DISCOVERING CODING LNCRNAS USING DEEP LEARNING TRAINING DYNAMICS

by AFSHAN NABI

Submitted to the Graduate School of Engineering and Natural Sciences in partial fulfilment of the requirements for the degree of Master of Science

> Sabancı University July 2021

DISCOVERING CODING LNCRNAS USING DEEP LEARNING TRAINING DYNAMICS

Approved by:



Date of Approval:

AFSHAN NABI 2021 $\ensuremath{\mathbb{O}}$

All Rights Reserved

ABSTRACT

DISCOVERING CODING LNCRNAS USING DEEP LEARNING TRAINING DYNAMICS

AFSHAN NABI

COMPUTER SCIENCE AND ENGINEERING M.A. THESIS, JULY 2021

Thesis Supervisor: Prof. Öznur Taştan Okan

Keywords: Noncoding RNA, coding lncRNAs, Deep learning, Training dynamics

Long non-coding RNAs (lncRNAs) are the largest class of non-coding RNAs (ncR-NAs). However, recent experimental evidence has shown that some lncRNAs contain small open reading frames (sORFs) that are translated into functional micropeptides. Current methods to detect misannotated lncRNAs rely on ribosome-profiling (riboseq) experiments, which are expensive and cell-type dependent. In addition, while very accurate machine learning models have been trained to distinguish between coding and non-coding sequences, little attention has been paid to the increasing evidence about the incorrect ground-truth labels of some lncRNAs in the underlying training datasets. We present a framework that leverages deep learning models' training dynamics to determine whether a given lncRNA transcript is misannotated. Our models achieve AUC scores > 91% and AUPR > 93% in classifying non-coding vs. coding sequences while allowing us to identify possible misannotated lncRNAs present in the dataset. Our results overlap significantly with a set of experimentally validated misannotated lncRNAs as well as with coding sORFs within lncRNAs found by a ribo-seq dataset. The general framework applied here offers promising potential for use in curating datasets used for training coding potential predictors and assisting experimental efforts in characterizing the hidden proteome encoded by misannotated lncRNAs.

ÖZET

DERIN ÖĞRENME EĞITIM DINAMIKLERINI KULLANARAK KODLAMA LNCRNA'LARINI KEŞFETME

AFSHAN NABI

BILGISAYAR BILIMI VE MÜHENDISLIĞI YÜKSEK LİSANS TEZİ, TEMMUZ 2021

Tez Danışmanı: Prof. Dr. Öznur Taştan Okan

Anahtar Kelimeler: kodlamayan RNA, kodlayan lncRNA, Derin Öğrenme, Eğitim Dinamikler

Uzun kodlamayan RNA'lar (lncRNA'lar), kodlamayan RNA'ların (ncRNA'lar) en büvük sınıfıdır. Bununla birlikte, son denevsel kanıtlar, bazı lncRNA'ların, fonksivonel mikropeptidlere çevrilen küçük açık okuma çerçeveleri (sORF'ler) içerdiğini göstermiştir. Yanlış yorumlanmış lncRNA'ları tespit etmek için mevcut yöntemler, pahalı ve hücre tipine bağlı olan ribozom profili oluşturma (ribo-seq) deneylerine dayanır. Ek olarak, kodlama yapan ve kodlamayan dizileri ayırt etmek için çok hassas makine öğrenimi modelleri eğitilmiş olsa da, temel eğitim veri kümelerinde bazı lncRNA'ların yanlış yer-gerçeği etiketleri hakkında artan kanıtlara çok az ilgi gösterilmiştir. Belirli bir lncRNA transkriptinin yanlış yorumlanıp işaretlenmediğini belirlemek için derin öğrenme modellerinin eğitim dinamiklerinden yararlanan bir cerceve sunuvoruz. Modellerimiz, veri kümesinde bulunan olası yanlış yorumlanmış lncRNA'ları belirlememize izin verirken, kodlama yapmayan ve kodlama dizilerini sınıflandırmada AUC puanları > 91% ve AUPR > 93% elde eder. Sonuçlarımız, bir ribo-seq veri kümesi tarafından bulunan lncRNA'lar içindeki sORF'leri kodlamanın yanı sıra, deneysel olarak doğrulanmış yanlış yorumlanmış bir dizi lncRNA ile önemli ölçüde örtüşmektedir. Burada uygulanan genel çerçeve, potansiyel tahmin edicileri kodlamak için kullanılan veri kümelerinin küratörlüğünde kullanım için umut verici bir potansiyel sunar ve yanlış yorumlanmış lncRNA'lar tarafından kodlanan gizli proteomun karakterize edilmesinde deneysel çabalara yardımcı olur.

ACKNOWLEDGEMENTS

Many thanks to:

- Öznur Taştan, and Ogün Adbali for interesting discussions.
- Sabanci University for material support and work opportunities.
- My family and friends for emotional support.

For *Öznur*, who always encourages me. For *Kazım*, who knows when I need gentle or tough love, and always obliges. For *Burhan*, who provides unconditional acceptance, support and friendship. For *Hajira*, who taught me love, loyalty and endurance. For *Ghulam*, who taught me hard-work, duty and patience. For *Hatice*, *Ramazan*, *Rümeysa*, *Furkan*, and *Kübra*, who shone lights in the dark.

TABLE OF CONTENTS

LIST OF TABLES									
\mathbf{LI}	ST (OF FIGURES	xi						
1.	INT	RODUCTION	1						
1.1. Some lncRNAs are misannotated									
	1.2.	. Ribosome Profiling can be used to identify misannotated lncRNAs							
		experimentally							
	1.3.	3. Computational Methods for Identifying misannotated lncRNAs							
	1.4.	Machine Learning for Sequence Classification	3						
	1.5.	Training Dynamics for Identifying lncRNAs	4						
2.	ME	THODS	5						
	2.1.	The Overall Framework	5						
	2.2.	Datasets	5						
		2.2.1. Pre-trained embeddings for deep learning models	6						
		2.2.2. Sequences for training deep learning models	7						
		2.2.3. Unsupervised clustering (t-SNE) of RNAs	7						
	2.3.	Deep Learning Model Architectures	7						
		2.3.1. Convolutional neural network:	8						
		2.3.2. LSTM:	8						
		2.3.3. Transformer:	8						
		2.3.4. Model Evaluation Set Up	8						
		2.3.5. Identifying misannotated lncRNAs using training dynamics	9						
		2.3.6. Unsupervised clustering (t-SNE) of RNA sequences	10						
		2.3.7. Comparison to cncRNAdb & ribo-seq	10						
3.	RES	SULTS	11						
	3.1.	Prediction performance of classifying coding vs. non-coding RNAs	11						
	3.2.	Training dynamics of deep learning models can be used to identify misannotated lncRNAs	13						

BIBLIOGRAPHY			
4.	DIS	CUSSION	25
	3.8.	Exploring features learnt by models	19
	3.7.	Misannotated lncRNAs exist in a continuous cluster with coding RNAs $% \mathcal{A}$	18
	3.6.	Candidate coding lncRNAs discovered using training dynamics	17
		those discovered by ribosome profiling	16
	3.5.	Misannotated lncRNAs from training dynamics overlap significantly	
		NAs discovered by ribo-seq	15
		experimentally validated coding lncRNAs & with misannotated lncR- $$	
	3.4.	Misannotated lncRNAs overlap significantly with manually curated,	
		lncRNAs	14
	3.3.	Different deep learning architectures find common misannotated	

LIST OF TABLES

	The test-data performances of the different models trained to	Table 3.1 .
	fy long non-coding RNAs and coding RNAs. AUC and AUPR	classif
12	nicro-averaged.	are m

LIST OF FIGURES

Figure 2.1. Workflow for identifying misannotated lncRNAs by examining the training dynamics of deep learning models. All RNA sequences are constrained to be between 200-4000 nucleotides long. From each RNA sequence 3-mer 'words' are obtained by using a window that slides by 1 nucleotide at each step. For each 3-mer 'word', 100-dimensional embeddings Ng (2017) are obtained. Each RNA also has an associated ground-truth label, i.e. each RNA is labelled as coding or non-coding. Deep learning models are trained using 100-dimensional embeddings for contiguous 3-mers from the sequences. At the end of each training epoch, the predicted probabilities for each RNA being coding or non-coding are saved. After training, the mean and standard deviation for the ground-truth label probability prediction are calculated and misannotated lncRNAs are identified. These are compared to lncRNAs containing translated sORFs determined from ribo-seq data.

6

- Figure 3.1. Analysis of the training dynamics of deep learning models. (a) Estimated coding probability across all training epochs shown for five RNAs. We expect coding and non-coding RNAs to have high and low coding probabilities respectively; this is the case for examples A and E. In this work, we are trying to identify noncoding RNAs that might be misannotated, when they are in fact coding. We are interested in lncRNAs- like B (ENST00000447563)which have consistently high estimated coding probability, despite having the ground truth-label 'Non-coding'. These are candidate coding lncRNA. In support of this observation, ENST00000447563 (also known as linc00689) was recently found to be protein coding Hartford & Lal (2020). C and D show examples of ambiguous samples, i.e. they show a large change in estimated coding probabilities as model training progresses, so we are not sure whether they are mislabelled or not. (b) Mean (y-axis) and standard deviation (x-axis) of ground truth class probability predictions across all training epochs can be used to determine mislabelled samples. Candidate misannotated RNAs are those in the Hard-to-Learn region i.e. RNAs with low mean and standard deviation for the ground truth class probability. .
- Figure 3.2. Training dynamics of deep learning models can be used to identify misannotated lncRNAs. To prove that training dynamics can be used to identify mislabelled RNAs, we sampled 5% of the data with $\mu_{ground\ truth}$ and $\sigma_{ground\ truth}$ less than 0.8 and 0.2 respectively. These are samples for which we have medium-high confidence that the ground truth labels are correct (before). After we flip the labels (coding RNAs become, non-coding RNAs and vice-versa), the samples move into the hard-to-classify region of the training dynamics summary plot.

20

21

Figure 3.3. Comparison with cncRNAdb (a) Comparison of the hard-to-learn ncRNAs obtained from by CNN, LSTM and Transformer models' training dynamics. (b) Comparison of the number of common RNAs that are present in the deep learning dataset and in cncRNAdb. (c) Intersection of the misannotated lncRNAs identified in the union of LSTM, CNN, and Transformer set with the common RNAs found in cncRNAdb. (d)Intersection of the misannotated lncRNAs identified in the intersection of at least 2 of LSTM, CNN, and Transformer set with the common RNAs found in cncRNAdb. (e) Intersection of the misannotated lncRNAs identified in the intersection of at least 2 of LSTM, CNN, and Transformer set with the common RNAs identified in the intersection of the misannotated lncRNAs. (e) Intersection of the misannotated lncRNAb. (e) Intersection of LSTM, CNN, and Transformer set with the common RNAs identified in the intersection of LSTM, CNN, and Transformer set with the common RNAs identified in the intersection of LSTM, CNN, and Transformer set with the common RNAs identified in the intersection of LSTM, CNN, and Transformer set with the common RNAs identified in the intersection of LSTM, CNN, and Transformer set with the common RNAs found in cncRNAdb. ... 22

Figure 3.4. Comparison to ribo-seq data based methods used to find misannotated lncRNAs from : (a) FLOSS (p-value ≈ 0), (b) ORFScore (p-value 4e - 320) and (c) PhyloP (p-value (1e - 32)) for the dataset from (Elkon, Loayza-Puch, Korkmaz, Lopes, Van Breugel, Bleijerveld, Altelaar, Wolf, Lorenzin, Eilers & others, 2015). Background set has 26857 lncRNAs..... 22Figure 3.5. Some of the candidate coding lncRNAs discovered in this work, with evidence for coding potential aggregated from other sources. Ribo-seq identifies RNAs associated with ribosomes, which are likely to be translated. Ribo-seq p-values (a combined score from FLOSS, ORFScore and PhyloP) show the likelihood of the identified RNA being a true-positive. PeptideShaker analyzes publicly available MS data and provides a confidence score for each peptide. The BLAST Hit is the top hit from running the query in BLASTx. CPC2, CPAT and RNASamba are tools for the coding potential prediction of a given RNA.) 23Figure 3.6. Misannotated lncRNAs exist in a continuous cluster with coding RNAs. t-SNE clusters obtained from hand-crafted features (Tong & Liu, 2019) generated from nucleotide sequences. Labels (Coding, lncRNA etc.) are only used for visualizing the clusters, not for generating the clusters. Putative misannotated lncRNAs lie in a cluster contiguous with coding RNAs. There are other clusters of ncRNAs that are well separated from coding RNAs and that do not contain any putative misannotated lncRNAs. 24Figure 3.7. Attention maps explain which parts of sequence are important for making a classification decision. An example attention map extracted from first convolutional layer of CNN model for ENST00000447563 (a misannotated lncRNA). The ground-truth for this RNA is 'Non-coding'. However, CNN model predicts that this is a non-coding RNA with probability 0.12. Attention visualization shows regions with contiguous 'A' nucleotides have high activation weights. 24

1. INTRODUCTION

1.1 Some lncRNAs are misannotated

Genome-wide transcriptome analyses have revealed that the vast majority of the human genome is transcribed; but only 2% of the human genome is annotated as protein coding (Djebali, Davis, Merkel, Dobin, Lassmann, Mortazavi, Tanzer, Lagarde, Lin, Schlesinger & others, 2012). A considerable fraction of transcripts are annotated as ncRNAs and lncRNAS constitute the largest category of ncRNAs (Derrien, Johnson, Bussotti, Tanzer, Djebali, Tilgner, Guernec, Martin, Merkel, Knowles & others, 2012). While lncRNAs studied are known to play vital roles in cellular processes such as regulation of translation, transcription, chromatin modification and mRNA stability (Batista & Chang, 2013; Rinn & Chang, 2012; Ulitsky & Bartel, 2013), functions of most lncRNAs remain unknown. Moreover, although lncRNAs by definition- do not code for proteins, recent studies have shown that short the open reading frames (sORFs) within some lncRNAs are translated into micropeptides of a median length of 23 amino acids (Andrews & Rothnagel, 2014; Choi, Kim & Nam, 2019; Couso & Patraquim, 2017; Hartford & Lal, 2020; Ji, Song, Regev & Struhl, 2015; Lu, Zhang, Lian, Sun, Meng, Chen, Sun, Yin, Li, Zhao & others, 2019). The translation events of lncRNAs were overlooked previously because the open reading frames (ORFs) present in lncRNAs do not meet the conventional criteria of an ORF: that it encodes at least 100 amino acids in eukaryotes (Hartford & Lal, 2020). Despite this, recent studies have shown that micropeptides translated from lncRNAs perform vital functions across species, including bacteria, flies and humans (Hartford & Lal, 2020; Matsumoto, Pasut, Matsumoto, Yamashita, Fung, Monteleone, Saghatelian, Nakayama, Clohessy & Pandolfi, 2017; Nelson, Makarewich, Anderson, Winders, Troupes, Wu, Reese, McAnally, Chen, Kavalali & others, 2016). Therefore, identifying misannotated lncRNAs is a necessary step towards the functional

characterization of this large class of transcripts.

1.2 Ribosome Profiling can be used to identify misannotated lncRNAs

experimentally

Experimental identification of coding transcripts is performed using ribosome profiling (ribo-seq), which involves capturing and sequencing RNA fragments protected by ribosomes (Ingolia, 2014). Use of ribo-seq data has revealed many unexpected micropeptides (Brar & Weissman, 2015), including sORFs within lncRNAs (Ingolia, Ghaemmaghami, Newman & Weissman, 2009). However, since ribo-seq data is known to contain false positives (Ingolia et al., 2009; Ingolia, Lareau & Weissman, 2011), several computational methods have been proposed to identify true translated ribo-seq fragments. These include FLOSS (Ingolia, 2014), ORFscore (Bazzini, Johnstone, Christiano, Mackowiak, Obermayer, Fleming, Vejnar, Lee, Rajewsky, Walther & others, 2014) and PhyloP (Miller, Rosenbloom, Hardison, Hou, Taylor, Raney, Burhans, King, Baertsch, Blankenberg & others, 2007; Olexiouk, Van Criekinge & Menschaert, 2018). FLOSS (Ingolia, 2014) relies on the typical length of ribo-seq fragments to determine truly coding ribo-seq fragments. ORFscore (Bazzini et al., 2014) relies on the property that translating ribosomes shift by three nucleotides (ribosome phasing), which leads to a characteristic pattern wherein true positive fragments have higher sequencing reads every third nucleotide. PhyloP is used to find true translated ribo-seq fragments by probing conservation across species (Miller et al., 2007; Olexiouk et al., 2018). These computational methods applied over riboseq data can be used to find sORFs that are both translated and located within lncRNAs. However, one major limitation of relying on ribo-seq data to identify misannotated lncRNAs is that not all transcripts are likely to be transcribed and translated at a given time point in a given cell. To obtain a complete picture of the misannotated lncRNAs in the genome, different cell types, at different developmental stages, under different environmental conditions need to be sequenced and analyzed. In contrast, the nucleotide sequence of a lncRNA transcript is unlikely to change across cell types and conditions. Therefore, methods to detect misannotated lncRNAs from nucleotide sequences will be useful in assisting experimental efforts and available ribo-seq based computational methods.

1.3 Computational Methods for Identifying misannotated lncRNAs

Once sufficient coding sORFs have been detected, a dataset containing positive (coding) and negative (non-coding) examples can be built to train models to predict the coding potential of a given sORF. These methods can then be used to assess the coding potential of a transcript. For instance, logistic regression (Zhu & Gribskov, 2019) and SVM (Tong, Hong, Xie & Liu, 2020) based models have been proposed to predict the coding potential of a given sORF with sequence length \leq 303 nucleotides. To determine whether a lncRNA is misannotated by using these methods requires first to extract all possible sORFs in the lncRNA and then assess the coding potential of each of these sORFs. However, while it is possible to predict the coding potential of lncRNA sORFs with these tools, it is impossible to evaluate the performance since the data on which lncRNA sORFs are truly coding is very sparse (Zhu & Gribskov, 2019).

1.4 Machine Learning for Sequence Classification

Several classical machine learning (Kang, Yang, Kong, Hou, Meng, Wei & Gao, 2017; Kong, Zhang, Ye, Liu, Zhao, Wei & Gao, 2007; Tong & Liu, 2019; Wang, Park, Dasari, Wang, Kocher & Li, 2013) and deep learning (Baek, Lee, Kwon & Yoon, 2018; Camargo, Sourkov, Pereira & Carazzolle, 2020; Hill, Kuintzle, Teegarden, Merrill III, Danaee & Hendrix, 2018) based models, which focus on longer length nucleotide sequences as input, have also been developed to predict the coding potential of a given RNA. Most of these methods demonstrate very high prediction performance. However, using these, it is not possible to identify lncRNAs that might be misannotated. This is because these models do not incorporate any strategy to deal with misannotated lncRNAs in the underlying training datasets. To ensure that we are not overfitting the models to learn biologically irrelevant decision boundaries, there is a need to find ways to determine possible misannotated RNAs in the underlying training datasets.

1.5 Training Dynamics for Identifying lncRNAs

We present a framework that leverages deep learning models' training dynamics to determine whether a given lncRNA transcript in the dataset is misannotated. In particular, we train convolutional neural network (CNN) (LeCun, Boser, Denker, Henderson, Howard, Hubbard & Jackel, 1989), long short term memory (LSTM) (Hochreiter & Schmidhuber, 1997), and Transformer (Vaswani, Shazeer, Parmar, Uszkoreit, Jones, Gomez, Kaiser & Polosukhin, 2017) architectures to predict whether a given nucleotide sequence is non-coding or coding and use the training dynamics to identify possible misannotated lncRNAs (Swayamdipta, Schwartz, Lourie, Wang, Hajishirzi, Smith & Choi, 2020). Our models learn biologically relevant features to distinguish between coding and non-coding RNAs with average AUC scores >91% and identify many misannotated lncRNAs. By generating unsupervised clusters of coding and non-coding RNAs, we observe that there might be a continuity in the embedded space between coding and misannotated lncRNAs. Finally, our results show a significant overlap with previous methods that use riboseq data to identify misannotated lncRNAs as well as with a set of experimentally validated misannotated lncRNAs. This work represents the first instance where deep learning model training dynamics are successfully applied to identify misannotated lncRNAs from nucleotide sequences. This approach can be applied to better curate datasets for training coding potential prediction models and can be applied alongside ribo-seq data to identify misannotated lncRNAs with high confidence.

2. METHODS

2.1 The Overall Framework

The workflow for determining misannotated lncRNAs is described in Figure 2.1. The main steps are as follows. We train deep learning based sequence classification models that can distinguish coding and non-coding RNAs. Once we establish that models can achieve good performance on the held-out test data, we retrain a final model on all the data and inspect its training dynamics to find possibly misanno-tated ncRNAs. By focusing especially on the union of the lncRNAs identified as misannotated by all the models, we arrive at a final list of putative misannotated lncRNAs. We compare this list to experimentally validated coding ncRNAs as well as to a ribo-seq dataset. We use unsupervised clustering to find where possibly misannotated lncRNAs are located within the broader RNA clusters. We also study the features used by the deep learning models to make classification decisions. In the following sections we detail the dataset, the sequence classification models trained and the other analysis we conducted.

2.2 Datasets

2.2.1 Pre-trained embeddings for deep learning models



Figure 2.1 Workflow for identifying misannotated lncRNAs by examining the training dynamics of deep learning models. All RNA sequences are constrained to be between 200-4000 nucleotides long. From each RNA sequence 3-mer 'words' are obtained by using a window that slides by 1 nucleotide at each step. For each 3-mer 'word', 100-dimensional embeddings Ng (2017) are obtained. Each RNA also has an associated ground-truth label, i.e. each RNA is labelled as coding or non-coding. Deep learning models are trained using 100-dimensional embeddings for contiguous 3-mers from the sequences. At the end of each training epoch, the predicted probabilities for each RNA being coding or non-coding are saved. After training, the mean and standard deviation for the ground-truth label probability prediction are calculated and misannotated lncRNAs are identified. These are compared to lncRNAs containing translated sORFs determined from ribo-seq data.

add

2.2.2 Sequences for training deep learning models

We use the dataset of human RNA nucleotide sequences compiled by (Tong & Liu, 2019) to train the sequence classification models. After filtering to remove noncoding RNA sequences < 200 nucleotides in length, the data comprises of 38,051 coding RNA and 19,472 non-coding RNA sequences. Filtering non-coding RNAs by length was necessary since the length distributions of coding and non-coding RNAs in the dataset was very different; non-coding RNAs are noticeably shorter than coding RNAs. For the deep learning models to learn biologically relevant features in order to distinguish between coding and non-coding RNAs, equalizing the sequence length distributions was necessary. If the sequence lengths of ncRNAs are significantly shorter than those of coding RNAs, then sequence length itself might be used by the models as a feature distinguishing between coding and non-coding RNAs.

2.2.3 Unsupervised clustering (t-SNE) of RNAs

2.3 Deep Learning Model Architectures

We train CNN (LeCun et al., 1989), LSTM (Hochreiter & Schmidhuber, 1997) and Transformer (Vaswani et al., 2017) models to classify non-coding and coding RNA sequences. Each input sequence is truncated to a length of 4000 nucleotides before being input to the deep learning models. The sequences are encoded as 1-nucleotide sliding window 3-mers using the 100-dimensional DNA-embeddings generated by (Ng, 2017). All three models are implemented using Keras (Chollet & others, 2015). We use ReLu as the activation function. We trained all models to minimize the sparse categorical cross-entropy loss using the Adam optimizer (Kingma & Ba, 2014). In all cases, we use a batch size of 64.

2.3.1 Convolutional neural network:

For the CNN, encoded sequences are fed into an embedding layer which is followed by 3 layers of 1-D convolution (each with 128 units and filter size 5) and max-pooling (5 units). These are followed by a dense layer of 128 units.

2.3.2 LSTM:

For the LSTM, encoded sequences are fed into an embedding layer which is followed by 2 layers of 1-D convolution (each with 128 units and filter size 5) and max-pooling (5 units), followed by a bi-directional LSTM layer. These are followed by a dense layer of 128 units.

2.3.3 Transformer:

Encoded sequences are added to a positional encoding and fed into a transformer block followed by global average pooling, dropout and a dense layer of 64 units. The transformer block comprises of a single-headed self-attention layer and a dense layer both followed by layer normalization.

2.3.4 Model Evaluation Set Up

We use the human coding & non-coding train and test datasets provided by (Tong & Liu, 2019). We set aside 20% of the training data as the validation data. We use Keras Tuner (O'Malley, Bursztein, Long, Chollet, Jin, Invernizzi & others, 2019) to find the optimal set of hyperparameters for the deep learning models. We created a hyperparameter search space for different model architecture and hyperparameter assignment values and used the Hyperband tuner (Li, Jamieson, DeSalvo, Rostamizadeh & Talwalkar, 2017) to find the optimal parameters based on validation loss. We tried the following choices for given hyperparameters: dense layer units 64, 128, and 256, 1-D convolutional filters (64 and 128, LSTM units 64, 128, and 256, dropout 0.2, 0.3, 0.4 and 0.5 and learning rate (logarithmic sampling between e-2 and e-4. We used the best model returned by the Hyperband tuner and retrain a model on the train-validation data to calculate and assess these models' performances on the held-out test data. Once the test performances are attained, we

rebuilt the models on all data to find the misannotated ncRNAs.

Since the training dataset is imbalanced in favor of coding RNA, we used classweights inversely proportional to the number of class samples to ensure learning. Moreover, since a coding RNA is unlikely to be misannotated, we penalized coding RNA misclassifications 5 times more than non-coding RNA misclassifications.

2.3.5 Identifying misannotated lncRNAs using training dynamics

We inspect the deep learning models' training dynamics to find possible misannotated lncRNAs. Swayamdipta et al. (Swayamdipta et al., 2020) report that it is possible to identify possibly mislabelled training samples in a given dataset by inspecting how model predictions for samples behave during training. We employ this strategy; at the end of each training epoch, the deep learning models are evaluated on the training examples and the predictions for the class probabilities are saved. Consider a training dataset of size $N, \mathcal{D} = \{(\boldsymbol{x}, y^*)_i\}_{i=1}^N$ where the *i* th instance consists of the observation, \boldsymbol{x}_i and its true label under the task, y_i^* . We calculate the mean and the standard deviation of the posterior probability of the ground-truth label for example *i* over *E* epochs as follows (Swayamdipta et al., 2020):

(2.1)
$$\hat{\mu}_{i} = \frac{1}{E} \sum_{e=1}^{E} p_{\theta^{(e)}}(y_{i}^{*} \mid \boldsymbol{x}_{i}), \qquad \hat{\sigma}_{i} = \sqrt{\frac{\sum_{e=1}^{E} \left(p_{\theta^{(e)}}(y_{i}^{*} \mid \boldsymbol{x}_{i}) - \hat{\mu}_{i} \right)^{2}}{E}}$$

where $p_{\theta^{(e)}}$ denotes the probability assigned at the end of the e^{th} epoch by the model parameterized with $\theta^{(e)}$. Using the mean and the standard deviation of the predicted probability of ground-truth class across all epochs, the training dataset can be divided into three groups: easy-to-learn, ambiguous, and hard-to-learn. The hard-to-learn samples are those with low mean and low standard deviation of the true class probabilities. In other words, the model consistently misclassifies these samples across training epochs. We retrain the models using both the training & test data and consider the lncRNAs within this hard-to-learn class as candidates for misannotation.

2.3.6 Unsupervised clustering (t-SNE) of RNA sequences

To analyze the coding and noncoding transcript distributions of the data, we calculated features on for all RNAs in the dataset, based on properties of the transcripts as in (Tong & Liu, 2019). These features include ORF length, ORF quality, nucleotide distribution, translated peptide stability etc. used by (Tong & Liu, 2019) see Supplementary Table 1 for more details. Using these features, we apply T-distributed stochastic neighbor embedding (t-SNE) (Maaten & Hinton, 2008) (SciKit-learn implementation (Pedregosa, Varoquaux, Gramfort, Michel, Thirion, Grisel, Blondel, Prettenhofer, Weiss, Dubourg, Vanderplas, Passos, Cournapeau, Brucher, Perrot & Duchesnay, 2011), perplexity=150, iterations=1000, learning rate=200) to reveal RNA clusters.

2.3.7 Comparison to cncRNAdb & ribo-seq

We downloaded data from the cncRNAdb (Huang, Wang, Zhao, Wang, Liu, Li, Cui, Li, Feng, Luo & others, 2020), a resource that provides a manually curated list of experimentally validated ncRNAs found to be coding. We filtered data to get lncRNAs found to be coding in *Homo sapiens* and compared the list to the misannotated lncRNA candidates generated from the deep learning models.

Next, we compared the list of misannotated lncRNAs generated by our models to a ribo-seq dataset. We downloaded data on sORFs identified in the ribo-seq data generated by (Elkon et al., 2015) from sORFs.org (Olexiouk et al., 2018). This database provides computations of values of FLOSS (Ingolia, 2014), ORFscore (Bazzini et al., 2014) and PhyloP (Miller et al., 2007) metrics for RNAs identified from the ribo-seq data. We used RNAs annotated as lncRNAs and present in both the sequence dataset (used to train deep learning models) and the Ribo-seq dataset in our analysis. According to previous considerations, to get the list of lncRNAs containing translated sORFs, we used the following cutoff values: 'Good' for the Floss-classification, ORFscore > 6 and PhyloP > 4 (Olexiouk et al., 2018).

3. RESULTS

3.1 Prediction performance of classifying coding vs. non-coding RNAs

Prediction performances calculated on the held-out test set for the models trained are provided in Table 3.1 and show that our models perform well on the classification task. The LSTM model achieves the highest classification performance with 94% AUC and 96% AUPR. The CNN model performs similarly well with 93% AUC and 95% AUPR, while the transformer achieves 91% AUC and 93% AUPR. Since our aim is to study the underlying dataset and find misannotated lncRNAs, higher prediction performance is not the primary focus. Instead, since we know that the training dataset contains lncRNAs that have incorrect ground-truth labels, we want to ensure that the models are not being overfitted to learn features that might not be relevant to learning the biological distinction between coding and non-coding RNAs as encoded in the nucleotide sequences. In the following sections, we detail how we employ these models to discover possibly misannotated lncRNAs in the underlying dataset.

Table 3.1 The test-data performances of the different models trained to classify long non-coding RNAs and coding RNAs. AUC and AUPR are micro-averaged.

	AUC	AUPR		Precision	Recall	F1-Score
ISTM	0.94	0.96	Non-Coding	0.93	0.95	0.94
			Coding	0.95	0.94	0.94
CNN	0.93	0.95	Non-Coding	0.93	0.92	0.93
CININ			Coding	0.93	0.94	0.94
Transformer	ner 0.91	0.93	Non-Coding	0.93	0.88	0.90
mansiormer			Coding	0.90	0.94	0.92

3.2 Training dynamics of deep learning models can be used to identify

misannotated lncRNAs

Having evaluated the CNN, LSTM and Transformer models to distinguish between coding RNA and non-coding RNA, we retrain the models using all data and inspect each instances training dynamics. During the training phase of each model, we track the coding probability predictions for each RNA. Figure 3.2a shows the predictions for the coding probability for three different RNAs across all training epochs for the LSTM model. For example, the coding probability predictions for ENST00000447563 (shown in orange) -an RNA annotated as long non-coding (ground-truth)- are consistently high. In other words, as model training progresses, this RNA is invariably classified as coding. It was recently shown that ENST00000447563 has been misannotated as lncRNA when it can, in fact, code for a protein (Hartford & Lal, 2020). Two other examples of correctly annotated coding and non-coding RNA are also shown in Figure 3.2a. By studying the predictions made by models as they are under training, it is possible to identify putative misannotated lncRNAs.

Figure 3.1b expands upon this idea: calculating the mean and standard deviation of predicted probability for the ground-truth class across all training epochs provides a measure of identifying misannotated lncRNAs. lncRNAs in the lower left quadrant of Figure 3.1b are considered putative misannotated lncRNAs; these samples have low mean and standard deviation for the predicted probability of the ground-truth class over all training epochs. In other words, these lncRNAs are consistently classified into the non-ground-truth class (coding) and therefore, are likely to be misannotated. It is interesting to note that the majority of the putative mislabelled samples have ground-truth label ncRNA. This observation supports the notion that the current method for identifying putative misannotated lncRNAs is reasonable. This is because an RNA with ground-truth 'coding' is unlikely to be misannotated. In conclusion, many lncRNAs might be misannotated and sequence information combined with training dynamics of deep learning based classifiers might help identify such misannotations.

3.3 Different deep learning architectures find common misannotated

lncRNAs

Figure 3.3a shows the overlap between the lists of misannotated lncRNAs generated by CNN, LSTM and Transformer models. It is interesting to note that despite the difference in network architectures, the intersection of possible misannotated lncRNAs is large. The CNN model identifies the smallest number of candidate misannotated ncRNAs. It is interesting to note that the number of common candidates identified by Transformer and LSTM but not by CNN (912 in total) is large as compared to the common candidates between CNN & Transformer only (78) and between LSTM & CNN only (167). 1271 candidates are identified by all 3 models.

3.4 Misannotated lncRNAs overlap significantly with manually curated,

experimentally validated coding lncRNAs & with misannotated lncRNAs discovered by ribo-seq

The cncRNA database provides a manually curated list of experimentally validated coding lncRNAs (Huang et al., 2020). We report the overlap of the misannotated ncRNAs identified by training dynamics with this experimentally validated set. Figure 3.3b shows that 442 RNAs that are present in both the datasets. Out of these 442 RNAs common with cncRNAdb, 229, 179, and 109 are discovered by the union, intersection of at lease two, and intersection of all three of LSTM, CNN and Transformer models. The first thing of note is that while the training dynamics have the ability to identify *some* misannotated RNA, they are not able to identify *all* of them. The second thing to note is that the union of the misannotated RNAs shows the largest overlap with experimentally validates misannotated RNAs. This suggests that different models is more effective at identifying misannotated RNAs. It might be that different architectures of the CNN, LSTM and Transformer predispose them towards finding certain misannotated RNAs. Clearly the different models find different decision boundaries for the classification task they are trained on. This also suggests that an ensemble of models could be trained to identify misannotated lncRNAs with more confidence. The third thing of note is that training dynamics find a lot more candidates for misannotated RNAs than the 442 that are present in the intersection with cncRNAdb. For instance, the union of LSTM, CNN and Transformer models finds 3374 misannotated RNAs that are not present in cncR-NAdb. This suggests that there are a lot more coding lncRNAs than are currently known, even though an unknown fraction of the 3374 might be false positives.

3.5 Misannotated lncRNAs from training dynamics overlap significantly

those discovered by ribosome profiling

Next, we compared the overlap between the misannotated lncRNAs discovered by our deep learning models with high-throughput ribo-seq dataset from several different cell line compiled by (Olexiouk et al., 2018). Figure S4 shows the counts for lncRNAs obtained by applying 3 different methods (FLOSS, ORFScore and PhyloP) to identify true positives from ribo-seq data. For FLOSS, lncRNAs with a classification of 'Good' are considered candidate misannotated lncRNAs; it is interesting to note that most of the lncRNAs have a 'Good' FLOSS score. In contrast, fewer lncRNAs are considered misannotated according to ORFScore and PhyloP. The overlap between these 3 methods to find sORFs from ribo-seq data is shown in Figure S3.

It is important to note that the dataset used in the current work is much smaller and contains fewer lncRNAs than those found from the (Olexiouk et al., 2018) ribo-seq dataset. In order to be able to compare the numbers of misannotated lncRNAs found by the different methods, we first generated a list of lncRNAs that were present both in the ribo-seq dataset and in the nucleotide sequence dataset used for training deep learning models. From this common lncRNAs master list, we calculated the overlap between misannotated lncRNAs found by different methods. Figure 3.4 shows that the overlap between our method and FLOSS (hypergeometric test, p-value ≈ 0), ORFScore (hypergeometric test, p-value 4*e*-320) and PhyloP (hypergeometric test, p-value 1*e*-32) significant. This shows that our method successfully identifies misannotated lncRNAs by learning relevant features from the lncRNA nucleotide sequences.

3.6 Candidate coding lncRNAs discovered using training dynamics

Figure 3.5 shows some candidate coding lncRNAs discovered using the models trained here, together with evidence for coding potential aggregated from other sources. Ribo-seq p-value shows the confidence that the given lncRNA contains a true positive sORF as discovered from ribosome profiling (Olexiouk et al., 2018). PeptideShaker is a tool to analyze publicly available mass-spectrometry data. The confidence score, PeptideShaker score, for the peptide generated from each lncRNA sequence is also provided by (Olexiouk et al., 2018). BLAST Hit shows the most significant top hit for each lncRNA (Gish & States, 1993). CPC2 (Kang et al., 2017), CPAT (Wang et al., 2013) and RNASamba (Camargo et al., 2020) are coding potential prediction tools. For each candidate lncRNA, the coding potential predicted by these tools is shown as well. Finally, Transformer, LSTM and CNN refer to the coding potential predictions averaged over all epochs, as predicted by models trained in this work.

All lncRNAs shown here have a high confidence sORF discovered by ribosomeprofiling, a high confidence PeptideShaker score for the peptide produced from that ORF and a significant BLAST hit. While some of the candidates like Figure 3.5a and Figure 3.5f show high coding potential prediction from most models, others like Figure 3.5d and 3.5e are only discovered by the training dynamics models trained in this work.

3.7 Misannotated lncRNAs exist in a continuous cluster with coding

RNAs

To analyze the coding and noncoding transcript distributions of the data, we calculated features on for all RNAs in the dataset, based on properties of the transcripts as in (Tong & Liu, 2019). These features include ORF length, ORF quality, nucleotide distribution, translated peptide stability etc. used by (Tong & Liu, 2019) (see Table ?? for more details). Using these features, we apply T-distributed stochastic neighbor embedding (t-SNE) (Maaten & Hinton, 2008) (SciKit-learn implementation (Pedregosa et al., 2011), perplexity=150, iterations=1000, learning rate=200) to reveal RNA clusters.

Figure 3.6 shows the clusters obtained by performing t-SNE (Maaten & Hinton, 2008) on these features generated from RNA sequences. The labels of the RNAs (coding, non-coding) are not used while generating the clusters. However, based on available coding and non-coding ground-truth labels, along with the biotype information for the ncRNAs, we label each individual RNA example. LncRNAs determined as misannotations by the different deep learning models are labeled in black; interestingly, putative misannotated lncRNAs lie in a cluster contiguous with coding RNAs. This suggests that there is indeed some continuity between coding and lncRNAs in this embedded space and that the categories might not be as mutually exclusive as we believe, which is consistent with recent research discovering that some lncRNAs encode micropeptides (Hartford & Lal, 2020). In support of this, there are clusters of non-coding RNAs (labelled Misc RNA) that are well separated from coding RNAs and that do not contain many putative misannotated lncRNAs.

3.8 Exploring features learnt by models

To understand which which regions of the sequence are useful for making classification decision, we visualize the activation weights of the model layers. These activation weights determine which sequence features are paid most attention to by the model. Figure 3.7 shows an example attention map of a misannotated lncRNA generated from the first convolutional layer of the CNN model. Supplementary Figure 1 shows the attention weights visualized for a coding and long non-coding RNA that are not misannotated according to the criteria described above. The CNN model appears to focus on continuous stretches of adenines in the sequence to make decisions about whether a given RNA is coding or long non-coding. This might be because the poly-adenylation sites are one of the major distinguishing features between coding and non-coding RNAs. Supplementary Figure ?? shows the average attention given to all codons for this sequence. Codons with high 'adenine' content have higher average attention, but codons ending with 'TA' like 'ATA', 'CTA', 'GTA' & 'TTA' also have high average attention. Studying these and comparing the average attention differences in codons between coding and non-coding RNAs might prove interesting.



(b)

Figure 3.1 Analysis of the training dynamics of deep learning models. (a) Estimated coding probability across all training epochs shown for five RNAs. We expect coding and non-coding RNAs to have high and low coding probabilities respectively; this is the case for examples A and E. In this work, we are trying to identify non-coding RNAs that might be misannotated, when they are in fact coding. We are interested in lncRNAs- like B (ENST00000447563)- which have consistently high estimated coding probability, despite having the ground truth-label 'Non-coding'. These are candidate coding lncRNA. In support of this observation, ENST00000447563 (also known as linc00689) was recently found to be protein coding Hartford & Lal (2020). C and D show examples of ambiguous samples, i.e. they show a large change in estimated coding probabilities as model training progresses, so we are not sure whether they are mislabelled or not. (b) Mean (y-axis) and standard deviation (x-axis) of ground truth class probability predictions across all training epochs can be used to determine mislabelled samples. Candidate misannotated RNAs are those in the Hard-to-Learn region i.e. RNAs with low mean and standard deviation for the ground truth class probability.



Figure 3.2 Training dynamics of deep learning models can be used to identify misannotated lncRNAs. To prove that training dynamics can be used to identify mislabelled RNAs, we sampled 5% of the data with $\mu_{ground\ truth}$ and $\sigma_{ground\ truth}$ less than 0.8 and 0.2 respectively. These are samples for which we have medium-high confidence that the ground truth labels are correct (before). After we flip the labels (coding RNAs become, non-coding RNAs and vice-versa), the samples move into the hard-to-classify region of the training dynamics summary plot.



Figure 3.3 Comparison with cncRNAdb (a) Comparison of the hard-to-learn ncRNAs obtained from by CNN, LSTM and Transformer models' training dynamics. (b) Comparison of the number of common RNAs that are present in the deep learning dataset and in cncRNAdb. (c) Intersection of the misannotated lncRNAs identified in the union of LSTM, CNN, and Transformer set with the common RNAs found in cncRNAdb. (d)Intersection of the misannotated lncRNAs identified in the intersection of at least 2 of LSTM, CNN, and Transformer set with the common RNAs found in cncRNAdb. (e) Intersection of the misannotated lncRNAs identified in the intersection of LSTM, CNN, and Transformer set with the common RNAs found in cncRNAdb. (e) Intersection of the misannotated lncRNAs identified in the intersection of LSTM, CNN, and Transformer set with the common RNAs found in cncRNAdb. (e) Intersection of the misannotated lncRNAs identified in the intersection of LSTM, CNN, and Transformer set with the common RNAs found in cncRNAdb. (e) Intersection of the misannotated lncRNAs identified in the intersection of LSTM, CNN, and Transformer set with the common RNAs found in cncRNAdb. (e) Intersection of the misannotated lncRNAs identified in the intersection of LSTM, CNN, and Transformer set with the common RNAs found in cncRNAdb.



Figure 3.4 Comparison to ribo-seq data based methods used to find misannotated lncRNAs from : (a) FLOSS (p-value ≈ 0), (b) ORFScore (p-value 4e - 320) and (c) PhyloP (p-value (1e - 32)) for the dataset from (Elkon et al., 2015). Background set has 26857 lncRNAs.



Figure 3.5 Some of the candidate coding lncRNAs discovered in this work, with evidence for coding potential aggregated from other sources. Riboseq identifies RNAs associated with ribosomes, which are likely to be translated. Ribo-seq p-values (a combined score from FLOSS, ORFScore and PhyloP) show the likelihood of the identified RNA being a true-positive. PeptideShaker analyzes publicly available MS data and provides a confidence score for each peptide. The BLAST Hit is the top hit from running the query in BLASTx. CPC2, CPAT and RNASamba are tools for the coding potential prediction of a given RNA.)



Figure 3.6 Misannotated lncRNAs exist in a continuous cluster with coding RNAs. t-SNE clusters obtained from hand-crafted features (Tong & Liu, 2019) generated from nucleotide sequences. Labels (Coding, lncRNA etc.) are only used for visualizing the clusters, not for generating the clusters. Putative misannotated lncRNAs lie in a cluster contiguous with coding RNAs. There are other clusters of ncRNAs that are well separated from coding RNAs and that do not contain any putative misannotated lncRNAs.



Figure 3.7 Attention maps explain which parts of sequence are important for making a classification decision. An example attention map extracted from first convolutional layer of CNN model for ENST00000447563 (a misannotated lncRNA). The ground-truth for this RNA is 'Non-coding'. However, CNN model predicts that this is a non-coding RNA with probability 0.12. Attention visualization shows regions with contiguous 'A' nucleotides have high activation weights.

4. DISCUSSION

In this work, we apply the general framework described by (Swayamdipta et al., 2020) for detecting mislabelled samples in a training dataset to detect misannotated lncRNAs. The training dataset, comprising of nucleotide sequences of coding and non-coding RNA, is used to train CNN, LSTM and Transformer models. At the end of each training epoch, coding and non-coding prediction probabilities for every RNA sample are saved. Calculating the mean and standard deviation of the ground-truth class helps determine whether a given RNA is possibly mislabelled. LncRNAs with low mean and standard deviation for the non-coding class are the possible misannotations.

A large number of misannotated lncRNAs are identified by all three different deep learning methods. This is significant since the algorithms to distinguish between coding and non-coding RNAs employed by the models are different. Moreover, when we compare the misannotations discovered here to previous methods to detect misannotated lncRNAs from ribo-seq data and manually curated dataset, we see a large overlap for two of the methods, suggesting that our method is successfully able to detect misannotated lncRNAs. It is also interesting to note that our method shows a high overlap with a manually curated list of misannotated lncRNAs. Therefore, we conclude that this approach offers promising potential for use in curating datasets used for training coding potential predictors and assisting experimental efforts in characterization of misannotated lncRNAs.

To our knowledge, this work also represents the first instance in which nucleotide sequence embeddings and transformer models are applied to the problem of building coding potential predictive models. Using nucleotide embeddings might be preferable to other representations like one-hot encoding (Hill et al., 2018) or integer encoding (Camargo et al., 2020) used previously. This is because embeddings are learnt from the complete human genome and incorporate the context in which a given codon is found in the DNA. All models are configured to have trainable embeddings; this helps to learn better representations of codons in RNA, since the original embeddings are learnt from DNA sequences. Future work to compare the original embeddings to the embeddings generated from models trained here might provide valuable insight into the differences between codons in DNA and RNA.

One limitation of the approach presented here is that it is computationally intensive since models need to be retrained on the complete dataset after evaluation of the test set performance. However, this approach represents the first method that can find possibly misannotated lncRNAs from the nucleotide sequence alone. In conjugation with ribo-seq data, it can be used to identify misannotated lncRNAs with high confidence. Moreover, it can be used for curating the training datasets used for training coding potential predictors. Future work that compares the misannotated lncRNAs obtained from models here with ribo-seq datasets from different cell-types will provide interesting results on the cell-line specificity of misannotated lncRNAs.

In the future, conducting the following analysis will provide interesting insights. First, a meta analysis of attention maps like the one shown in Figure 3.7 to compare different groups of RNAs (coding, non-coding, misannotated) to see if there are any significant differences. One approach is to calculate the distribution of average attention paid to every 3-mer in the sequence, and then to make a correspondence to the amino acid encoded. Second, since it is widely known that disease-related variants are located in the non-coding regions of the human genome, it might be interesting to look for genetic variants present in the misannotated lncRNAs identified here. Third, integrating conservation information from different species together with the analysis here might provide insights into the evolution of coding and non-coding RNA.

BIBLIOGRAPHY

- Andrews, S. J. & Rothnagel, J. A. (2014). Emerging evidence for functional peptides encoded by short open reading frames. *Nature Reviews Genetics*, 15(3), 193– 204.
- Baek, J., Lee, B., Kwon, S., & Yoon, S. (2018). Lncrnanet: long non-coding rna identification using deep learning. *Bioinformatics*, 34 (22), 3889–3897.
- Batista, P. J. & Chang, H. Y. (2013). Long noncoding rnas: cellular address codes in development and disease. *Cell*, 152(6), 1298–1307.
- Bazzini, A. A., Johnstone, T. G., Christiano, R., Mackowiak, S. D., Obermayer, B., Fleming, E. S., Vejnar, C. E., Lee, M. T., Rajewsky, N., Walther, T. C., et al. (2014). Identification of small orfs in vertebrates using ribosome footprinting and evolutionary conservation. *The EMBO journal*, 33(9), 981–993.
- Brar, G. A. & Weissman, J. S. (2015). Ribosome profiling reveals the what, when, where and how of protein synthesis. *Nature reviews Molecular cell biology*, 16(11), 651–664.
- Camargo, A. P., Sourkov, V., Pereira, G. A. G., & Carazzolle, M. F. (2020). Rnasamba: neural network-based assessment of the protein-coding potential of rna sequences. NAR Genomics and Bioinformatics, 2(1), lqz024.
- Choi, S.-W., Kim, H.-W., & Nam, J.-W. (2019). The small peptide world in long noncoding rnas. Briefings in bioinformatics, 20(5), 1853–1864.
- Chollet, F. et al. (2015). Keras. https://keras.io.
- Couso, J.-P. & Patraquim, P. (2017). Classification and function of small open reading frames. Nature reviews Molecular cell biology, 18(9), 575.
- Derrien, T., Johnson, R., Bussotti, G., Tanzer, A., Djebali, S., Tilgner, H., Guernec, G., Martin, D., Merkel, A., Knowles, D. G., et al. (2012). The gencode v7 catalog of human long noncoding rnas: analysis of their gene structure, evolution, and expression. *Genome research*, 22(9), 1775–1789.
- Djebali, S., Davis, C. A., Merkel, A., Dobin, A., Lassmann, T., Mortazavi, A., Tanzer, A., Lagarde, J., Lin, W., Schlesinger, F., et al. (2012). Landscape of transcription in human cells. *Nature*, 489(7414), 101–108.
- Elkon, R., Loayza-Puch, F., Korkmaz, G., Lopes, R., Van Breugel, P. C., Bleijerveld, O. B., Altelaar, A. M., Wolf, E., Lorenzin, F., Eilers, M., et al. (2015). Myc coordinates transcription and translation to enhance transformation and suppress invasiveness. *EMBO reports*, 16(12), 1723–1736.
- Gish, W. & States, D. J. (1993). Identification of protein coding regions by database similarity search. Nature genetics, 3(3), 266–272.
- Hartford, C. C. R. & Lal, A. (2020). When long noncoding becomes protein coding. Molecular and Cellular Biology, 40(6).
- Hill, S. T., Kuintzle, R., Teegarden, A., Merrill III, E., Danaee, P., & Hendrix, D. A. (2018). A deep recurrent neural network discovers complex biological rules to decipher rna protein-coding potential. *Nucleic acids research*, 46(16), 8105–8113.
- Hochreiter, S. & Schmidhuber, J. (1997). Long short-term memory. Neural computation, 9(8), 1735–1780.
- Huang, Y., Wang, J., Zhao, Y., Wang, H., Liu, T., Li, Y., Cui, T., Li, W., Feng, Y.,

Luo, J., et al. (2020). cncrnadb: a manually curated resource of experimentally supported rnas with both protein-coding and noncoding function. *Nucleic Acids Research*.

- Ingolia, N. T. (2014). Ribosome profiling: new views of translation, from single codons to genome scale. Nature Reviews Genetics, 15(3), 205–213.
- Ingolia, N. T., Ghaemmaghami, S., Newman, J. R., & Weissman, J. S. (2009). Genome-wide analysis in vivo of translation with nucleotide resolution using ribosome profiling. *science*, 324 (5924), 218–223.
- Ingolia, N. T., Lareau, L. F., & Weissman, J. S. (2011). Ribosome profiling of mouse embryonic stem cells reveals the complexity and dynamics of mammalian proteomes. *Cell*, 147(4), 789–802.
- Ji, Z., Song, R., Regev, A., & Struhl, K. (2015). Many lncrnas, 5'utrs, and pseudogenes are translated and some are likely to express functional proteins. *elife*, 4, e08890.
- Kang, Y.-J., Yang, D.-C., Kong, L., Hou, M., Meng, Y.-Q., Wei, L., & Gao, G. (2017). Cpc2: a fast and accurate coding potential calculator based on sequence intrinsic features. *Nucleic acids research*, 45 (W1), W12–W16.
- Kingma, D. P. & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.
- Kong, L., Zhang, Y., Ye, Z.-Q., Liu, X.-Q., Zhao, S.-Q., Wei, L., & Gao, G. (2007). Cpc: assess the protein-coding potential of transcripts using sequence features and support vector machine. *Nucleic acids research*, 35(suppl_2), W345– W349.
- LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., & Jackel, L. D. (1989). Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4), 541–551.
- Li, L., Jamieson, K., DeSalvo, G., Rostamizadeh, A., & Talwalkar, A. (2017). Hyperband: A novel bandit-based approach to hyperparameter optimization. *The Journal of Machine Learning Research*, 18(1), 6765–6816.
- Lu, S., Zhang, J., Lian, X., Sun, L., Meng, K., Chen, Y., Sun, Z., Yin, X., Li, Y., Zhao, J., et al. (2019). A hidden human proteome encoded by 'noncoding'genes. *Nucleic acids research*, 47(15), 8111–8125.
- Maaten, L. v. d. & Hinton, G. (2008). Visualizing data using t-sne. Journal of machine learning research, 9(Nov), 2579–2605.
- Matsumoto, A., Pasut, A., Matsumoto, M., Yamashita, R., Fung, J., Monteleone, E., Saghatelian, A., Nakayama, K. I., Clohessy, J. G., & Pandolfi, P. P. (2017). mtorc1 and muscle regeneration are regulated by the linc00961-encoded spar polypeptide. *Nature*, 541(7636), 228–232.
- Miller, W., Rosenbloom, K., Hardison, R. C., Hou, M., Taylor, J., Raney, B., Burhans, R., King, D. C., Baertsch, R., Blankenberg, D., et al. (2007). 28way vertebrate alignment and conservation track in the ucsc genome browser. *Genome research*, 17(12), 1797–1808.
- Nelson, B. R., Makarewich, C. A., Anderson, D. M., Winders, B. R., Troupes, C. D., Wu, F., Reese, A. L., McAnally, J. R., Chen, X., Kavalali, E. T., et al. (2016). A peptide encoded by a transcript annotated as long noncoding rna enhances serca activity in muscle. *Science*, 351(6270), 271–275.
- Ng, P. (2017). dna2vec: Consistent vector representations of variable-length k-mers. arXiv preprint arXiv:1701.06279.

- Olexiouk, V., Van Criekinge, W., & Menschaert, G. (2018). An update on sorfs. org: a repository of small orfs identified by ribosome profiling. *Nucleic acids research*, 46(D1), D497–D502.
- O'Malley, T., Bursztein, E., Long, J., Chollet, F., Jin, H., Invernizzi, L., et al. (2019). Keras Tuner. https://github.com/keras-team/keras-tuner.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, E. (2011). Scikitlearn: Machine learning in Python. Journal of Machine Learning Research, 12, 2825–2830.
- Rinn, J. L. & Chang, H. Y. (2012). Genome regulation by long noncoding rnas. Annual review of biochemistry, 81, 145–166.
- Swayamdipta, S., Schwartz, R., Lourie, N., Wang, Y., Hajishirzi, H., Smith, N. A., & Choi, Y. (2020). Dataset cartography: Mapping and diagnosing datasets with training dynamics.
- Tong, X., Hong, X., Xie, J., & Liu, S. (2020). Cppred-sorf: Coding potential prediction of sorf based on non-aug. *BioRxiv*.
- Tong, X. & Liu, S. (2019). Cppred: coding potential prediction based on the global description of rna sequence. Nucleic acids research, 47(8), e43–e43.
- Ulitsky, I. & Bartel, D. P. (2013). lincrnas: genomics, evolution, and mechanisms. Cell, 154(1), 26–46.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In Advances in neural information processing systems, (pp. 5998–6008).
- Wang, L., Park, H. J., Dasari, S., Wang, S., Kocher, J.-P., & Li, W. (2013). Cpat: Coding-potential assessment tool using an alignment-free logistic regression model. *Nucleic acids research*, 41(6), e74–e74.
- Zhu, M. & Gribskov, M. (2019). Mipepid: Micropeptide identification tool using machine learning. BMC bioinformatics, 20(1), 559.