RNA Alternative Splicing Prediction with Discrete Compositional Energy Network

Alvin Chan*
guoweial001@e.ntu.edu.sg
Nanyang Technological University
Singapore

Anna Korsakova*
kors0001@e.ntu.edu.sg
Nanyang Technological University
Singapore

Yew-Soon Ong
asyong@ntu.edu.sg
Nanyang Technological University
Singapore

Fernaldo Richtia Winnerdy
fernando.winnerdy@ntu.edu.sg
Nanyang Technological University
Singapore

Kah Wai Lim
kwlim@ntu.edu.sg
Nanyang Technological University
Singapore

Anh Tuan Phan
phantuan@ntu.edu.sg
Nanyang Technological University
Singapore

ABSTRACT

A single gene can encode for different protein versions through a process called alternative splicing. Since proteins play major roles in cellular functions, aberrant splicing profiles can result in a variety of diseases, including cancers. Alternative splicing is determined by the gene's primary sequence and other regulatory factors such as RNA-binding protein levels. With these as input, we formulate the prediction of RNA splicing as a regression task and build a new training dataset (CAPD) to benchmark learned models. We propose discrete compositional energy network (DCEN) which leverages the hierarchical relationships between splice sites, junctions and transcripts to approach this task. In the case of alternative splicing prediction, DCEN models mRNA transcript probabilities through its constituent splice junctions' energy values. These transcript probabilities are subsequently mapped to relative abundance values of key nucleotides and trained with ground-truth experimental measurements. Through our experiments on CAPD¹, we show that DCEN outperforms baselines and ablation variants.²

CCS CONCEPTS

- **Applied computing** → **Bioinformatics; Health informatics;**
- **Computing methodologies** → **Neural networks.**

KEYWORDS

energy-based models, machine learning, deep neural networks, splicing prediction

ACM Reference Format:

Alvin Chan, Anna Korsakova, Yew-Soon Ong, Fernaldo Richtia Winnerdy, Kah Wai Lim, and Anh Tuan Phan. 2021. RNA Alternative Splicing Prediction with Discrete Compositional Energy Network. In *ACM Conference on Health, Inference, and Learning (ACM CHIL '21)*, April 8–10, 2021, Virtual Event, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3450439.3451857>

*Both authors contributed equally to this research.

¹Dataset available at: <https://doi.org/10.21979/N9/FFN0XH>

²Codes and models are released at: https://github.com/alvinchangw/DCEN_CHIL2021

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ACM CHIL '21, April 8–10, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8359-2/21/04.

<https://doi.org/10.1145/3450439.3451857>

Inference, and Learning (ACM CHIL '21), April 8–10, 2021, Virtual Event, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3450439.3451857>

1 INTRODUCTION

RNA plays a key role in the human body and other organisms as a precursor of proteins. RNA alternative splicing (AS) is a process where a single gene may encode for more than one protein isoforms (or mRNA transcripts) by removing selected regions in the initial pre-mRNA sequence. In the human genome, up to 94% of genes undergo alternative splicing [36]. AS not only serves as a regulatory mechanism for controlling levels of protein isoforms suitable for different tissue types but is also responsible for many biological states involved in disease [35], cell development and differentiation [14]. While advances in RNA-sequencing technologies [5] have made quantification of AS in patients' tissues more accessible, an AS prediction model will alleviate the burden from experimental RNA profiling and open doors for more scalable *in-silico* studies of alternatively spliced genes.

Previous studies on AS prediction mostly either approach it as a classification task or study a subset of AS scenarios. Training models that can predict strengths of AS in a continuous range and are applicable for all AS cases across the genome would allow wider applications in studying AS and factors affecting this important biological mechanism. To this end, we propose AS prediction as a regression task and curate Context Augmented Psi Dataset (CAPD) to benchmark learned models. CAPD is constructed using high-quality transcript counts from the ARCHS4 database [21]. The data in CAPD encompass genes from all 23 pairs of human chromosomes and include 14 tissue types. In this regression task, given inputs such as the gene sequence and an array of tissue-wide RNA regulatory factors, a model would predict, for key positions on the gene sequence, these positions' relative abundance (ψ) in the mRNAs found in a patient's tissue.

Each mRNA transcript may contain one or more splice junctions, locations where splicing occurred in the gene sequence. We hypothesize that a model design that considers these splice junctions would be key for good AS prediction. More specifically, these splice junctions are produced through a series of molecular processes during splicing, so modeling the energy involved at each junction may offer an avenue to model the whole AS process. This is the intuition behind our proposed discrete compositional energy network (DCEN) which models the energy of each splice junction,

composes candidate mRNA transcripts' energy values through their constituent splice junctions and predicts the transcript probabilities. The final component maps transcript probabilities to each known exon start and end's ψ values. DCEN is trained end-to-end in our experiments with ground-truth ψ labels. Through our experiments on CAPD, DCEN outperforms other baselines that lack the hierarchical design to model relationships between splice site, junctions and transcript. To the best of our knowledge, DCEN is the first approach to model the AS process through this hierarchical design. While DCEN is evaluated on AS prediction here, it can potentially be used for other applications where the compositionality of objects (e.g., splice junctions/transcripts) applies. All in all, the prime contributions of our paper are as follows:

- We construct Context Augmented Psi Dataset (CAPD) to serve as a benchmark for machine learning models in alternative splicing (AS) prediction.
- To predict AS outcomes, we propose discrete compositional energy network (DCEN) to output ψ by modeling transcript probabilities and energy levels through their constituent splice junctions.
- Through experiments on CAPD, we show that DCEN outperforms baselines and ablation variants in AS outcome prediction, generalizing to genes from withheld chromosomes and of much larger lengths.

2 BACKGROUND: RNA ALTERNATIVE SPLICING

RNA alternative splicing (AS) is a process where a single gene (DNA / pre-mRNA) can produce multiple mRNAs, and consequently proteins, increasing the biodiversity of proteins encoded by the human genome. Pre-mRNAs contain two kinds of nucleotide segments, introns and exons. Each post-splicing mRNA transcript would only have a subset of exons while the introns and remaining exons are removed. A molecular machine called spliceosome joins the upstream exon's end with the downstream exon's start nucleotides to form a splice junction and removes the intronic segment between these two sites. In an example of an exon-skipping AS event in Figure 1, a single pre-mRNA molecule can be spliced into more than one possible mRNA transcripts (T_α, T_β) with different probabilities ($P_{T_\alpha}, P_{T_\beta} \in [0, 1]$). These probabilities are largely determined by local features surrounding the splice sites (exon starts/ends) such as the presence of key motifs on the exonic and intronic regions surrounding the splice sites nucleotide. Global contextual regulatory factors such as RNA-binding proteins and small molecular signals [3, 34, 37] can also influence the transcript probabilities, creating variability for AS outcomes in cells from different tissue types or patients. While exon-skipping is the most common form of AS, there are others such as alternative exon start/end positions and intron retention.

2.1 Measurement of Alternative Splicing Outcome

One standard way to quantify AS outcome for a group of cells is through **percent spliced-in** ($\psi \in [0, 1]$). Essentially, ψ is defined as the ratio of relative abundance of an exon over all mRNA products within a single gene. An exon with $\psi = 1$ means that it is included

in all mRNAs found from experimental RNA-sequencing measurements while $\psi = 0$ means the exon is missing in that particular gene. ψ can also be annotated onto exon's key positions such as its start and end locations. This allows one to approach the AS prediction as a regression task of predicting ψ for each nucleotide of interest.

Figure 1: Mechanism of alternative splicing and its relationship with ψ annotations. Introns in the gene sequences are colored gray while exons (e_1, e_2, e_3) are colored otherwise. Each patient sample has a unique set of regulatory factors, leading to a variety of alternative splicing outcomes in the population. In the training of DCEN, the DNA sequences of patients are assumed to be the same.

3 RELATED WORK

We review prior art on RNA splicing prediction and energy-based models, highlighting those most similar to our work.

3.1 Splice Site Classification

The earliest task of machine learning on RNA splicing involves classification of splicing sites such as exon start and end positions in a given gene sequence, first using models such as decision trees [28] and support vector machines [9]. As deep learning gains wider adoption, a line of works uses neural networks for splice site prediction from raw sequence [19, 24, 43, 45]. In a recent example, [19] used a 1-D Resnet model to classify individual nucleotides in a pre-mRNA sequence into 3 categories: 1) exon's start, 2) exon's end or 3) none of the two classes. Unlike these models that only classify splice sites, we propose DCEN to predict ψ levels of splice sites which involve the consideration of patient-specific input such as levels of RNA regulatory factors on top of just primary gene sequences.

3.2 Alternative Splicing Prediction

The prior work in alternative splicing prediction can be categorized into two distinct groups. The first group framed the prediction as

a classification task, whether an alternative splicing event would occur given input or change in input. The earliest examples involved using a Bayesian regression [1] and Bayesian neural network [40] to predict whether an exon would be skipped or included in a transcript. [23] used a neural network with dense layers to predict the type of AS event. Using a classification framework [39] to predict one of three classes (high/medium/low), the relative ψ value can also be inferred. Another deep learning-based approach [24] utilized a CNN-based framework to classify between four AS event classes (exon skipping, alternative 3', alternative 5' or constitutive exon).

The second group, which includes DCEN, addresses the prediction as a regression rather than a classification task. This formulation gives higher resolution in the AS event since predicted values correlate with the strength of the AS outcome. [4] proposed a deep learning model to predict which site is most likely to be spliced given the raw sequence input of 80-nt around the site. The neural networks in [6, 7] also use the primary sequences around candidate splice sites as inputs to infer their ψ values. These approaches predict ψ values of splice sites using only their individual representations without modeling the relationship between these splice sites and their parent transcripts, and are conceptually similar to the 2000-nt SpliceAI baseline here which is outperformed by DCEN in our experiments.

Since cellular signals such as RNA-binding proteins (RBPs) are observed to affect RNA splicing [37, 41], models such as [20] predicts ψ values with RBP information and generated genomic features used as input. while [18, 44] have emerged to incorporate both primary sequence features and RBP levels to better predict exon inclusion levels given a small number of experimental read counts. While also considering regulatory factors such as RBPs, our approach differs from [37, 41] as we do not assume the availability of experimental read counts for the gene of interest. To the best of our knowledge, DCEN is the first approach to model whole transcript constructs (through energy levels) on top of the immediate neighborhood around the nucleotide of interest when predicting its ψ value in the splicing process.

3.3 Energy-Based Models

Most recent work in energy-based models (EBM) [22] focused on the application of image generative modeling. Neural networks were trained to assign low energy to real samples [12, 15, 26, 31, 38] so that realistic-looking samples can be sampled from the low-energy regions of the EBM's energy landscape. Instead of synthesizing new samples, our goal here is to predict RNA splicing outcomes. Other applications of EBMs include anomaly detection [31], protein conformation prediction [11] and reinforcement learning [16]. Previous compositional EBMs such as [10, 16] considered high dimensional continuous spaces in their applications which makes sampling from the model intractable. In contrast, since genes consist of a finite number of known transcripts, DCEN considers discrete space where the probabilities of the transcripts are tractable through importance sampling with a uniform distribution.

4 CAPD: CONTEXT AUGMENTED PSI DATASET

The core aim of Context Augmented Psi Dataset (CAPD) is to construct matching pairs of sample-specific inputs and labels to frame the alternative splicing prediction as a regression task and facilitate future benchmarking of machine learning splicing prediction models. Each CAPD sample is a unique AS profile of a gene from the cells of a particular tissue type, from an individual patient. Its annotations contain $\psi \in [0, 1]$ of all the known exon starts and ends for the particular gene. Apart from the ψ labels, each data sample also contains the following as inputs: a) full sequence of the gene ($\mathbf{x}$), b) nucleotide positions of all the known transcripts ($\mathcal{T}$) on the full gene sequence and c) levels of RNA-regulatory factors ($\mathbf{x}_{\text{reg}}$).

4.1 Construction of CAPD

We mine transcript abundance data from the publicly available ARCHS4 database v.8 [21]. ARCHS4 database v.8 contains expression data for 238522 publicly available human RNA-seq samples that were retrieved from Gene Expression Omnibus (GEO) and aligned to human transcriptome Ensembl 90 cDNA [42] to produce count numbers for each transcript in each sample. A simple keyword search was used to find $250 \times N_T$ samples for each of the 16 tissue types ($N_T = 16$) selected for training: adipose tissue, blood, brain, breast, colon, heart, kidney, liver, lung, lymph node, prostate gland, skeletal muscle tissue, testes, thyroid gland. Underrepresented tissues, i.e. with the number of samples below a threshold ($N = 100$) were not included (see the pipeline in Figure 2A). Each sample row contains transcript-centric raw RNA counts. Standard normalization of RNA read counts to RPKM [25] was then performed. To exclude samples with significantly different expression patterns, a z-score outlier removal procedure described in [27] for similar tasks was applied to samples from each tissue separately. Expression patterns of samples belonging to one tissue type are varying, and $\approx 2 - 5\%$ of the sample population is usually considered outlying by the algorithm. To homogenize the data, batch effects were removed with *edgeR* R library [29] (Figure 2B). Heterogeneity of the data, however, is still expected as the samples belong to different sources and were chosen randomly. This gives the normalized transcript count values (c_{T_α}) for each gene transcript (T_α).

To construct the levels of RNA splicing regulatory factors ($\mathbf{x}_{\text{reg}}$), we extract the expression levels of RBPs corresponding to the RBPDB database [8] and RNA chemically modifying corresponding to [2] in a sample-wise manner from the normalized transcript count matrix (Figure 2C). This gives a 3971-dimensional $\mathbf{x}_{\text{reg}}$. These, together with universal pre-mRNA transcript sequences (same for each sample), compose model input.

To generate primary pre-mRNA transcript sequences, we follow the procedure described in [19]: pre-mRNA (synonymous to DNA, with $T \rightarrow U$) sequences are extracted with flanking ends of 1000 nt on each side, while intergenic sequences are discarded. Pseudogenes, genes with sequence assembly gaps and genes with paralogs are excluded from the data. Exon coordinate information is retrieved from GENCODE Release 26 for GRCh38 [13] comprehensive set, downloaded from the UCSC table browser³. We omit

³<https://genome.ucsc.edu/cgi-bin/hgTables>

Figure 2: Pipeline of the CAPD dataset preparation with the steps as follows: A) a simple keyword search in the ARCHS4 v.8 database is performed to find a given number of samples for the specified tissues. Tissues for which the number of samples found is less than a specified threshold ($N=100$) are not included. B) Transcript libraries are normalized to RPKM, processing by *edgeR* library and outlier removal are performed. C) Normalized transcript counts are then used to construct training labels y and auxiliary regulatory vector inputs x_{reg} for each sample. The latter is then used together with pre-mRNA transcript sequence as model input. D) After model training, a sample input for ψ prediction consists of pre-mRNA transcript sequence and a vector of abundance of sample-specific regulatory factors.

genes with missing matching GENCODE ID, resulting in a total of 19399 unique human gene sequences. The coordinates of the gene pre-mRNA sequence start and end are determined by the left- or right-most position among all the transcripts for that gene, further extended with flanking ends of 1000 nt for context on each side. The CAPD dataset is split into train and test according to the chromosome number and length of the genes. All genes in chromosome 1, 3, 5, 7, and 9 are withheld as test samples (Test-Chr), similar to [19]. To test that models can generalize to genes of longer lengths, we further withhold all genes from the other chromosomes that have >100K nt (Test-Long) and group them in the test set. The remaining genes are used for training. Key statistics of the CAPD is summarized in Table 1.

4.2 Label Annotation

The ψ labels are constructed as follows (Figure 2C): 1) the count values c_i for all known exon start and end (position i) are initialized to zero. 2) enumerating through all transcripts, each transcript count is added onto the counts of its constituent exons' start/end $c_i \leftarrow c_i + c_{T_m}, \forall i \in T_m$. 3) To compute ψ values, each count value

is divided by the sum of transcript counts to normalize its value to $[0, 1]$, i.e., $\psi_i = c_i / \sum_m c_{T_m}$.

Table 1: CAPD data statistics.

	Train	Test-Chr	Test-Long
# of unique genes	11,472	5,604	2,323
mean pre-mRNA length (nt)	26,196	73,355	247,041
mean # of exons	6.7	7.0	10.5

5 DCEN: DISCRETE COMPOSITIONAL ENERGY NETWORK

In § 2, we learn that final mRNA splice isoforms may comprise one or more splice junctions, points where upstream exon's end and downstream exon's start meet. Inspired by the creation of splice junctions by spliceosomes, the first stage of DCEN models the energy values of the splicing process at splice junctions. As the splice junctions are typically far enough (~300 nt) and assumed to be independent of one another, the energy of a final mRNA transcript

can be composed by the summation of its constituent splice junctions' energy. The second stage of DCEN derives probabilities for the formation of each transcript from their energy values. These transcript probabilities are then mapped to the relative abundance (ψ) of exon starts/ends at its corresponding splice sites. Since a particular splice junction may appear in more than one mRNA isoform, we design DCEN to be invariant to the splice junctions. In the following § 5.1, we discuss the key components of DCEN while § 5.5 details its training process.

5.1 Model Architecture

Here, we detail the DCEN learned energy functions and how transcript probabilities can be derived from their energy levels through Boltzmann distribution and importance sampling. A summary of DCEN model architecture is shown in Figure 3.

Figure 3: Model architecture of proposed discrete compositional energy network (DCEN). A) Learned representations of the DNA sequence are extracted by f_{nt} while the representations of the regulatory factors by f_{reg} . B) The representation of a particular splice junction is the concatenation of acceptor and donor sites' representations. The learned energy value of a particular splice junction is computed by passing its representation through f_E . The energy level of each transcript is the sum of its junctions' energies. C) Through a softmax operation, the probability of each transcript in a gene is computed. By pooling the probabilities of all transcripts that contain a particular nucleotide, we get the nucleotide's predicted ψ value.

5.2 Learned Energy Functions

The weights of DCEN's learned energy functions consist of a 1) feature extractor, 2) regulatory factors encoder and 3) junction energy network. The feature extractor (f_{nt}) takes the pre-mRNA sequence of length l as its input ($x \in \mathbb{R}^{l \times 4}$) where 4 is the number of possible nucleotides and outputs a hidden representation ($h_{nt} \in \mathbb{R}^{l \times d}$) for each nucleotide position (x_i) while the regulatory factors encoder (f_{reg}) takes in the levels of regulatory factors (x_{reg}) to compute a gene-wide hidden states $h_{reg} \in \mathbb{R}^d$:

$$h_{nt} = f_{nt}(x), \quad h_{reg} = f_{reg}(x_{reg}) \quad (1)$$

We concatenate h_{reg} to all the position-specific h_{nt} to form a new position-wise hidden state (h_i) that is dependent on regulatory factors. The representation of a particular splice junction ($J_k = (i, j)$) is the concatenation between the hidden states of upstream exon end's and downstream exon start's hidden states:

$$h_{J_k} = [h_i; h_j], \quad h_i = [h_{nt_i}; h_{reg}] \quad (2)$$

If an exon start/end is the first/last nucleotide of a transcript, its hidden state is concatenated with a learned start/end token instead (h_{start} or h_{end} respectively). To model the energy ($E_{J_k} \in \mathbb{R}$) of producing splice junction J_k , we feed its representation into the energy network (f_E). We sum up the energy values of all splice junctions (J_k) inside a mRNA transcript (T_α) to compose the total energy ($E_{T_\alpha} \in \mathbb{R}$) involved in producing the transcript from a splicing event:

$$E_{T_\alpha} = \sum_k E_{J_k}, \quad \forall J_k \in T_\alpha, \quad E_{J_k} = f_E(h_{J_k}) \quad (3)$$

5.3 Transcript Probabilities from Energy Values

After obtaining the energy levels (E_{T_α}) of all mRNA transcript candidates for a particular gene, we can compute the probabilities of these transcripts via a softmax operation through the theorem below.

THEOREM 5.1. *Given the energy levels of all the possible discrete states of a system, the probability of a particular state T_i is the softmax output of its energy E_{T_i} with respect to those of all other possible states in the system, i.e.,*

$$P_i = \frac{\exp(-E_{T_i})}{\sum_j \exp(-E_{T_j})} = \text{Softmax}_i(E) \quad (4)$$

Its proof, deferred to the Appendix § A.2, can be derived through Boltzmann distribution and importance sampling. Since each mRNA transcript T_α can be interpreted as a discrete state of the alternative splicing event for a particular gene (system) in Theorem 5.1, we can compute its probability from its energy value.

It is important to also consider a null state with energy E_{null} where none of the gene's mRNA transcripts is produced. In DCEN, E_{null} is a learned parameter. In summary, the probability of producing transcript T_α in an gene splicing event is

$$P_{T_\alpha} = \text{Softmax}_\alpha(E), \quad E = [E_{T_\alpha}, E_{T_\beta}, \dots, E_{null}] \quad (5)$$

If the null state were not considered, the model would incorrectly assume that a particular gene is always transcribed since the sum of transcripts' probabilities in a gene would be $\sum_i P_{T_i} = 1$.

5.4 Exon Start/End Inclusion Levels from Transcript Probabilities

We can compute the probability of a particular (exon start/end) nucleotide of position i by summing up the probabilities of all transcripts that contain that nucleotide.

$$P_i = \sum_m P_{T_m}, \forall T_m \in \mathcal{T}_i \quad (6)$$

where $\mathcal{T}_i$ is the set of transcripts containing the nucleotide of interest. By inferring the (exon start/end) nucleotide inclusion levels through transcript probabilities, our model has the advantage of additional access to the relative transcription level of the gene's transcripts over baselines that infer directly at the nucleotide positions (such as baselines in § 6.1).

5.5 Training Algorithm

5.5.1 Regression Loss. Since experimental ψ levels from CAPD are essentially the empirical observations of nucleotide present in the final mRNA transcript products, its normalized values ($y_i \in [0, 1]$) can be used as the ground-truth label for the predicted nucleotide inclusion levels. This allows us to train DCEN as a regression task by minimizing the mean squared error (MSE) between the predicted nucleotide ψ and the normalized experimental ψ values:

$$L_\psi = \sum_i \|\tilde{y}_i - y_i\|_2^2, \tilde{y}_i = P_i \quad (7)$$

In our experiments, only nucleotide positions that are either an exon start or end are involved in this regression training objective.

5.5.2 Classification Loss. We also include a classification objective, similar to [19], where a classification head f_{cls} takes the nucleotide hidden states ($\mathbf{h}_{nt}$) as input to predict probability of every nucleotide as one of the 3 classes (exon start, end and neither) to give the classification loss:

$$L_{cls} = -\mathbf{y}_{cls}^\top \log f_{cls}(\mathbf{h}_{nt}) \quad (8)$$

where $\mathbf{y}_{cls}$ is the ground-truth labels for each nucleotide. This helps the DCEN learn features on the gene primary sequence that are important for RNA splicing. A summary of the training phase is shown in Algorithm 1.

6 EXPERIMENTS

We evaluate DCEN and baselines on the prediction of the ψ values of exon starts and ends in our new CAPD dataset. In the following, we describe the baseline models and DCEN ablation variants before discussing the experimental setup and how well these models generalize to withheld test samples.

6.1 Baselines and Ablation Variants

SpliceAI is a 1D convolutional Resnet [17] trained to predict splice sites on pre-mRNA sequences. We train three variants of SpliceAI

Algorithm 1: Discrete Compositional Energy Network Training

```

1 Input: Training data  $\mathcal{D}_{\text{train}}$ , Learning rate  $\gamma$ ,
2 for each training iteration do
3   Sample  $(\mathbf{x}, \mathbf{x}_{\text{reg}}, \mathbf{y}, \mathcal{T}, \mathcal{J}) \sim \mathcal{D}_{\text{train}}$ 
4    $\mathbf{h}_{nt} \leftarrow f_{nt}(\mathbf{x})$ ,
5    $L_{cls} \leftarrow -\mathbf{y}_{cls}^\top \log f_{cls}(\mathbf{h}_{nt})$   $\triangleright$  Compute exon start/end
      classification cross-entropy loss
6    $\mathbf{h}_{reg} \leftarrow f_{reg}(\mathbf{x}_{reg})$ 
7    $\mathbf{h}_i \leftarrow [\mathbf{h}_{nt}; \mathbf{h}_{reg}]$ 
8    $\mathbf{h}_{J_k} \leftarrow [\mathbf{h}_i; \mathbf{h}_j]$ ,  $J_k = (i, j)$   $\triangleright$  Get junction state from
      upstream exon end  $i$  and downstream exon start  $j$ 
9    $E_{J_k} \leftarrow f_E(\mathbf{h}_{J_k})$   $\triangleright$  Compute splice junction energy
10   $E_{T_\alpha} \leftarrow \sum_k E_{J_k}$ ,  $\forall J_k \in \mathcal{T}_\alpha$ ,  $\forall T_\alpha \in \mathcal{T}$   $\triangleright$  Compose
      transcript energy
11   $P_{T_\alpha} \leftarrow \text{Softmax}_\alpha(E)$ ,  $E = [E_{T_\alpha}, E_{T_\beta}, \dots]$   $\triangleright$  Compute
      transcript probabilities
12   $\tilde{y}_i \leftarrow \sum_m P_{T_m}$ ,  $\forall T_m \in \mathcal{T}_i$ 
13   $L_\psi \leftarrow \sum_i \|\tilde{y}_i - y_i\|_2^2$   $\triangleright$  Compute exon start/end
      inclusion regression loss
14   $\theta \leftarrow \theta + \gamma \nabla_\theta (\lambda_{reg} L_\psi + \lambda_{cls} L_{cls})$ 
```

to compare as baselines in our experiments: the first (**SpliceAI-cls**, Figure 4a) is trained only on the classification objective, similar to the original paper, to predict whether a nucleotide is an exon start, end or neither of them. The second (**SpliceAI-reg**, Figure 4b) is trained only on a regression objective like DCEN to directly predict the ψ levels of nucleotides while the third variant (**SpliceAI-cls+reg**) is trained on both the classification and regression objectives. Both SpliceAI-reg and SpliceAI-cls+reg also have a regulatory factors encoder (f_{reg}) similar to DCEN's to compute $\mathbf{h}_i$ and a regression head (f_ψ) to output ψ level for each nucleotide. For a direct comparison, DCEN's feature extractor f_{nt} takes the same architecture as the SpliceAI Resnet. Two DCEN ablation variants are also evaluated: The **Junction-psi** model (Figure 4c) predicts the psi levels of a particular splice junction directly rather than its energy level in the case of DCEN. A simpler ablation variant (**SpliceAI-ML** or SpliceAI-match layers, Figure 4b) substitutes DCEN's f_E with a position-wise feedforward MLP containing the same number of parameters to verify that DCEN's better performance is not due to more learned parameters.

6.2 Data & Models

We use the CAPD dataset (§ 4) for the training and evaluation of all models. 10% of the CAPD training genes are randomly selected as the validation set for early-stopping while the rest are used as training samples for the models. For DCEN's f_{nt} and the SpliceAI baselines, we follow the same setup as the SpliceAI-2K model in [19] which is a Resnet made up of 1-D convolutional layers with a perceptive window of 2K nucleotides, 1K on each flanking sides. The SpliceAI Resnet model has a total of 12 residual units and hidden states of size 32. The number of channels in $\mathbf{h}_{reg}$ and $\mathbf{h}_{nt}$ are 32 while $\mathbf{h}$ has a size of 64 channels. We use a 3-layer MLP for f_{reg} . The regression head f_ψ in SpliceAI-reg and SpliceAI-cls+reg is

a 3-layer MLP with a sigmoid activation to outputs a scalar ψ value. DCEN's f_E is a 4-layer MLP and outputs a scalar energy value for each splice junction. Intermediate hidden states of f_{reg} and f_E all have dimension of 32.

6.3 Training & Evaluation

All models are trained with Adam optimizer with a learning rate of 0.001 in our experiments. Due to the data's large size, the training is early-stopped when the model's validation performances plateau: less than 25% of the full train dataset for all models in our experiments. For the training of SpliceAI models, samples are fed into the model with a batch size of 8 sequences with a maximum length of 7K nucleotides (5K labeled + 2K flanking). For training of DCEN and its ablation variants, a SpliceAI-cl® model pretrained on CAPD was used as the weights of f_{reg} , f_{nt} and only the parameters of f_E is trained to reduce training time. In each training iteration of DCEN and its ablation variants, a batch of 16 genes was used to train the weights. We evaluate all the models on withheld test samples with two standard regression metrics: Spearman rank correlation and Pearson correlation. Pearson correlation measures the linear relationship between the ground-truth and predicted exon start/end inclusion levels while Spearman rank correlation is based on the ranked order of the prediction and ground-truth values.

6.4 Results

The DCEN outperforms all baselines and ablation variants for both regression metrics when evaluated on the withheld test samples, as shown in Table 2. Even with the same number of parameters, the DCEN shows better performance than the SpliceAI-ML and Junction-psi model. Even with more trained parameters, Junction-psi model performs worse than the SpliceAI-reg and SpliceAI-cl® baselines while the SpliceAI-ML does not show clear improvement over these two baselines.

A key difference between these baselines and DCEN lies in DCEN's inductive bias that models the hierarchical relationships between splice sites and junctions, as well as between splice junctions and their parent transcripts. Even with a matching number of parameters as DCEN, SpliceAI-ML still underperforms DCEN indicating DCEN's model size is not the key contributor to its performance. These observations indicate that DCEN's design to compose transcripts' energy through splice junctions and infer their probabilities from energy values is key for better prediction. The correlation results reported in Table 2, 3 and 5 all have p-values of zero in working precision due to the large number of gene samples.

Training DCEN for 5 times the training steps at the early stopped point does not result in better (Table 6 in S A.1) nor much worse performance on the test samples. This suggests that DCEN is underfitting the training dataset and may benefit from a larger model size.

6.4.1 Excluded chromosomes. When evaluated separately test samples from chromosomes (1, 3, 5, 7, and 9) not seen during the training phase, DCEN maintains its superior performance (Table 3) when compared to the ablation baselines, showing that it generalizes across novel gene sequences.

Table 2: Performance of DCEN and baselines on all withheld test samples.

Model	Spearman Cor.	Pearson Cor.
SpliceAI-cl®	0.399	0.349
SpliceAI-reg	0.579	0.574
SpliceAI-cl®	0.577	0.571
SpliceAI-ML	0.572	0.575
Junction-psi	0.559	0.510
DCEN (ours)	0.623	0.651

6.4.2 Long gene sequences. DCEN was trained only on genes sequences of length less than 100K nucleotides. From Table 5, we observe that DCEN still outperforms the other baselines by a substantial margin and retains most of its performance on these samples when evaluated on genes with long sequences (>100K nucleotides). Compared to shorter introns in genes of shorter length, splicing of pre-mRNA with very large introns was observed to occur in a more nested and sequential manner [30, 33]. Since genes with long sequences contain more large introns, this difference in the splicing mechanism may explain the slight drop in DCEN's performance on genes with longer sequences.

Table 3: Performance of DCEN and baselines on Test-Chr, test samples from chromosomes (1, 3, 5, 7, and 9) different from the training set.

Model	Spearman Cor.	Pearson Cor.
SpliceAI-cl®	0.400	0.353
SpliceAI-reg	0.579	0.586
SpliceAI-cl®	0.577	0.583
SpliceAI-ML	0.570	0.587
Junction-psi	0.560	0.525
DCEN (ours)	0.622	0.665

6.4.3 Performance differs across tissue types. We observe that the performance of DCEN's prediction differs across test samples of different tissues types (Table 4). The best performing tissue (heart) has Spearman rank correlation of 0.694 and Pearson correlation of 0.716 while the lowest (testes) achieves 0.415 and 0.423 respectively.

6.4.4 Integration of regulatory features allows for sample-specific predictions. An example of a set of model predictions in comparison with the ground truth values for BCL2 gene acceptor sites is shown in Figure 5. Box plots of experimental values (left of Figure 5) represent distributions of ψ in 50 samples for three types of tissues: blood, skeletal muscle and lung. The predictions of DCEN vs SpliceAI-cl® for these acceptor sites are shown on the right. We observe that DCEN's predictions can partially recover the variances in ψ values unique to tissue types (e.g., acceptor site 158259 in Figure 5). When comparing the variance of ψ prediction within each tissue type with the variance of ground-truth ψ , the Pearson correlation is 0.177 while the Spearman rank correlation is 0.198. Figure 6 in § A.1 shows the comparison for BCL2 donor sites.

Figure 4: Ablation models used to compare with our proposed DCEN model. (a) SpliceAI-cls [19] is trained to predict whether a particular nucleotide in the DNA sequence is an acceptor, donor site or neither of them from only the primary DNA sequence, without input of regulatory factors (x_{reg}). The predicted probability of the acceptor/donor class is assumed to be predicted ψ values during the comparison. (b) SpliceAI-reg and SpliceAI-ML are both trained to predict the ψ value of a particular nucleotide with a regression objective. These two ablation models, unlike DCEN (Figure 3), infer ψ values from individual nucleotide's representation and does not rely on the compositionality of transcripts and their splice junctions. (c) Junction-psi infers the ψ value of a particular splice junction from the joint representations of its acceptor and donor sites, without modeling the relationship between junctions and their parent transcripts.

Figure 5: Comparison of the BCL2 gene acceptor sites' ψ , experimental vs predicted. Left: box plots of experimental values (grounds truths) of BCL2 acceptor sites in three different tissues: blood, skeletal muscle and lung, denoted as BLD, SKM and LNG respectively. Right: box plots of predicted ψ for the same group of samples using DCEN; SpliceAI-cls prediction (black triangles) shown for comparison. Please note that SpliceAI-cls prediction uses sequence information only and the prediction will be the same for every sample. Acceptor site coordinates (y-axis) are counted relative to the genomic coordinate of the first exon start of BCL2. Data from 50 samples from each tissue are shown.

Table 4: DCEN tissue-specific performance.

Tissue	Spearman Cor.	Pearson Cor.
Testes	0.415	0.423
Brain	0.497	0.582
Lung	0.585	0.624
Lymph	0.586	0.621
Liver	0.596	0.619
Blood	0.598	0.619
Colon	0.616	0.645
Breast	0.626	0.656
Adipose	0.653	0.674
Muscle	0.648	0.678
Kidney	0.654	0.684
Prostate	0.658	0.688
Thyroid	0.674	0.691
Heart	0.694	0.716

Table 5: Performance of DCEN and baselines on Test-Long, withheld samples with long gene sequences (>100K nucleotides).

Model	Spearman Cor.	Pearson Cor.
SpliceAI-cls	0.400	0.344
SpliceAI-reg	0.585	0.564
SpliceAI-cls+reg	0.582	0.561
SpliceAI-ML	0.577	0.565
Junction-psi	0.560	0.495
DCEN (ours)	0.626	0.639

7 LIMITATIONS & FUTURE DIRECTIONS

ARCHS4 database contains data from short-read RNA-seq, and alignment of short reads to quantify transcript isoforms and their

respective exons is challenging and an active area of ongoing research [32]. Inaccurate attribution of short reads might result in misleading exon counts between transcript isoforms and skewed ψ levels, which limits the prediction usability in a clinical setting.

Another possible issue arises from the fact that samples from the ARCHS4 database are heterogeneous. We speculate that this heterogeneity contributes to the variance in DCEN's performance across the different tissue types (§ 6.4.3), since the experimental parameters might differ across different labs while collecting the data from these different tissues.

Although we performed batch effect removal with a standard *edgeR* procedure, batch effects might remain in the data. Homogeneous datasets, as Genotype-Tissue Expression (GTEx)⁴ may be used instead, but using ARCHS4 database offers advantages in 1) the diversity of tissue types (healthy and diseased) and 2) its amendable form of data to be further processed for downstream applications.

In the training and evaluation here, the predictions are inferred based on a universal DNA sequence assumed to be the same for all patients. Genomic variations, such as mutations and single nucleotide variants, often lead to aberrant splicing outcomes. One future direction would be to utilize DCEN to study the roles of such genomic variations in the alternative outcomes. Through DCEN's hierarchical approach of modeling whole transcripts' probabilities, it is possible to not only draw insights into how mutations can affect inclusion levels of individual splice sites [19] but also into the relative expression of transcripts.

8 CONCLUSIONS

We curate CAPD to benchmark learning models on alternative splicing (AS) prediction as a regression task, to facilitate future work in this key biological process. By exploiting the compositionality of discrete components, we propose DCEN to predict the AS outcome by modeling mRNA transcripts' probabilities through its constituent splice junctions' energy levels. Through our experiments on CAPD, we show that DCEN outperforms baselines and other ablation variants in predicting AS outcomes. Our work shows that deconstructing a task into a hierarchy of discrete components can improve performance in learning models. We hope that DCEN can be used in future work to study RNA regulatory factors' role in aberrant splicing events.

9 ACKNOWLEDGMENTS

This work is supported by the Data Science and Artificial Intelligence Research Center (DSAIR), the School of Computer Science and Engineering at Nanyang Technological University and the Singapore National Research Foundation Investigatorship (NRF-NRFI2017-09).

REFERENCES

- [1] Yoseph Barash, John A. Calarco, Weijun Gao, Qun Pan, Xincheng Wang, Ofer Shai, Benjamin J. Blencowe, and Brendan J. Frey. 2010. Deciphering the splicing code. *Nature* (2010). <https://doi.org/10.1038/nature09000>
- [2] Georgeta N. Basturea. 2013. Research Methods for Detection and Quantitation of RNA Modifications. *Materials and Methods* (2013). <https://doi.org/10.13070/mm.en.3.186>
- [3] Sung Ho Boo and Yoon Ki Kim. 2020. The emerging role of RNA modifications in the regulation of mRNA stability. <https://doi.org/10.1038/s12276-020-0407-z>
- [4] Hannes Bretschneider, Shreshth Gandhi, Amit G. Deshwar, Khalid Zuberi, and Brendan J. Frey. 2018. COSSMO: Predicting competitive alternative splice site selection using deep learning. In *Bioinformatics*. <https://doi.org/10.1093/bioinformatics/bty244>
- [5] Douglas W. Bryant, Henry D. Priest, and Todd C. Mockler. 2012. Detection and quantification of alternative splicing variants using RNA-seq. *Methods in Molecular Biology* (2012). https://doi.org/10.1007/978-1-61779-839-9_7
- [6] Jun Cheng, Muhammed Hasan Celik, Anshul Kundaje, and Julien Gagneur. 2020. MTSplICE predicts effects of genetic variants on tissue-specific splicing. *bioRxiv* (2020).
- [7] Jun Cheng, Thi Yen Duong Nguyen, Kamil J Cygan, Muhammed Hasan Celik, William G Fairbrother, Julien Gagneur, et al. 2019. MMSplICE: modular modeling improves the predictions of genetic variant effects on splicing. *Genome biology* 20, 1 (2019), 1–15.
- [8] Kate B. Cook, Hilal Kazan, Khalid Zuberi, Quaid Morris, and Timothy R. Hughes. 2011. RBPDB: A database of RNA-binding specificities. *Nucleic Acids Research* (2011). <https://doi.org/10.1093/nar/gkq1069>
- [9] Sven Degroev, Yvan Saeys, Bernard De Baets, Pierre Rouzé, and Yves Van de Peer. 2005. SpliceMachine: Predicting splice sites from high-dimensional local context representations. *Bioinformatics* (2005). <https://doi.org/10.1093/bioinformatics/bti166>
- [10] Yilun Du, Shuang Li, and Igor Mordatch. 2019. Compositional Visual Generation with Energy Based Models. (2019).
- [11] Yilun Du, Joshua Meier, Jerry Ma, Rob Fergus, and Alexander Rives. 2020. Energy-based models for atomic-resolution protein conformations. *arXiv preprint arXiv:2004.13167* (2020).
- [12] Yilun Du and Igor Mordatch. 2019. Implicit generation and generalization in energy-based models. *arXiv preprint arXiv:1903.08689* (2019).
- [13] Adam Frankish, Mark Diekhans, Anne Maud Ferreira, Rory Johnson, Irwin Jun-greis, Jane Loveland, Jonathan M. Mudge, Cristina Sisu, James Wright, Joel Armstrong, If Barnes, Andrew Berry, Alexandra Bignell, Silvia Carbonell Sala, Jacqueline Chrast, Fiona Cunningham, Tomás Di Domenico, Sarah Donaldson, Ian T. Fiddes, Carlos García Giron, Jose Manuel Gonzalez, Tiago Grego, Matthew Hardy, Thibaut Hourlier, Toby Hunt, Osagie G. Izuogu, Julien Lagarde, Fergal J. Martin, Laura Martinez, Shamika Mohanan, Paul Muir, Fabio C.P. Navarro, Anne Parker, Baikang Pei, Fernando Pozo, Magali Ruffier, Bianca M. Schmitt, Eloise Stapleton, Marie Marthe Suner, Irina Sycheva, Barbara Uszczyńska-Ratajczak, Jinuri Xu, Andrew Yates, Daniel Zerbino, Yan Zhang, Bronwen Aken, Jyoti S. Choudhary, Mark Gerstein, Roderic Guigó, Tim J.P. Hubbard, Manolis Kellis, Benedict Paten, Alexandre Reymond, Michael L. Tress, and Paul Flicek. 2019. GENCODE reference annotation for the human and mouse genomes. *Nucleic Acids Research* (2019). <https://doi.org/10.1093/nar/gky955>
- [14] L. M. Gallego-Paez, M. C. Bordone, A. C. Leote, N. Saraiva-Agostinho, M. Ascensão-Ferreira, and N. L. Barbosa-Morais. 2017. Alternative splicing: the pledge, the turn, and the prestige: The key role of alternative splicing in human biological systems. <https://doi.org/10.1007/s00439-017-1790-y>
- [15] Will Grathwohl, Kuan-Chieh Wang, Jörn-Henrik Jacobsen, David Duvenaud, Mohammad Norouzi, and Kevin Swersky. 2019. Your classifier is secretly an energy based model and you should treat it like one. *arXiv preprint arXiv:1912.03263* (2019).
- [16] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. 2017. Reinforcement learning with deep energy-based policies. *arXiv preprint arXiv:1702.08165* (2017).
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [18] Yuanhua Huang and Guido Sanguinetti. 2017. BRIE: Transcriptome-wide splicing quantification in single cells. *Genome Biology* (2017). <https://doi.org/10.1186/s13059-017-1248-5>
- [19] Kishore Jagannathan, Sofia Kyriazopoulou Panagiotopoulou, Jeremy F. McRae, Siavash Fazel Darbandi, David Knowles, Yang I. Li, Jack A. Kosmicki, Juan Arbelaez, Wenwu Cui, Grace B. Schwartz, Eric D. Chow, Efsthios Kanterakis, Hong Gao, Amirali Kia, Serafim Batzoglou, Stephan J. Sanders, and Kyle Kai How Farh. 2019. Predicting Splicing from Primary Sequence with Deep Learning. *Cell* (2019). <https://doi.org/10.1016/j.cell.2018.12.015>
- [20] Anupama Jha, Matthew R Gazzara, and Yoseph Barash. 2017. Integrative deep models for alternative splicing. *Bioinformatics* 33, 14 (2017), i274–i282.
- [21] Alexander Lachmann, Denis Torre, Alexandra B. Keenan, Kathleen M. Jagodnik, Hoyjin J. Lee, Lily Wang, Moshe C. Silverstein, and Avi Ma'ayan. 2018. Massive mining of publicly available RNA-seq data from human and mouse. *Nature Communications* (2018). <https://doi.org/10.1038/s41467-018-03751-6>
- [22] Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, and F Huang. 2006. A tutorial on energy-based learning. *Predicting structured data* 1, 0 (2006).
- [23] Michael K.K. Leung, Hui Yuan Xiong, Leo J. Lee, and Brendan J. Frey. 2014. Deep learning of the tissue-regulated splicing code. *Bioinformatics* (2014). <https://doi.org/10.1093/bioinformatics/btu277>

⁴<https://gtexportal.org>

- [24] Zakaria Louadi, Mhaned Oubounyt, Hilal Tayara, and Kil To Chong. 2019. Deep splicing code: Classifying alternative splicing events using deep learning. *Genes* (2019). <https://doi.org/10.3390/genes10080587>
- [25] Ali Mortazavi, Brian A. Williams, Kenneth McCue, Lorian Schaeffer, and Barbara Wold. 2008. Mapping and quantifying mammalian transcriptomes by RNA-Seq. *Nature Methods* (2008). <https://doi.org/10.1038/nmeth.1226>
- [26] Erik Nijkamp, Mitch Hill, Song-Chun Zhu, and Ying Nian Wu. 2019. Learning non-convergent non-persistent short-run MCMC toward energy-based model. In *Advances in Neural Information Processing Systems*. 5232–5242.
- [27] Michael C. Oldham, Genevieve Konopka, Kazuya Iwamoto, Peter Langfelder, Tadafumi Kato, Steve Horvath, and Daniel H. Geschwind. 2008. Functional organization of the transcriptome in human brain. *Nature Neuroscience* (2008). <https://doi.org/10.1038/nn.2207>
- [28] M. Pertea, X. Lin, and S. L. Salzberg. 2001. GeneSplicer: A new computational method for splice site prediction. *Nucleic Acids Research* (2001). <https://doi.org/10.1093/nar/29.5.1185>
- [29] Mark D. Robinson, Davis J. McCarthy, and Gordon K. Smyth. 2009. edgeR: A Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics* (2009). <https://doi.org/10.1093/bioinformatics/btp616>
- [30] Christopher R. Sibley, Warren Emmett, Lorea Blazquez, Ana Faro, Nejc Haberman, Michael Briesse, Daniah Trabzuni, Mina Ryten, Michael E. Weale, John Hardy, Miha Modic, Tomaž Curk, Stephen W. Wilson, Vincent Plagnol, and Jernej Ule. 2015. Recursive splicing in long vertebrate genes. *Nature* (2015). <https://doi.org/10.1038/nature14466>
- [31] Yunfu Song and Zhijian Ou. 2018. Learning neural random fields with inclusive auxiliary generators. *arXiv preprint arXiv:1806.00271* (2018).
- [32] Tamara Steijger, Josep F. Abril, Pär G. Engström, Felix Kokocinski, Martin Akerman, Tyler Alioto, Giovanna Ambrosini, Stylianos E. Antonarakis, Jonas Behr, Paul Bertone, Regina Bohnert, Philipp Bucher, Nicole Cloonan, Thomas Derrien, Sarah Djebali, Jiang Du, Sandrine Dudoit, Mark Gerstein, Thomas R. Gingeras, David Gonzalez, Sean M. Grimmond, Roderic Guigó, Lukas Habegger, Jennifer Harrow, Tim J. Hubbard, Christian Iseli, Géraldine Jean, André Kahles, Julien Lagarde, Jing Leng, Gregory Lefebvre, Suzanna Lewis, Ali Mortazavi, Peter Niermann, Gunnar Ratsch, Alexandre Reymond, Paolo Ribeca, Hugues Richard, Jacques Rougemont, Joel Rozowsky, Michael Sammeth, Andrea Sboner, Marcel H. Schulz, Steven M.J. Searle, Naryttza Diaz Solorzano, Victor Solovyyev, Mario Stanke, Brian J. Stevenson, Heinz Stockinger, Armand Valsesia, David Weese, Simon White, Barbara J. Wold, Jie Wu, Thomas D. Wu, Georg Zeller, Daniel Zerbino, and Michael Q. Zhang. 2013. Assessment of transcript reconstruction methods for RNA-seq. *Nature Methods* (2013). <https://doi.org/10.1038/nmeth.2714>
- [33] Hitoshi Suzuki, Toshiki Kameyama, Kenji Ohe, Toshifumi Tsukahara, and Akila Mayeda. 2013. Nested introns in an intron: Evidence of multi-step splicing in a large intron of the human dystrophin pre-mRNA. *FEBS Letters* (2013). <https://doi.org/10.1016/j.febslet.2013.01.057>
- [34] Andrea Taladriz-Sender, Emma Campbell, and Glenn A. Burley. 2019. Splice-switching small molecules: A new therapeutic approach to modulate gene expression. *Methods* (2019). <https://doi.org/10.1016/j.ymeth.2019.06.011>
- [35] Jamal Tazi, Nadia Bakkour, and Stefan Stamm. 2009. Alternative splicing and disease. <https://doi.org/10.1016/j.bbdis.2008.09.017>
- [36] Eric T. Wang, Rickard Sandberg, Shujun Luo, Irina Khrebtkova, Lu Zhang, Christine Mayr, Stephen F. Kingsmore, Gary P. Schroth, and Christopher B. Burge. 2008. Alternative isoform regulation in human tissue transcriptomes. *Nature* (2008). <https://doi.org/10.1038/nature07509>
- [37] Joshua T. Witten and Jernej Ule. 2011. Understanding splicing regulation through RNA splicing maps. <https://doi.org/10.1016/j.jtig.2010.12.001>
- [38] Jianwen Xie, Yang Lu, Song-Chun Zhu, and Yingnian Wu. 2016. A theory of generative convnet. In *International Conference on Machine Learning*. 2635–2644.
- [39] Hui Y. Xiong, Babak Alipanahi, Leo J. Lee, Hannes Bretschneider, Daniele Merico, Ryan K.C. Yuen, Yimin Hua, Serge Gueroussov, Hamed S. Najafabadi, Timothy R. Hughes, Quaid Morris, Yoseph Barash, Adrian R. Krainer, Nebojsa Jojic, Stephen W. Scherer, Benjamin J. Blencowe, and Brendan J. Frey. 2015. The human splicing code reveals new insights into the genetic determinants of disease. *Science* (2015). <https://doi.org/10.1126/science.1254806>
- [40] Hui Yuan Xiong, Yoseph Barash, and Brendan J. Frey. 2011. Bayesian prediction of tissue-regulated splicing using RNA sequence and cellular context. *Bioinformatics* (2011). <https://doi.org/10.1093/bioinformatics/btr444>
- [41] Brian A. Yee, Gabriel A. Pratt, Brenton R. Graveley, Eric L. van Nostrand, and Gene W. Yeo. 2019. RBP-Maps enables robust generation of splicing regulatory maps. *RNA* (2019). <https://doi.org/10.1261/rna.069237.118>
- [42] Daniel R. et al. Zerbino. 2018. Ensembl 2018. *Nucleic Acids Research* (2018). <https://doi.org/10.1093/nar/gkx1098>
- [43] Yi Zhang, Xinan Liu, James MacLeod, and Jinze Liu. 2018. Discerning novel splice junctions derived from RNA-seq alignment: A deep learning approach. *BMC Genomics* (2018). <https://doi.org/10.1186/s12864-018-5350-1>
- [44] Zijun Zhang, Zhicheng Pan, Yi Ying, Zhijie Xie, Samir Adhikari, John Phillips, Russ P. Carstens, Douglas L. Black, Yingnian Wu, and Yi Xing. 2019. Deep-learning augmented RNA-seq analysis of transcript splicing. *Nature Methods* (2019). <https://doi.org/10.1038/s41592-019-0351-9>

- [45] Jasper Zuallaert, Frédéric Godin, Mijung Kim, Arne Soete, Yvan Saeys, and Wesley De Neve. 2018. Splicerover: Interpretable convolutional neural networks for improved splice site prediction. *Bioinformatics* (2018). <https://doi.org/10.1093/bioinformatics/bty497>

A APPENDIX

A.1 Additional Results

Table 6: Performance of DCEN with early-stopping versus 5x training steps.

Model	Spearman Cor.	Pearson Cor.
Early-Stopped	0.623	0.651
5x Training steps	0.620	0.650

Figure 6: Comparison of the BCL2 gene donor sites' ψ , experimental vs predicted. Left: box plots of experimental values (grounds truths) of BCL2 donor sites in three different tissues: blood, skeletal muscle and lung, denoted as BLD, SKM and LNG respectively. Right: box plots of predicted ψ for the same group of samples using DCEN; SpliceAI-cl prediction (black triangles) shown for comparison. Please note that SpliceAI-cl prediction uses sequence information only and the prediction will be the same for every sample. Donor site coordinates (y axis) are counted relative to the genomic coordinate of the first exon start of BCL2. Data from 50 samples from each tissue are shown.

A.2 Proof

THEOREM A.1. *Given the energy levels of all the possible discrete states of a system, the probability of a particular state T_i is the softmax output of its energy E_{T_i} with respect to those of all other possible states in the system, i.e.,*

$$P_i = \frac{\exp(-E_{T_i})}{\sum_j \exp(-E_{T_j})} = \text{Softmax}_i(E) \quad (9)$$

PROOF. From Boltzmann distribution, the probability that a system takes on a particular state (x) can be expressed as:

$$\begin{aligned} p_\theta(x) &= \frac{\exp(-E_\theta(x))}{Z(\theta)} \\ &= \frac{h(x)}{Z(\theta)} \end{aligned} \quad (10)$$

where

$$\begin{aligned} Z(\theta) &= \int \exp(-E_\theta(x)) dx \\ &= \int h(x) dx \end{aligned} \quad (11)$$

is known as the partition function.

Since the probabilities of all possible states sum to 1, we have

$$1 = \mathbb{E}_{x \sim p_\theta}[1] = \sum_x \frac{h(x)}{Z} \quad (12)$$

which gives

$$Z = \sum_x h(x) . \quad (13)$$

Through importance sampling with another probability distribution q , we can express Z as

$$\begin{aligned} Z &= \sum_x \frac{h(x)}{q(x)} q(x) \\ &= \frac{1}{n} \sum_i \frac{h(x_i)}{q(x_i)} , \quad x_i \sim q . \end{aligned} \quad (14)$$

Using an uniform discrete distribution as q where all k possible states (x_i) have the same probability $q(x_i) = (1/k)$, we get

$$\begin{aligned} Z &= \frac{1}{k} \sum_i \frac{h(x_i)}{(1/k)} \\ &= \sum_i h(x_i) \end{aligned} \quad (15)$$

Combining Eq. (11) and (15), this gives

$$\begin{aligned} p_\theta(x_i) &= \frac{h(x_i)}{\sum_j h(x_j)} \\ &= \frac{\exp(-E_\theta(x_i))}{\sum_j \exp(-E_\theta(x_j))} \\ &= \text{Softmax}_i(E) \end{aligned} \quad (16)$$

□