposed in this work (ConvNeXtCNNs, ConvNeXtLSTMs, ConvNeXtDCNNs, and ConvNeXtTransformers) for the ATAC-seq peak regions stratified by cell lines and primary tumor samples. Each point depicts a test chromosome. The black dashed line shows the average performance of a model across all datasets and chromosomes. (C) Improvements of the new ConvNeXt-based methods proposed in this work for the ATAC-seq peaks as compared to existing methods. The significance is calculated with a two-sided Mann-Whitney U test on Pearson's R calculated for each test chromosome. \*\*\*\*:  $P \leq 0.0001$ , \*\*\*:  $P \leq 0.001$ , \*\*:  $P \leq 0.01$ , \*:  $P \leq 0.05$ , ns:  $P > 0.05$ . The  $\Delta$  Pearson's R is calculated as the difference between mean Pearson's R across all chromosomes for a ConvNeXt-based method and the corresponding existing method. (D) The total training time (in hours) of each method on a single RTX2080Ti GPU averaged across eight datasets. The error bar shows 95% confidence interval. Median training time is mentioned over each bar. (E) The F1 score is calculated for all the methods to evaluate the retrieval of ATAC-seq peak calls from the predictions. Each point represents one of the eight datasets.  $P$ -values are calculated using the Wilcoxon signed-rank test. (F) The average shape of the predicted ATAC-seq signal is compared against the experimental shape in the ATAC-seq peak regions of the GM12878 cell line. (G) Cosine similarity between experimental ATAC-seq and predicted ATAC-seq in the peak regions of the GM12878 cell line. (H) An example of predicted signal vs. the experimental signal around the *FMNL1* gene in the HepG2 cell line. A distal enhancer, discovered as a peak only based on the predictions of the ConvNeXtDCNN method, is highlighted.

117 Among the architectures employed in existing methods (CNNs, LSTMs, dCNNs),  
118 dCNNs consistently outperformed others in both peak-focused and genome-wide set-  
119 tings (Fig. 2A, Supplementary Fig. S2). However, upon incorporating ConvNeXt V2  
120 blocks in the stem of these architectures, all models demonstrated similar high corre-  
121 lations with the experimental ATAC-seq values in the accessible regions as well as on  
122 whole genome (average Pearson's  $R \geq 0.7$ ) (Fig. 2B, Supplementary Fig. S2).

123 The novel transformer-based architecture, ConvNeXtTransformer, that we also  
124 introduced in this evaluation, performed comparably to the ConvNeXtCNN and  
125 ConvNeXtDCNN models. We also explored the xLSTM architecture, a more recent  
126 development than the transformer, which offers improved memory efficiency for  
127 sequential data processing (Beck et al., 2024). However, when tested in our frame-  
128 work, ConvNeXt-xLSTM was outperformed by ConvNeXtLSTM, suggesting that its  
129 recurrent, long-range dependency modeling impaired the capture of local context crit-  
130 ical for ATAC-seq prediction compared to ConvNeXtTransformer's attention-based  
131 approach (Supplementary Fig. S3). Overall, incorporating ConvNeXt blocks into the  
132 stem consistently improved the performance of the existing models across all datasets,  
133 although the effect size was modest when comparing ConvNeXtDCNN to the dilated  
134 CNN (Fig. 2C, Supplementary Fig. S2). This underscores the better effectiveness of  
135 ConvNeXt blocks in capturing genomic features with local and global context.

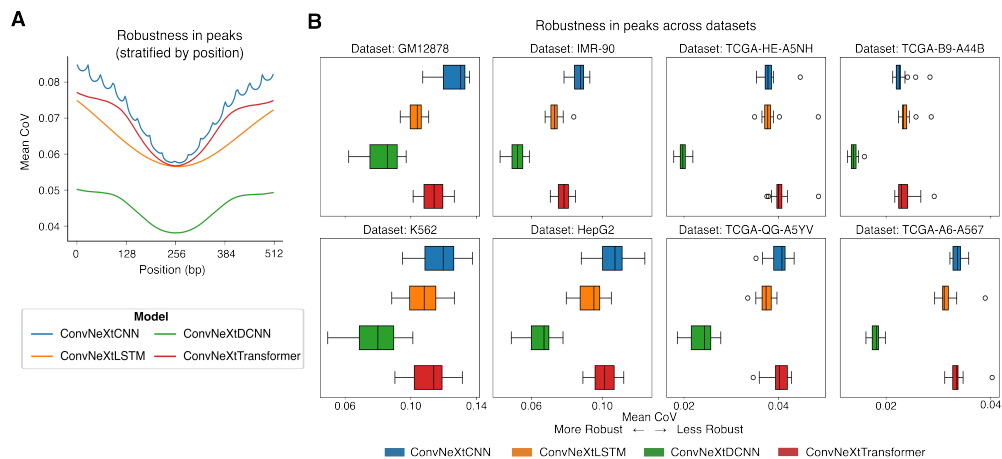
136 Adding ConvNeXt blocks to the stem, however, increased training times, with Con-  
137 vNeXtDCNN requiring approximately 1.7 more hours as compared to dCNN (Fig. 2D).  
138 Despite this, all models exhibited fast inference, predicting ATAC-seq signal across  
139 the entire Chromosome 17 in under 2 minutes, underscoring their practical utility for  
140 variant-effect prediction tasks (Supplementary Fig. S2).

## 141 **Assessing peak detection and signal shape fidelity in predicted** 142 **ATAC-seq tracks**

143 Accurate identification of ATAC-seq peaks is essential for deep learning models to  
144 generate outputs suitable for downstream tasks. We performed peak calling on the  
145 predicted ATAC-seq signals from each model (Methods) and compared the resulting  
146 peaks against experimental data. For robust evaluation, we used the F1 score, which  
147 accounts for class imbalance. ConvNeXt-based models outperformed existing baselines,  
148 with ConvNeXtDCNN achieving the highest F1 score (Fig. 2E).

149 To assess how well these models capture signal shapes around the peak maxima, we  
150 plotted the min-max normalized predicted ATAC-seq signal averaged across all peaks,  
151 datasets, and training folds (Fig. 2F). While ConvNeXtLSTM tended to over-predict  
152 the signal at peak boundaries, producing smoother profiles, other ConvNeXt-based  
153 methods closely matched the experimental ATAC-seq signal, accurately capturing  
154 nucleosome positioning around peak centers. This qualitative observation was further  
155 supported quantitatively, as the other ConvNeXt-based methods showed higher cosine  
156 similarity with the experimental ATAC-seq signal than ConvNeXtLSTM (Fig. 2G).

157 Finally, we visualized the chromatin accessibility signals generated by each  
158 ConvNeXt-based model on test chromosomes and their corresponding peaks (Fig. 2H).  
159 The predicted signals and peaks co-localized with many known promoters and  
160 enhancers (Moore et al., 2020), corroborating quantitative overlap analyses with  
161 ground-truth accessible regions and confirming the high fidelity of predicted peak  
162 shapes (Fig. 2E-H). Notably, some enhancers were uniquely identified by specific meth-  
163 ods; for example, an enhancer in the gene body of *FMNL1-DT* in HepG2 cells was  
164 marked as accessible only by ConvNeXtDCNN (Fig. 2H). Overall, ConvNeXtDCNN  
165 was the top-performing method based on its peak-calling accuracy and correlation  
166 with experimental data, with other ConvNeXt-based models also showing strong  
167 performance.



**Fig. 3 Evaluation of model robustness to input shifts. (A)** Position-stratified coefficient of variation (CoV) is calculated for all the methods for the ATAC-seq peak regions. The variation is averaged across all the peaks in unseen chromosomes of the 8 datasets. **(B)** The mean CoV measured for ATAC-seq peak regions across the eight datasets used in this study. Each point represents mean CoV for a chromosome unseen by the model during training.

## 168 Evaluating model robustness to shifts in the DNA sequence 169 inputs

170 A critical property of reliable genomic models is robustness to input shifts, *i.e.*, the  
171 ability to generate consistent predictions regardless of whether the input sequence is  
172 precisely centered on a peak maximum or slightly offset. To quantify this property, we  
173 performed a systematic robustness test on all ConvNeXt-based models, measuring the  
174 variation in predicted ATAC-seq signal across common genomic regions when the input  
175 window is shifted (Toneyan et al., 2022) (Supplementary Fig. S4). For each ATAC-seq  
176 peak across all test chromosomes and training folds, we generated 17 different input  
177 shifts and calculated the coefficient of variation across each 4 bp bin in the commonly  
178 predicted region. Lower coefficient of variation corresponded to higher robustness,  
179 indicating more consistent predictions despite slight shifts in input positioning.

180 The position-stratified coefficient of variation, averaged across all peaks in differ-  
181 ent datasets, showed that all methods displayed higher robustness at peak centers  
182 (Fig. 3A). The periodic variation pattern observed in ConvNeXtCNN can be attributed

183 to the inherent sensitivity of convolutional networks to random input shifts (R. Zhang,  
184 2019).

185 Among all the methods, ConvNeXtDCNN consistently emerged as the most robust  
186 method across all datasets, followed by ConvNeXtLSTM (Fig. 3B-C). The robustness  
187 of ConvNeXtDCNN was further confirmed by a low variation score observed using the  
188 entire unseen chromosomes of each dataset (Supplementary Fig. S4).

## 189 Interpreting sequence features underlying model predictions

190 To gain mechanistic insights into ConvNeXt-based architectures, we performed feature  
191 attribution using integrated gradients (Sundararajan et al., 2017) followed by motif  
192 discovery (Methods). This analysis revealed that the ConvNeXt-based models consis-  
193 tently learned sequence motifs corresponding to transcription factors with known roles  
194 in regulating chromatin accessibility in each cellular context (Fig. 4A-D). For exam-  
195 ple, in K562, motifs of erythroid regulators such as GATA1 and MAFK were highly  
196 enriched, while in HepG2, motifs of hepatic factors, including HNF4G and FOXA2  
197 predominated. Similar enrichment patterns were observed in GM12878 and IMR-90,  
198 where motifs of transcription factors such as SP1 and CTCF emerged prominently.  
199 The majorly identified transcription factors for each cell type were consistent with  
200 the findings of ChromBPNet (Pampari et al., 2024) (Supplementary Fig. S5). These  
201 findings indicate that the ConvNeXt V2 block effectively captures cell type-specific  
202 regulatory features that underlie chromatin accessibility.

203 We next examined attribution profiles at individual genomic loci to evaluate  
204 how these motifs contribute to local predictions. At a particular locus on Chromo-  
205 some 4 in K562 cells, ConvNeXtCNN, ConvNeXtDCNN, and ConvNeXtTransformer  
206 strongly highlighted GATA1 and MAFK motifs in close agreement with experimen-  
207 tal ChIP-seq binding, whereas ConvNeXtLSTM produced weaker and less localized  
208 attributions (Fig. 4E). Similarly, in HepG2, ConvNeXt-based methods accurately

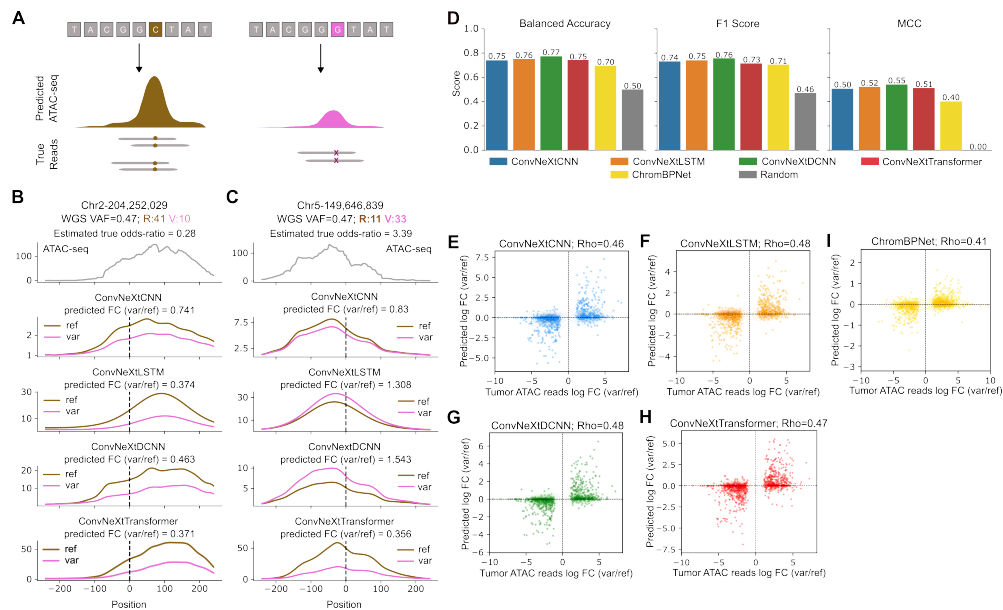


## 217 **Evaluating model sensitivity to genetic point variants**

218 Experimentally profiling chromatin accessibility for every genetic variant is pro-  
219 hibitively resource-intensive, highlighting the need for accurate in silico prediction of  
220 variant effects. To evaluate the ability of benchmarked models to capture the effects  
221 of genomic variants, we obtained SNVs from four cancer patients, two with colon  
222 adenocarcinoma and two with kidney renal papillary cell carcinoma, using data from  
223 the ICGC (J. Zhang et al., 2019a), and estimated variant-induced changes on chro-  
224 matin accessibility based on ATAC-seq read counts (Fig. 5A). A total of 1,025 SNVs  
225 were analyzed after filtering out variants with minimal effect, defined as those with  
226 odds ratios between 0.67 and 1.33. For this task, all models were retrained on the full  
227 genome, excluding only regions containing SNVs. Model predictions were then gener-  
228 ated using 2,048 bp SNV-centered input sequences, where the reference and variant  
229 sequences differed by a single nucleotide.

230 We first qualitatively observed the predicted ATAC-seq profiles by each of the  
231 ConvNeXt-based models using reference and alternate alleles as input. While all mod-  
232 els generally produced consistent changes in accessibility between the two alleles,  
233 ConvNeXtDCNN most accurately recapitulated the fine-grained shape of the predicted  
234 signal (Fig. 5B, Supplementary Fig. S6). An exception was observed at an SNV near  
235 149.6 Mb on Chromosome 5, where only the ConvNeXtLSTM and ConvNeXtDCNN  
236 models correctly predicted the direction of change, with ConvNeXtDCNN providing  
237 a prediction closer to the ground-truth odds ratio (Fig. 5C).

238 We then assessed the ability of each model to correctly predict the directionality of  
239 change in chromatin accessibility between the reference and variant alleles across all  
240 SNVs. All models outperformed the state-of-the-art method ChromBPNet (Pampari  
241 et al., 2024) and the random baseline, where for each SNV, a value of 0 or 1 was  
242 randomly assigned according to the proportion of SNVs with positive and negative  
243 effects on experimental ATAC-seq (Fig. 5D). Moreover, ConvNeXtDCNN achieved the



**Fig. 5 Evaluating model ability to accurately predict effects of single-nucleotide variants on chromatin accessibility.** (A) Experimental setup to test each model's ability to predict the change in chromatin accessibility between the reference DNA sequence (left) and a genomic variant (right). (B-C) Examples of allele-specific chromatin accessibility predictions for the reference allele and a genomic variant in Chromosomes 2 and 5, respectively. WGS VAF: Whole genome sequencing variant allele frequency; R, V: ATAC-seq total reference reads and variant reads, respectively. The model predictions are smoothed using a 1D Gaussian kernel with  $\sigma = 8$ . (D) ConvNeXt-based models are compared against ChromBPNet based on their ability to predict the correct directionality of change in accessibility between reference and variant. A total of 1,025 variants are considered across all cancer patients included in this study. Balanced accuracy, F1 score, and Matthews correlation coefficient (MCC) are used as metrics. A random baseline is included to demonstrate the improvements in metrics across all models proposed in this work. (E-H) Scatter plots showing the quantitative change in accessibility by each of the ConvNeXt-based models. (I) Performance of ChromBPNet is also shown for reference. Log-odds-ratio based on ATAC-seq read counts for reference and variants are compared against log-fold change (FC) between model predictions with reference and variant as model inputs and Spearman's correlation ( $Rho$ ) is calculated.

244 highest performance with balanced accuracy of 0.77, F1 score of 0.76, and Matthews  
 245 correlation coefficient (MCC) of 0.55 (Fig. 5D). Both F1 score and MCC account for  
 246 class imbalance and are, therefore, more robust metrics for this task.

247 We next tackled the more challenging task of quantitatively predicting the change  
 248 in chromatin accessibility induced by an SNV. For this evaluation, the estimated log-  
 249 odds ratio from ATAC-seq read counts for the reference and variant alleles served  
 250 as the ground truth. This was compared against the log-fold-change between the

251 ATAC-seq signal predicted by each model for the variant and the reference nucleotide  
252 (Methods). ConvNeXtLSTM and ConvNeXtDCNN achieved the highest Spearman's  
253 correlation with the true odds ratio ( $Rho = 0.48$ ) (Fig. 5E-H), closely followed by  
254 ConvNeXtTransformer ( $Rho = 0.47$ ) and ConvNeXtCNN ( $Rho = 0.46$ ). All methods  
255 outperformed the ChromBPNet method (Fig. 5I), demonstrating that incorporating  
256 ConvNeXt-based feature extraction improves sensitivity to sequence variation.

257 In conclusion, while all models performed well on the task of predicting the effect  
258 of SNVs on chromatin accessibility, ConvNeXtDCNN had a slight edge over the other  
259 methods.

## 260 Discussion

261 In this work, we compared existing deep learning architectures - CNNs, recurrent  
262 neural networks (specifically, LSTMs), and dCNNs, on their ability to predict high-  
263 resolution chromatin accessibility signals from genome sequences (Fig. 1A). To improve  
264 model performance, we proposed new ConvNeXt-based models: ConvNeXtCNN, Con-  
265 vNeXtLSTM, and ConvNeXtDCNN. Additionally, we introduced a transformer-based  
266 model, ConvNeXtTransformer, to create a more comprehensive set of architectures for  
267 this task (Fig. 1B).

268 The ConvNeXt-based models outperformed their traditional counterparts in accu-  
269 rately predicting ATAC-seq signals at 4 bp resolution across eight different cell types  
270 and primary tumor samples. Among them, ConvNeXtCNN, ConvNeXtDCNN, and  
271 ConvNeXtTransformer most accurately captured the shape of the ATAC-seq peaks  
272 (Fig. 2). We further evaluated the robustness of these models to shifts in the DNA  
273 sequence inputs. ConvNeXtDCNN emerged as the most robust method, efficiently  
274 integrating local and global sequence features through its dilated CNN layers (Fig. 3).

275 Using the feature attribution measures, we observed that the ConvNeXt-based  
276 models (ConvNeXtCNN, ConvNeXtLSTM, ConvNeXtDCNN, and ConvNeXtTrans-  
277 former) learn biologically meaningful sequence features (Fig. 4). The previously  
278 observed Tn5 insertion bias in signal profile prediction was observed only in the Con-  
279 vNeXtDCNN model when trained and evaluated on the ATAC-seq from HepG2 cells  
280 (Pampari et al., 2024). In this experiment, out of all motifs that were considered impor-  
281 tant by the model based on the attribution scores (Methods), only 3.5% corresponded  
282 to the Tn5 motif (Fig. 4D).

283 Finally, we assessed the model’s ability to predict changes in chromatin accessibility  
284 caused by single-nucleotide genomic variants. ConvNeXtLSTM and ConvNeXtDCNN  
285 achieved the highest accuracy in predicting the log-odds-ratio of ATAC-seq reads  
286 between variant and reference alleles. ConvNeXtDCNN also best captured the direc-  
287 tionality of accessibility changes, accurately distinguishing between increases and  
288 decreases in chromatin accessibility upon SNV introductions (Fig. 5). This highlights  
289 the sensitivity of the ConvNeXtDCNN model to genomic variants and suggests its  
290 potential utility in studying the functional effects of non-coding genetic variants.

291 While our ConvNeXt-based models demonstrated strong performance in predicting  
292 chromatin accessibility and the effects of SNVs, our models were trained in a cell-type-  
293 specific manner. Therefore, they were optimized for individual cellular contexts rather  
294 than generalized across different cell types. While this design ensures high accuracy  
295 within a given cell type, it also means that a new model must be trained for each  
296 additional cell type of interest. Future work could explore data-efficient fine-tuning  
297 approaches to accelerate model adaptation to new cellular contexts without requiring  
298 full retraining.

299 Overall, our study highlighted the advantages of ConvNeXt-based models in  
300 predicting high-resolution chromatin accessibility and regulatory effects of genetic  
301 variants. By combining ConvNeXt blocks as genomic feature extractors with CNNs,

302 LSTMs, dilated CNNs, and transformers, we established a comprehensive framework of  
303 sequence-based chromatin modeling. Our results demonstrated that ConvNeXtDCNN  
304 is not only highly accurate but also robust to input shifts, capture biologically mean-  
305 ingful features, and effective at capturing effects of SNVs on chromatin accessibility,  
306 making ConvNeXtDCNN a promising method for genome-wide regulatory analysis.

## 307 **Methods**

### 308 **ATAC-seq bigWig data**

309 We downloaded the ATAC-seq signal *P*-value tracks for four cell lines - GM12878,  
310 IMR-90, HepG2, and K562, from the ENCODE Consortium (Kagda et al., 2023) -  
311 where the experimentally measured ATAC-seq has been preprocessed with the open  
312 source processing pipeline (Lee et al., 2016). These cell lines consist of healthy and  
313 cancer cells, thereby creating a diverse dataset.

314 To further diversify our validation datasets, we used ATAC-seq data of two colon  
315 adenocarcinoma (COAD) and two kidney renal papillary cell carcinoma (KIRP)  
316 tumors from The Cancer Genome Atlas (TCGA) (Corces et al., 2018). These datasets  
317 have been made available by the Genome Data Commons (GDC) Data Portal (Gross-  
318 man et al., 2016). The choice of these patients was made due to the availability of the  
319 data on their somatic mutations and single-nucleotide polymorphisms (SNPs) in the  
320 International Cancer Genome Consortium (ICGC) and GDC data portals, respectively  
321 (Grossman et al., 2016; J. Zhang et al., 2019b). Moreover, these patients had the high  
322 number of combined SNV calls ( $> 175,000$ ) before applying SNV-filtering (described  
323 below) and high tumor purity (Consensus Purity Estimate  $> 0.8$ ) (Aran et al., 2015).

324 The ATAC-seq BAM files for the primary tumor samples were downloaded from  
325 National Cancer Institute's Genomic Data Commons portal (Heath et al., 2021). Signal  
326 *P*-value bigWig files were then generated using the ENCODE's ATAC-seq processing

327 pipeline (Lee et al., 2016). All datasets used in this study are listed in Supplementary  
328 Table S1.

## 329 Preprocessing

330 For the ATAC-seq of each cell line, we composed a dataset for each cell line or can-  
331 cer patient where each item is a pair of one-hot encoded DNA sequence of shape  
332 (2048, 4) and the corresponding 4 bp resolution ATAC-seq of shape (512). To reduce  
333 the ATAC-seq resolution from 1 bp to 4 bp in the dataset, the values were binned  
334 using the maximum value.

335 Each model was trained on whole-genome data from a single cell line or cancer  
336 patient sample, excluding input DNA sequences corresponding to genomic locations  
337 with unreliable accessibility measurements to reduce noise and biases. Specifically, we  
338 excluded input DNA sequences overlapping with blacklisted regions ([https://storage.  
339 googleapis.com/basenji\\_barnyard2/hg38.blacklist.rep.bed](https://storage.googleapis.com/basenji_barnyard2/hg38.blacklist.rep.bed)) or where less than or equal  
340 to 35% of the 2048 bp input overlapped with unmappable regions ([https://storage.  
341 googleapis.com/basenji\\_barnyard2/umap\\_k36\\_t10\\_l32\\_hg38.bed](https://storage.googleapis.com/basenji_barnyard2/umap_k36_t10_l32_hg38.bed)). The annotations for  
342 blacklisted and unmappable regions of the hg38 genome were sourced from Kelley et  
343 al., 2018. Reverse-complemented DNA sequences were augmented to the training data  
344 to enhance model robustness.

345 During the evaluation, we imposed a tighter constraint on the mappability and  
346 only evaluated the input DNA sequences that did not overlap with any unmappable  
347 region. Additionally, we used only the predictions corresponding to the central 1024bp  
348 genomic section. This ensured that there was at least 512 bp of genomic context present  
349 on either side of any region that we were trying to predict. Each model was trained  
350 on the full 512-length predicted output vector, but for evaluation, we restricted the  
351 analysis to the central 256 entries to avoid edge effects.

352 For each model trained on a particular cell line or cancer patient sample, we eval-  
353 uate the model on whole genome data and on a dataset composed of peaks by taking  
354 input DNA sequences centered on the MACS2-called ATAC-seq peaks to compare how  
355 models perform on high signal regions.

## 356 **Training details**

357 We did a 5-fold cross validation where the genome was randomly split into 5 sets of  
358 chromosome – [1, 11, 13, 20], [2, 10, 14, 19, 21], [3, 12, 16, 17, 22], [6, 7, 9, 15], [4,  
359 5, 8, 18]. For each fold, a model was trained for a maximum of 70 epochs with early  
360 stopping, with a patience of 5 epochs based on Pearson’s correlation on validation  
361 data.

362 Poisson negative log-likelihood loss was used to train all models with the AdamW  
363 optimizer (Loshchilov, Hutter, et al., 2017). We began by linearly increasing the learn-  
364 ing rate during the first epoch until we reached a maximum of  $1 \times 10^{-3}$ . We then  
365 used exponential decay coupled with cosine annealing (Loshchilov and Hutter, 2016),  
366 restarting at the beginning of every epoch.

## 367 **Model architectures**

368 All models processed one-hot encoded DNA sequences of 2048 bp and predicted ATAC-  
369 seq signals corresponding to the central 1024 bp at 4 bp resolution. Each model is  
370 illustrated in Supplementary Fig. S1. The runtime and memory requirements of each  
371 model are described in Fig. 2D and Supplementary Fig. S2.

## 372 **CNN**

373 The CNN, inspired by DeepSEA (Zhou and Troyanskaya, 2015), was adapted for 2048  
374 bp inputs and outputs. It comprised three convolutional layers with 512, 640, and 768  
375 filters (kernel size 8), each followed by max pooling (window 4) and dropout (rate  
376 0.2) to reduce dimensionality and prevent overfitting. The flattened output fed into

377 two dense layers with 1024 and 512 ReLU-activated units. This architecture excels at  
378 detecting local motif patterns, such as transcription factor binding sites, indicative of  
379 chromatin accessibility.

### 380 **LSTM**

381 The recurrent model (LSTM), based on DanQ (Quang and Xie, 2016), processed inputs  
382 through a convolutional stem with 512 filters (kernel size 26), followed by max pooling  
383 (window 4) and dropout (rate 0.2). A bidirectional LSTM layer (Schuster and Paliwal,  
384 1997) with 512 units captured long-range dependencies in the sequence. Two dense  
385 layers (1024 and 512 ReLU-activated units) produced the final output. The LSTM's  
386 memory enables modeling sequential relationships across the DNA.

### 387 **dCNN**

388 The dilated CNN (dCNN), inspired by Basenji (Kelley et al., 2018) and adapted by  
389 ChromBPNet (Pampari et al., 2024) for 2048 bp inputs, used 11 residual blocks. Each  
390 block contained two convolutional layers, the first with a dilation rate increasing by  
391 a factor of 1.5 across blocks, followed by GELU activation (Hendrycks and Gimpel,  
392 2016). A dense layer reduced the channel dimension to 1, and a softplus activation  
393 ensured positive outputs. This design captures long-range genomic interactions.

394

395 To enhance feature extraction, we replaced traditional convolutional blocks with  
396 ConvNeXt V2 blocks (Woo et al., 2023), which offer improved efficiency and  
397 expressiveness via depthwise convolutions and layer normalization.

### 398 **ConvNeXtCNN**

399 This model used a ConvNeXt stem with 128 filters (kernel size 15), followed by three  
400 ConvNeXt blocks with 256, 512, and 512 filters (kernel size 15), each applying depth-  
401 wise convolutions (Chollet, 2017). Max pooling (window 4) and dropout (rate 0.2)

402 followed each block. Two dense layers (512 and 32 ReLU-activated units) processed  
403 the feature matrix, with a softplus activation ensuring positive outputs.

#### 404 **ConvNeXtLSTM**

405 The ConvNeXtLSTM employed a ConvNeXt stem with 256 filters (kernel size 26),  
406 followed by a bidirectional LSTM layer (Schuster and Paliwal, 1997) with 256 units  
407 to model distal genomic interactions. The final LSTM state fed into two dense layers  
408 (256 and 512 ReLU-activated units), with a softplus activation for the output.

#### 409 **ConvNeXtDCNN**

410 Similar to dCNN, the ConvNeXtDCNN used 11 dilated residual blocks but replaced  
411 the initial convolutional layer with a ConvNeXt stem (256 filters, kernel size 15). The  
412 remaining architecture, including dilation rates and softplus output, mirrored dCNN.

#### 413 **ConvNeXtTransformer**

414 We introduced the first transformer-based model for ATAC-seq prediction. Five Con-  
415 vNeXt blocks (256 filters, kernel size 15) and pooling layers tokenized the input into  
416 512 sequences at 4 bp resolution. This sequence of 512 tokens was then passed through  
417 four transformer blocks. Mirroring the behavior seen in current large-language mod-  
418 els (Jiang et al., 2023; Touvron et al., 2023), we observed better performance when  
419 using rotary positional encodings (RoPE) (Su et al., 2024) compared to conventional  
420 sinusoidal encodings proposed in Vaswani et al., 2017. A dense layer aggregated the  
421 (512, C) output across channels, followed by a softplus activation.

#### 422 **ConvNeXt-xLSTM**

423 Lastly, we built an extension of the above methods by using xLSTM in the core of  
424 the model. xLSTM is a newly built model that extends LSTM by retaining memory  
425 across different states, thereby capturing both local and global contexts (Beck et al.,

426 2024). The input DNA sequence was passed through 5 ConvNeXt blocks followed by  
427 an xLSTM block consisting of three mLSTM and one sLSTM blocks with a kernel  
428 size of 4 and 4 heads (Supplementary Fig. S3). The output was then passed through  
429 a dense layer with 128 hidden units and softplus activation function.

## 430 Hyperparameter tuning

431 To optimize model performance for ATAC-seq signal prediction, we tuned hyperpa-  
432 rameters for each method by maximizing Pearson’s correlation on the validation set  
433 of fold 0, comprising held-out chromosomes. For all models, we tested batch sizes of  
434 64, 128, and 256, initial learning rates of 0.005, 0.001, 0.0005, and 0.0001, and a cosine  
435 annealing scheduler with warmup ratios of 0.1, 0.25, or 0.5 of total epochs (Loshchilov  
436 and Hutter, 2016).

437 For CNN-based methods (CNN and ConvNeXtCNN), we tuned the number of  
438 convolutional layers in the core block (2, 3, or 4), the kernel size in the core block (3,  
439 7, or 15), and the number of dense layers in the output block. All combinations of  
440 these parameters were tested. Increasing the dense layers from 2 to 3 did not improve  
441 performance, so we retained 2 layers.

442 For LSTM-based methods (LSTM and ConvNeXtLSTM), we tuned the kernel size  
443 of the initial convolutional stem (3, 7, 15, or 26, inspired by DanQ (Quang and Xie,  
444 2016)), the number of bidirectional LSTM layers (1 or 2), the number of dense layers  
445 (2 or 3), and their hidden sizes (256, 512, or 1024).

446 For dilated-CNN-based methods (dCNN and ConvNeXtdCNN), we explored archi-  
447 tectures with 7 to 15 convolutional residual blocks, inspired by Enformer (Avsec,  
448 Agarwal, et al., 2021) and ChromBPNet (Pampari et al., 2024). We varied the dila-  
449 tion rate (1 to 2, in steps of 0.1) and tested three convolutional block configurations  
450 within each residual block. These configurations used kernel counts of 128 to 512 (step

451 size 64), 64 to 256 (step size 64), or 256 to 4096 (step size 256), with corresponding  
452 kernel sizes of 5 to 20, 2 to 5, or 1 to 4, respectively.

453 For the ConvNeXtTransformer, we tuned the ConvNeXt stem with 1 to 4 layers,  
454 kernel sizes of 7, 11, 15, or 26, and filter counts of 64, 128, or 256. The core ConvNeXt  
455 blocks were tested with 2 to 6 layers, kernel sizes of 3 to 11, and inverted bottleneck  
456 scales of 0.5, 1, 2, or 4, which adjust channel expansion in ConvNeXt’s depthwise  
457 convolutions (Woo et al., 2023). The transformer core was optimized with 2 to 6 layers  
458 and feedforward dimension scales of 0.5, 1, 2, or 4.

## 459 **Adding mappability information to models**

460 Experimental ATAC-seq signals are influenced by the mappability of the region of  
461 interest. While deep learning models can implicitly learn mappability from DNA  
462 sequences, explicitly providing this information can enhance certain models. To achieve  
463 this, we predicted mappability values from the initial layers of each model alongside  
464 the 4 bp resolution ATAC-seq signal (Supplementary Fig. S7).

465 The mappability values were calculated as a binary vector based on overlap with  
466 Basenji’s unmappable regions (Kelley et al., 2018). Zeros were assigned to the genomic  
467 segments overlapping with unmappable regions, and ones were assigned otherwise.

468 Adding this mappability information stabilized CNN and LSTM models but had  
469 no significant effect on dCNN or ConvNeXt-based models (Supplementary Fig. S7).  
470 This suggests that architectures capturing long-range contextual information, such as  
471 ConvNeXt, may implicitly encode mappability due to their larger receptive fields.

## 472 **Peak calling**

473 Predicted ATAC-seq signals from each model were saved as bigWig files using the  
474 pyBigWig Python package (Ryan, Roberts, et al., 2023). These were then converted  
475 to bedGraphs using `bigWigToBedGraph` from UCSC Genome Browser utilities (Kent

476 et al., 2010) and sorted by chromosome and start position. Finally, peaks were called  
477 from the sorted bedGraphs using MACS2 `bdgpeakcall` command (version 2.2.9.1)  
478 (Y. Zhang et al., 2008) with a minimum score threshold of 1 to filter low-confidence  
479 regions.

## 480 **Variation score**

To calculate the robustness of each model to input shifts, we used the variation score as described in Toneyan et al., 2022. For each 512 bp region in our test dataset, we constructed 17 distinct 2048 bp long input sequences such that the output corresponding to each input contains the predictions for the selected 512 bp region. Using these 17 predictions, the coefficient of variation was calculated for each 4 bp bin, *i.e.*, for each bin we calculated the mean and standard deviation of the predictions. The coefficient of variation (CV) for a given bin was then defined as

$$CV_i = \frac{\sigma_i}{\mu_i},$$

481 where  $\sigma_i$  and  $\mu_i$  denote the standard deviation and mean, respectively, of the 17  
482 predictions for bin  $i$ .

483 Finally, the variation score was obtained by taking the average coefficient of vari-  
484 ation across all the bins. For the whole-genome evaluation, this was repeated for all  
485 the 512 bp regions of the test chromosomes whereas for peak-centered evaluation, only  
486 those 512 bp regions were selected that contained an ATAC-seq peak.

## 487 **Motif discovery**

To interpret sequence features underlying model predictions, we computed base-resolution feature attributions using integrated gradients (Sundararajan et al., 2017). Attribution scores were calculated with respect to the input DNA sequence for each prediction. Following the approach described in Pampari et al., 2024, the attribution

of each input feature was defined as its effect on the weighted sum of the predicted profile. Specifically, for the predicted profile  $\hat{y} \in \mathbb{R}^d$ , the attribution of each input nucleotide was calculated for the target

$$y = \sum_{j=1}^d \frac{e^{\hat{y}_j}}{\sum_k e^{\hat{y}_k}} \hat{y}_j$$

488 For computational efficiency, attributions were calculated on 30,000 randomly sampled  
 489 peaks per cell line. Integrated gradients were implemented with the PyTorch Captum  
 490 library (Kokhlikyan et al., 2020), using 50 interpolation steps and 20 dinucleotide-  
 491 shuffled sequences as the baseline.

492 The resulting attribution or importance maps were processed with TF-MoDISCO  
 493 Lite (Shrikumar et al., 2018) using “-n 100,000 -w 500” parameters to cluster  
 494 recurring 100,000 high-attribution sequence patterns into representative motifs. Con-  
 495 sensus motifs were generated with the `tfmodisco-lite report` workflow, using  
 496 representative motifs with Tn5 variants obtained from the ChromBPNet repos-  
 497 itory ([https://github.com/kundajelab/chrombpnet/blob/master/chrombpnet/data/](https://github.com/kundajelab/chrombpnet/blob/master/chrombpnet/data/motifs.meme.txt)  
 498 [motifs.meme.txt](https://github.com/kundajelab/chrombpnet/blob/master/chrombpnet/data/motifs.meme.txt)). The discovered motifs and their associated transcription factors  
 499 were then reported to highlight the biological relevance of the features captured by  
 500 each model.

## 501 SNV filtering

502 We define a single-nucleotide variant (SNV) as either a somatic mutation or a single-  
 503 nucleotide polymorphism (SNP). For the four primary tumor samples, we obtained  
 504 somatic mutation calls from the study of pan-cancer analysis of whole genome  
 505 (PCAWG) (Goldman et al., 2020). The SNP calls were obtained from GDC data portal  
 506 (Grossman et al., 2016).

507 For each patient, the SNVs were filtered to retain only those SNVs that signifi-  
508 cantly changed the chromatin accessibility profiles. First, somatic mutations that did  
509 not pass the quality control measures of PCAWG were removed. Similarly, we only  
510 selected those SNPs that had a heterozygous genotype and a high confidence score  
511 (uncertainty value  $< 0.1$ ). The filtered files were then combined, and total read counts  
512 from the whole genome sequencing (WGS) and ATAC-seq experiments were calcu-  
513 lated for the reference and variant alleles for each SNV. Using these values, the variant  
514 allele frequency (VAF) for each SNV was calculated, giving us WGS VAF and ATAC  
515 VAF. As we estimated the expected change in ATAC-seq signal for each SNV using  
516 the ATAC-seq read counts, we expected the SNVs to be unbiased towards either allele.  
517 Therefore, we selected SNVs with WGS VAF in the range of  $[0.4, 0.6]$ . Finally, to obtain  
518 SNVs that have a significant change in ATAC-seq, we calculated the significance of  
519 change using the Fisher's exact test and the odds ratio. Each of these quantities was  
520 calculated using the reference and alternate allele read counts in WGS and ATAC-seq  
521 experiments. SNVs that had a  $P$ -value  $< 0.01$  and odds-ratio  $\in [0.67, 1.33]$  were con-  
522 sidered to have a significant change in ATAC-seq signal and therefore, were selected  
523 for SNV analysis.

524 This filtering led to a total SNV count of 1,025, with 80 somatic mutations and  
525 945 SNPs across the four cancer patients.

## 526 Software availability

527 ASAP is available as a Python package at <https://pypi.org/project/atac-asap/>,  
528 released under the MIT license. Detailed documentation and installation instructions  
529 are provided at <https://github.com/BoevaLab/ASAP/wiki>.

530 All results presented in this manuscript can be reproduced using the source code  
531 and scripts, which are provided as Supplemental Code and deposited on Zenodo at  
532 <https://doi.org/10.5281/zenodo.15377297>.

## 533 **Competing interest statement**

534 The authors declare no competing interests.

## 535 **Acknowledgments**

536 This project is partially funded by the Swiss Data Science Center (SDSC) collabora-  
537 tive projects grant (C22-09); Swiss Government Excellence Scholarship (ESKAS-Nr:  
538 2021.0468) to A.G.

539 **Author contributions:** V.B. and A.G. conceived the study. V.B. and E.K. super-  
540 vised the project. A.G. and L.K. processed the datasets. T.M. and L.K. developed the  
541 different neural network architectures with the assistance of L.Z. and A.G. The exper-  
542 iments were devised and interpreted collectively by all authors. L.K., T.M., and A.G.  
543 implemented the experiments and analyzed the results. A.G. composed the figures.  
544 A.G., T.M., and L.K. wrote the manuscript while all other authors provided feedback.

## **References**

- Aran, D., Sirota, M., & Butte, A. J. (2015). Systematic pan-cancer analysis of tumour purity. *Nature communications*, *6*(1), 8971.
- Avsec, Ž., Agarwal, V., Visentin, D., Ledsam, J. R., Grabska-Barwinska, A., Taylor, K. R., Assael, Y., Jumper, J., Kohli, P., & Kelley, D. R. (2021). Effective gene expression prediction from sequence by integrating long-range interactions. *Nature methods*, *18*(10), 1196–1203.
- Avsec, Ž., Weilert, M., Shrikumar, A., Krueger, S., Alexandari, A., Dalal, K., Froepf, R., McAnany, C., Gagneur, J., Kundaje, A., et al. (2021). Base-resolution models of transcription-factor binding reveal soft motif syntax. *Nature genetics*, *53*(3), 354–366.

- Beck, M., Pöppel, K., Spanring, M., Auer, A., Prudnikova, O., Kopp, M., Klambauer, G., Brandstetter, J., & Hochreiter, S. (2024). Xlstm: Extended long short-term memory. *arXiv preprint arXiv:2405.04517*.
- Buenrostro, J. D., Wu, B., Chang, H. Y., & Greenleaf, W. J. (2015). Atac-seq: A method for assaying chromatin accessibility genome-wide. *Current protocols in molecular biology*, *109*(1), 21–29.
- Cazares, T. A., Rizvi, F. W., Iyer, B., Chen, X., Kotliar, M., Bejjani, A. T., Wayman, J. A., Donmez, O., Wronowski, B., Parameswaran, S., et al. (2023). Maxatac: Genome-scale transcription-factor binding prediction from atac-seq with deep neural networks. *PLOS Computational Biology*, *19*(1), e1010863.
- Chollet, F. (2017). Xception: Deep learning with depthwise separable convolutions. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1251–1258.
- Corces, M. R., Granja, J. M., Shams, S., Louie, B. H., Seoane, J. A., Zhou, W., Silva, T. C., Groeneveld, C., Wong, C. K., Cho, S. W., et al. (2018). The chromatin accessibility landscape of primary human cancers. *Science*, *362*(6413), eaav1898.
- Dalla-Torre, H., Gonzalez, L., Mendoza-Revilla, J., Lopez Carranza, N., Grzywaczewski, A. H., Oteri, F., Dallago, C., Trop, E., de Almeida, B. P., Sirelkhatim, H., et al. (2024). Nucleotide transformer: Building and evaluating robust foundation models for human genomics. *Nature Methods*, 1–11.
- 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. (2024). Chromafold predicts the 3d contact map from single-cell chromatin accessibility. *Nature Communications*, *15*(1), 9432.

- Giresi, P. G., Kim, J., McDaniel, R. M., Iyer, V. R., & Lieb, J. D. (2007). Faire (formaldehyde-assisted isolation of regulatory elements) isolates active regulatory elements from human chromatin. *Genome research*, *17*(6), 877–885.
- Goldman, M. J., Zhang, J., Fonseca, N. A., Cortés-Ciriano, I., Xiang, Q., Craft, B., Piñeiro-Yáñez, E., O’Connor, B. D., Bazant, W., Barrera, E., et al. (2020). A user guide for the online exploration and visualization of pcawg data. *Nature communications*, *11*(1), 3400.
- Grossman, R. L., Heath, A. P., Ferretti, V., Varmus, H. E., Lowy, D. R., Kibbe, W. A., & Staudt, L. M. (2016). Toward a shared vision for cancer genomic data. *New England Journal of Medicine*, *375*(12), 1109–1112.
- Grover, A., Zhang, L., Muser, T., Häffiger, S., Wang, M., Yates, J., Van Allen, E. M., Theis, F. J., Ibarra, I. L., Krymova, E., et al. (2024). Universalepi: Harnessing attention mechanisms to decode chromatin interactions in rare and unexplored cell types. *bioRxiv*, 2024–11.
- Guo, Y., Zhou, D., Nie, R., Ruan, X., & Li, W. (2020). Deepanf: A deep attentive neural framework with distributed representation for chromatin accessibility prediction. *Neurocomputing*, *379*, 305–318.
- Heath, A. P., Ferretti, V., Agrawal, S., An, M., Angelakos, J. C., Arya, R., Bajari, R., Baqar, B., Barnowski, J. H., Burt, J., et al. (2021). The nci genomic data commons. *Nature genetics*, *53*(3), 257–262.
- Hendrycks, D., & Gimpel, K. (2016). Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Ji, Y., Zhou, Z., Liu, H., & Davuluri, R. V. (2021). Dnabert: Pre-trained bidirectional encoder representations from transformers model for dna-language in genome. *Bioinformatics*, *37*(15), 2112–2120.

- Jiang, A. Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D. S., Casas, D. d. l., Bressand, F., Lengyel, G., Lample, G., Saulnier, L., et al. (2023). Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Kagda, M. S., Lam, B., Litton, C., Small, C., Sloan, C. A., Spragins, E., Tanaka, F., Whaling, I., Gabdank, I., Youngworth, I., et al. (2023). Data navigation on the encode portal. *arXiv preprint arXiv:2305.00006*.
- Karbalayghareh, A., Sahin, M., & Leslie, C. S. (2022). Chromatin interaction-aware gene regulatory modeling with graph attention networks. *Genome Research*, *32*(5), 930–944.
- Kelley, D. R., Reshef, Y. A., Bileschi, M., Belanger, D., McLean, C. Y., & Snoek, J. (2018). Sequential regulatory activity prediction across chromosomes with convolutional neural networks. *Genome research*, *28*(5), 739–750.
- Kelley, D. R., Snoek, J., & Rinn, J. L. (2016). Basset: Learning the regulatory code of the accessible genome with deep convolutional neural networks. *Genome research*, *26*(7), 990–999.
- Kent, W. J., Zweig, A. S., Barber, G., Hinrichs, A. S., & Karolchik, D. (2010). Bigwig and bigbed: Enabling browsing of large distributed datasets. *Bioinformatics*, *26*(17), 2204–2207.
- Kokhlikyan, N., Miglani, V., Martin, M., Wang, E., Alsallakh, B., Reynolds, J., Melnikov, A., Kliushkina, N., Araya, C., Yan, S., & Reblitz-Richardson, O. (2020). Captum: A unified and generic model interpretability library for pytorch.
- Lee, J., Christoforo, G., Foo, C., Probert, C., Kundaje, A., Boley, N., Kim, D., Dacre, M., et al. (2016). Kundajelab/atac\_dnase\_pipelines: 0.3. 3. *Zenodo*.
- Liu, Q., Xia, F., Yin, Q., & Jiang, R. (2018). Chromatin accessibility prediction via a hybrid deep convolutional neural network. *Bioinformatics*, *34*(5), 732–738.
- Loshchilov, I., & Hutter, F. (2016). Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*.

- Loshchilov, I., Hutter, F., et al. (2017). Fixing weight decay regularization in adam. *arXiv preprint arXiv:1711.05101*, 5.
- Min, X., Zeng, W., Chen, N., Chen, T., & Jiang, R. (2017). Chromatin accessibility prediction via convolutional long short-term memory networks with k-mer embedding. *Bioinformatics*, *33*(14), i92–i101.
- Moore, J., Purcaro, M., Pratt, H., Epstein, C., Shores, N., Adrian, J., Kawli, T., Davis, C., Dobin, A., Kaul, R., et al. (2020). Expanded encyclopaedias of dna elements in the human and mouse genomes. *Nature*, *583*(7818), 699–710.
- Pampari, A., Shcherbina, A., Kvon, E. Z., Kosicki, M., Nair, S., Kundu, S., Kathiria, A. S., Risca, V. I., Kuningas, K., Alasoo, K., et al. (2024). Chrombpnet: Bias factorized, base-resolution deep learning models of chromatin accessibility reveal cis-regulatory sequence syntax, transcription factor footprints and regulatory variants. *bioRxiv*, 2024–12.
- Quang, D., & Xie, X. (2016). Danq: A hybrid convolutional and recurrent deep neural network for quantifying the function of dna sequences. *Nucleic acids research*, *44*(11), e107–e107.
- Ryan, D., Roberts, E., et al. (2023). Deeptools/pybigwig: 0.3. 22. *Zenodo* <https://doi.org/10.5281/zenodo.7809144>.
- Schuster, M., & Paliwal, K. K. (1997). Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, *45*(11), 2673–2681.
- Shrikumar, A., Tian, K., Avsec, Ž., Shcherbina, A., Banerjee, A., Sharmin, M., Nair, S., & Kundaje, A. (2018). Technical note on transcription factor motif discovery from importance scores (tf-modisco) version 0.5. 6.5. *arXiv preprint arXiv:1811.00416*.
- Song, L., & Crawford, G. E. (2010). Dnase-seq: A high-resolution technique for mapping active gene regulatory elements across the genome from mammalian cells. *Cold Spring Harbor Protocols*, *2010*(2), pdb–prot5384.

- Stewart, A. J., Hannehalli, S., & Plotkin, J. B. (2012). Why transcription factor binding sites are ten nucleotides long. *Genetics*, *192*(3), 973–985.
- Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., & Liu, Y. (2024). Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, *568*, 127063.
- Sundararajan, M., Taly, A., & Yan, Q. (2017). Axiomatic attribution for deep networks. *International conference on machine learning*, 3319–3328.
- Tan, J., Shenker-Tauris, N., Rodriguez-Hernaez, J., Wang, E., Sakellaropoulos, T., Boccalatte, F., Thandapani, P., Skok, J., Aifantis, I., Fenyő, D., et al. (2023). Cell-type-specific prediction of 3d chromatin organization enables high-throughput in silico genetic screening. *Nature biotechnology*, *41*(8), 1140–1150.
- Thibodeau, A., Khetan, S., Eroglu, A., Tewhey, R., Stitzel, M. L., & Ucar, D. (2021). Core-atac: A deep learning model for the functional classification of regulatory elements from single cell and bulk atac-seq data. *PLoS computational biology*, *17*(12), e1009670.
- Toneyan, S., Tang, Z., & Koo, P. K. (2022). Evaluating deep learning for predicting epigenomic profiles. *Nature machine intelligence*, *4*(12), 1088–1100.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al. (2023). Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, *30*.
- Woo, S., Debnath, S., Hu, R., Chen, X., Liu, Z., Kweon, I. S., & Xie, S. (2023). Convnext v2: Co-designing and scaling convnets with masked autoencoders. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16133–16142.

- Yang, M., & Ma, J. (2023). Unadon: Transformer-based model to predict genome-wide chromosome spatial position. *Bioinformatics*, *39*(Supplement\_1), i553–i562.
- Yang, R., Das, A., Gao, V. R., Karbalayghareh, A., Noble, W. S., Bilmes, J. A., & Leslie, C. S. (2023). Epiphany: Predicting hi-c contact maps from 1d epigenomic signals. *Genome Biology*, *24*(1), 134.
- Zhang, J., Bajari, R., Andric, D., Gerthoffert, F., Lepsa, A., Nahal-Bose, H., Stein, L. D., & Ferretti, V. (2019a). The international cancer genome consortium data portal. *Nature biotechnology*, *37*(4), 367–369.
- Zhang, J., Bajari, R., Andric, D., Gerthoffert, F., Lepsa, A., Nahal-Bose, H., Stein, L. D., & Ferretti, V. (2019b). The international cancer genome consortium data portal. *Nature biotechnology*, *37*(4), 367–369.
- Zhang, R. (2019). Making convolutional networks shift-invariant again. *International conference on machine learning*, 7324–7334.
- Zhang, Y., Liu, T., Meyer, C. A., Eeckhoute, J., Johnson, D. S., Bernstein, B. E., Nusbaum, C., Myers, R. M., Brown, M., Li, W., et al. (2008). Model-based analysis of chip-seq (macs). *Genome biology*, *9*, 1–9.
- Zhang, Z., Feng, F., Qiu, Y., & Liu, J. (2023). A generalizable framework to comprehensively predict epigenome, chromatin organization, and transcriptome. *Nucleic Acids Research*, *51*(12), 5931–5947.
- Zhou, J., & Troyanskaya, O. G. (2015). Predicting effects of noncoding variants with deep learning-based sequence model. *Nature methods*, *12*(10), 931–934.