Learning Composition Rules for Mammalian Circuits with Neural Attention

Jenhan Tao, Gregory J. Fonseca, Christopher K. Glass
Department of Cellular and Molecular Medicine
University of California, San Diego

Christopher Benner
Department of Medicine
University of California, San Diego

Abstract

The expression of each gene in mammalian cells is controlled by regulatory sequences called enhancers. Regulatory logic encoded at enhancers is interpreted by transcription factors (TFs), which recognize individual 'words' in the genome. Here we describe a neural network with an attention mechanism that learns to discriminate enhancers from random genomic sequences. We distill the parameters learned by our model to identify combinations of TF target sequences that signal activation of a gene in response to specific cellular stimuli.

Introduction

Regulation of gene expression in mammalian cells is mediated in part by hundreds of sequence specific TFs that bind to their individual binding motifs at enhancers, which are distal regulatory elements located thousands of basepairs away from a gene. The binding of TFs at enhancers mediates the recruitment of machinery necessary for transcription such as RNA polymerase. Prior studies have suggested two classes of TFs: 1) lineage determining TFs (LDTFs) and and 2) signal dependent TFs (SDTFs). LDTFs bind to cell type specific enhancers while SDTFs bind at enhancers bound by LDTFs in response to a cellular stimuli, resulting in cell type specific activation of an enhancer in response to stimuli (Heinz et al. 2013) (Figure 1). These studies suggest that context specific gene expression in a cell type is genetically encoded by combinations of TF binding motifs at millions of enhancers scattered throughout the genome (Consortium 2012).

Given the evidence that TFs act collaboratively to activate enhancers, it follows that individual TF motifs are poor predictors of whether or not a sequence is an enhancer. The biological activity of an enhancer may depend on the composition of TF motifs - arrangement and spacing between TF motifs, as well as the degeneracy of each motif (Farley et al. 2016). And so, we endeavored to teach an attentive neural network (ANN) to distinguish accessible enhancer elements from background genomic sequences in macrophage cells, a cell of the innate immune system. Our ANN uses a neural mechanism, which is at the heart of current models for modeling sequences of words in natural language processing applications such as language translation and sentiment analysis (Vaswani et al. 2017; Cheng, Dong, and Lapata 2016). In our ANN, neural attention allows the model to focus on the TF motifs that are functional at an enhancer and ignore dozens of other nonfunctional motifs (Figure 1). By extracting information learned by our network, we can identify combinations of TF motifs that signal the activation of an enhancer in response to cellular stimuli (Figure 1).

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

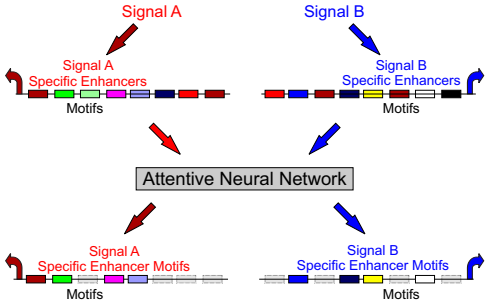

Figure 1: Attentive neural network learns to ignore nonfunctional motifs (faded gray boxes) thereby revealing TF motifs that control activation in response to signal A and B respectively

Model Design

The architecture of our ANN model is shown in (Figure 2). To learn words recognized by TFs, Our model applies a 1-dimensional convolution, $conv_m$ over 4 channels to the input sequence, s, encoded as a one hot vector (Alipanahi et al. 2015; Kelley, Snoek, and Rinn 2016). To learn relationships between words, we eschew recurrent neurons, which require many parameters that are hard to interpret, and use neural dot product self attention only. The rectified convolution output, $R = rect(conv_m(s)$ is then fed to the attention layer. Using the notation of Vaswani et al, we project R using 3 separate sets of weights, forming RW^Q, RW^K, RW^V. The product $A = (RW^Q)(RW^K)^T$ forms an attention matrix, which can be used to identify interactions between positions within a sequence. The output of the attention layer, $(RW^Q)(RW^K)^T(RW^V)$, a weighted sequence of words

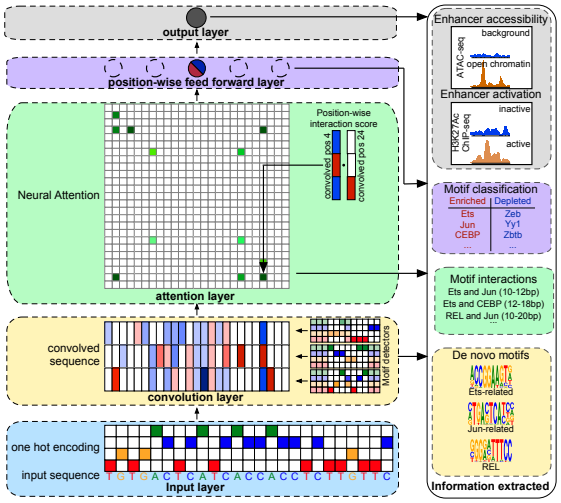

Figure 2: Overview of attentive neural network model

(higher weight indicate that a word is more important) is then fed to a single neuron that re-weights each motif at each position to indicate whether a motif is enriched or depleted. To predict whether a sequence is an enhancer or genomic sequence, the output layer simply counts the number of enriched motifs that are present within a sequence. Notably, our model does not apply any pooling or dense layers, allowing our model to learn position specific information for each individual position in a sequence.

Results

We used ChIP-seq targeting H3K27Ac and ATAC-seq , we profiled the active enhancers in mouse macrophage cells stimulated with 3 treatments - Vehicle, KLA, and IL4. These experiments identified tens of thousands of active enhancers in each treatment context. To assess the performance of our model architecture, we compared the performance of our model against the current state of the art, a convolutional network. We trained our model and an implementation of DeepBind, a previously described convolutional network (Alipanahi et al. 2015), to distinguish active enhancers from random genomic sequences. The performance of our model exceeded that of the convolutional model, in terms of model accuracy and precision (Figure 1). To ensure that the improvement in the performance of our model is not due to the greater number of free parameters (Figure 1), we also trained a large convolutional network (with 54 convolution kernels and 108 dense neurons versus 16 convolution kernels and 32 neurons in the original model). The improved performance of our model suggests that the attention mechanism is capable of extracting additional useful information.

Future Work

While we are encouraged by the performance of our model, we believe the insights we can extract from the network are of greater importance. Using graph clustering approaches,

			Att	Conv	LargeConv
		Params	10753	2162	10850
Tx	Veh	Acc.	0.854	0.822	0.846
		Prec.	0.838	0.804	0.830
	KLA	Acc.	0.859	0.807	0.839
		Prec.	0.857	0.791	0.826
	IL4	Acc.	0.862	0.832	0.847
		Prec.	0.858	0.809	0.836

Table 1: Performance metrics (n=3), accuracy and precision, of our attentive neural network (Att.), a convolutional network (Conv), and a large convolutional network (LargeConv) are shown for 3 treatment conditions (Veh, KLA, IL4)

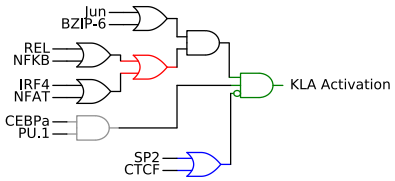

Figure 3: Circuit diagram of TFs that need to be present at an enhancer that activates in cells treated with KLA

we can distill the parameters learned by the model into logical circuits detailing the TF motifs that need to be present in an enhancer that responds to a specific stimulus such as KLA (Figure 3). Using this information, we will be able to design artificial transcriptional units which are cell type and signal specific. This would allow us to use well characterized, general delivery systems such as adenovirus, which target cells indiscriminately. We can specify treatment activity by using a synthetic target specific transcriptional unit designed using information from our ANN that drives the expression of a therapeutic transgene. As an example, a transcriptional unit specific to cancer driving the expression of a immune reactive transgene would allow for immune specific targeting of the cancer.

References

Alipanahi, B.; Delong, A.; Weirauch, M. T.; and Frey, B. J. 2015. Predicting the sequence specificities of DNA- and RNA-binding proteins by deep learning. *Nat Biotechnol* 33(8):831–838.

Cheng, J.; Dong, L.; and Lapata, M. 2016. Long Short-Term Memory-Networks for Machine Reading. *arXiv*.

Consortium, T. E. P. 2012. An integrated encyclopedia of DNA elements in the human genome. *Nature* 489(7414):57–74.

Farley, E. K.; Olson, K. M.; Zhang, W.; Rokhsar, D. S.; and Levine, M. S. 2016. Syntax compensates for poor binding sites to encode tissue specificity of developmental enhancers. *Proceedings of the National Academy of Sciences* 113(23):6508–6513.

Heinz, S.; Romanoski, C. E.; Benner, C.; Allison, K. A.; Kaikkonen, M. U.; Orozco, L. D.; and Glass, C. K. 2013. Effect of natural genetic variation on enhancer selection and function. *Nature* 503(7477):487–92.

Kelley, D. R.; Snoek, J.; and Rinn, J. L. 2016. Basset: learning the regulatory code of the accessible genome with deep convolutional neural networks. *Genome research* 26(7):990–9.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention Is All You Need. *Nips* 40:e253.