

Network Representation of Large-Scale Heterogeneous RNA Sequences with Integration of Diverse Multi-omics, Interactions, and Annotations Data

Nhat Tran and Jean Gao

*Department of Computer Science & Engineering, The University of Texas at Arlington,
Arlington, TX 76010, USA*

E-mail: gao@uta.edu

Long non-coding RNA (lncRNA), microRNA, and messenger RNA enable key regulations of various biological processes through a variety of diverse interaction mechanisms. Identifying the interactions and cross-talk between these heterogeneous RNA classes is essential in order to uncover the functional role of individual RNA transcripts, especially for unannotated and sparsely discovered RNA sequences with no known interactions. Recently, sequence-based deep learning and network embedding methods are gaining traction as high-performing and flexible approaches that can either predict RNA-RNA interactions from sequence or infer missing interactions from patterns that may exist in the network topology. However, most of the current methods have several limitations, e.g., the inability to perform inductive predictions, to distinguish the directionality of interactions, or to integrate various sequence, interaction, expression, and genomic annotation datasets. We proposed a novel deep learning framework, rna2rna, which learns from RNA sequences to produce a low-dimensional embedding that preserves proximities in both the interaction topology and the functional affinity topology. In this proposed embedding space, the two-part “source and target contexts” capture the receptive fields of each RNA transcript to encapsulate heterogeneous cross-talk interactions between lncRNAs and microRNAs. The proximity between RNAs in this embedding space also uncovers the second-order relationships that allow for accurate inference of novel directed interactions or functional similarities between any two RNA sequences. In a prospective evaluation, our method exhibits superior performance compared to state-of-art approaches at predicting missing interactions from several RNA-RNA interaction databases. Additional results suggest that our proposed framework can capture a manifold for heterogeneous RNA sequences to discover novel functional annotations.

Keywords: network embedding; lncRNA; functional similarity; interactions; siamese network.

1. Introduction

Regulatory long non-coding RNAs (lncRNAs) and microRNAs (miRNAs) that influences gene expression post-transcriptionally by interacting to target messenger RNAs (mRNA) form a complex network of transcriptomic interactions. These heterogeneous families of non-coding RNAs are associated with nearly all cellular processes, and recently, lncRNAs are gaining considerable attention as the largest and most diverse class of non-coding RNA, encompassing nearly 30,000 discovered transcripts in human. Among many of their known functional interaction mechanisms, lncRNAs are known to act as miRNA decoys, derepress

gene expression by competing with miRNAs for shared mRNA targets, or directly regulate gene expression.¹ Determining the biological functions of the individual lncRNAs remains a challenge as most of these RNA transcripts are currently unannotated, and their known interactions are sparse. Recent advances in RNA sequencing (RNA-Seq), deep sequencing (CLIP-seq, LIGR-Seq), and computational methods allow for an unprecedented analysis of such transcripts and have enabled researchers to generate large-scale interaction and functional annotation databases. However, the interaction networks generated from such data are often scant and incomplete in the number of lncRNAs covered. Furthermore, although a large number of lncRNAs have been identified, only a few hundreds have had functional and molecular mechanisms determined to date, as observed in lncRNAdb.² In other avenues, a growing number of lncRNAs are being assigned biological functions based on their cell-specific expressions³ or disease associations.⁴ Still, a vast majority of other lncRNAs only has basic genomic information such as locus biotype and transcript sequence assigned. These transcripts might support important biological cell functions and could potentially serve as targets for genomic, diagnostic, or therapeutic studies. Thus, in the effort to functionally characterize these hypothetical lncRNAs, the essential tasks are integrating the various -omics and annotation attributes between the known RNAs, while inferring the functions and interactions to the novel RNAs directly from sequence. In order to address these challenges, this paper proposes a “network embedding” method, which aims to provide a network representation of existing knowledge in the interaction topology and attributes of transcriptome-wide RNAs, all while predicting novel functional interactions.

Many graph-theoretic methods have been applied to biological networks with the intuition that RNAs close together in the interaction topology are more likely to be involved in many of the same functions.⁵ Due to the extreme sparsity of the known interaction network among lncRNAs, it is pertinent to unravel the functional association between lncRNAs by considering its gene/transcript genomic annotations, tissue-specific expressions, gene-disease associations, and sequence similarity.⁴ In recent studies, Kishan *et al.*⁶ uncovered the second-order proximity relationship between interacting genes by integrating the gene regulatory network and gene expression as side information. On the other hand, several structure-free sequence-based methods⁷⁻⁹ have also been proposed for the prediction of protein binding sites, family classification, structure prediction, or interaction prediction from sequence. Motivated by these works, we propose an integrative method to extract a functional representation from RNA sequences in order to accurately predict new interactions and functional annotations, while simultaneously integrate various existing biological multi-omics and annotation databases.

2. Methods

2.1. Defining the Heterogeneous lncRNA-miRNA-mRNA Network

We formally define the heterogeneous network of lncRNA, miRNA, and mRNA interactions and functional similarity as two networks of directed and undirected edges. We denote the two networks $G_1(V, E^d)$ and $G_2(V, E^u)$ having the same set of nodes V and two sets of edges E^d and E^u . The set of nodes $V = \{v_1, \dots, v_n\}$ can also be expressed as $V = \{L, M, N\}$ s.t. $|L| + |M| + |N| = n$, where L, M, N are the sets containing the lncRNA, microRNA, and mRNA heterogeneous nodes, respectively. The set of directed edges $E^d = \{e_{ij}^d\}_{i,j=1}^n$ represent

directed regulatory interactions that each specify a source and a target. The undirected edges $E^u = \{e_{ij}^u\}_{i,j=1}^n$ s.t. $e_{ji}^u = e_{ij}^u$ represent the functional affinities among the heterogeneous RNA nodes. Each edge e_{ij} is associated with a weight s.t. $0 \leq e_{ij} \leq 1$, indicating the strength of the connection between RNA i and RNA j . If $e_{ij} > 0$, we consider the edge a positive interaction/affinity, and if $e_{ij} = 0$, we consider the edge a negative (non) interaction/affinity. In this paper, we mainly consider the edge weights e_{ij} to be binary.

2.2. Materials and Data

The directed edges represent the directed regulatory interactions between lncRNAs, miRNAs, and mRNAs. We consider the directionality in the regulatory interactions by encoding directed edges, i.e., encoding miRNA-lncRNA interactions (e.g., miRNA inducing lncRNA decay) to be separate from lncRNA-miRNA interactions (e.g., lncRNAs acting as miRNA decoys). In this study, the various types of interaction collected from various experimentally-verified databases in the lncRNA-miRNA-mRNA interactome considered are:

- **microRNA-mRNA** post-transcriptional interactions from miRTarBase¹⁰ v7.0 containing a total of 442,067 interactions matched between 1,630 microRNAs and 14,666 mRNAs.
- **microRNA-lncRNA** interactions from DIANA-lncBase Experimental v2¹¹ containing a total of 53,926 matched interactions between 631 miRNAs and 2530 lncRNAs.
- **ncRNA-RNA** regulatory interactions from NPInter v3.0,¹² where we filtered only lncRNA-miRNA, lncRNA-mRNA, and miRNA-lncRNA interactions, which resulted in 208,841 interactions between 669 matching lncRNAs and 7,369 mRNAs.
- **lncRNA-mRNA** post-transcriptional regulatory interactions from lncRNA2Target v2.0¹³ which contained a total of 65,655 interactions between 1037 lncRNAs and 28,866 mRNAs.
- **mRNA-mRNA** gene regulatory interactions obtained from the BioGRID v3.5 database,¹⁴ which included more than 347,246 matched interactions among 19,607 mRNAs.

These heterogeneous interactions are combined into an integrated network, and the associated set of edges is E^d , where the binary edge weight $e_{ij}^d \in \{0, 1\}$ indicates whether a regulatory interaction from RNA node v_i to RNA node v_j has been observed in the literature.

We also outline the various multi-omics, sequence, and annotation databases utilized to provide functional attributes to individual RNAs. In order to maximize the number of transcripts and genes matched between different annotation systems, the RNA transcripts are indexed by standard gene symbols.

- **Tissue-specific median expressions** obtained from the GTEx Portal,¹⁵ which provides tissue-specific RNA-Seq expressions for 15,598 non-diseased samples of various human tissue types. From GTEx Analysis v7, median TPM values by tissue were calculated for 18653 mRNA genes, 1181 miRNA and 12706 lncRNA transcripts over 44 tissue types.
- **Genomic annotations** obtained from the GENCODE Release 29¹⁶ which contains the transcript biotype annotation and genomic location for lncRNAs, miRNAs, and mRNAs.
- **Transcript sequences** comprised of genome-wide human lncRNA and mRNA primary sequences from GENCODE Release 29¹⁶ and miRNA hairpin sequences from miRBase.¹⁷

- **Functional annotations** for mRNA were obtained from the HUGO Gene Names database.¹⁸ In addition, GO terms for 162 matched lncRNAs were obtained from RNAcentral¹⁹ which aggregated data from NONCODE²⁰ and lncipedia.²¹ For miRNAs, family classifications obtained from the TargetScan Release 7.2 (March 2018),²² RNA structure family annotation from Rfam 13.0,²³ and GO terms for 487 matched miRNAs from RNAcentral were included.
- **Disease associations** from DisGeNet²⁴ covered 7577 mRNA genes. The HMDD miRNA-disease²⁵ database includes 553 miRNAs, and LncRNADisease²⁶ includes 150 lncRNAs.

Each of the RNA functional attribute data types, excluding transcript sequences, is encoded to a data matrix, $A^k \in \mathbb{R}^{|V| \times m}$ for the annotation type k , which can either be real-valued or categorical. If an attribute k is categorically-valued, A^k is an indicator matrix where entries are binary-valued, and m is the number of all unique values. Considering that functional annotations for the non-coding RNAs are sparse, each node can have up to K attribute types associated with it.

2.3. Undirected RNA-RNA Functional Affinity Edges

We aim to capture the functional similarity between two RNA nodes of the same class by calculating as a similarity measure, suggesting a resemblance in RNA function or structure. For any categorical text-valued annotation (e.g., disease association, transcript biotype, RNA structure family, or GO terms) that two RNA nodes v_i and v_j both have been annotated, the attributes in this annotation are first transformed to binary feature representation. For a categorical annotation denoted as k , the binary 1-D feature vector associated with RNA node v_i is denoted as a_i^k . Using the Sørensen-Dice coefficient score,²⁷ a similarity score between two binary vectors for node v_i and node v_j for annotation k can be obtained by $s_{ij}^k = 2(a_i^k \cdot a_j^k) / (2(a_i^k \cdot a_j^k) + |a_i^k|_1 + |a_j^k|_1)$. This similarity measure ranges $[0, 1]$ and gives higher weight to the common attributes present in both RNAs than by the attributes present in only one RNA. Since most RNAs have null annotations, the Dice coefficient score is only computed between pairs of RNA nodes that have both been annotated, while other pairs remain a null value. For the real-valued attribute data, i.e., tissue-specific expressions, we instead calculate the absolute correlation value between the two expression vectors a_i^k and a_j^k to obtain s_{ij}^k .

To obtain an aggregate affinity score between a pair of RNA nodes across all K similarity values, we utilize a modified version of the Gower's similarity score.²⁸ For each RNA-RNA pair, Gower's similarity aggregates the similarity scores across all the annotation types and performs a weighted average. Typically, a similarity score for a pair of RNA nodes that do not have any associated annotation would be considered a 0, but in this study, we remove these null pairwise similarity from consideration. Thus, Gower's similarity will only aggregate the pairwise similarity scores from annotations that exist in both nodes to compute the average.

After computing for all pairs of RNA nodes, the resulting pairwise affinity matrix is \mathbf{S} where entries $S_{ij} = \sum_k^{K_{ij}} s_{ij}^k / |K_{ij}|$ with K_{ij} being the set of attribute types present in both RNA nodes v_i and v_j . The entries S_{ij} will then be selected as edge weights e_{ij}^u that represent the functional similarity edges between nodes. Since our model currently only considers unweighted binary edges, we selected undirected edges to have an affinity score close to 1.0, or higher than a chosen hard-threshold, to be considered as a positive edge. In our experiments, the

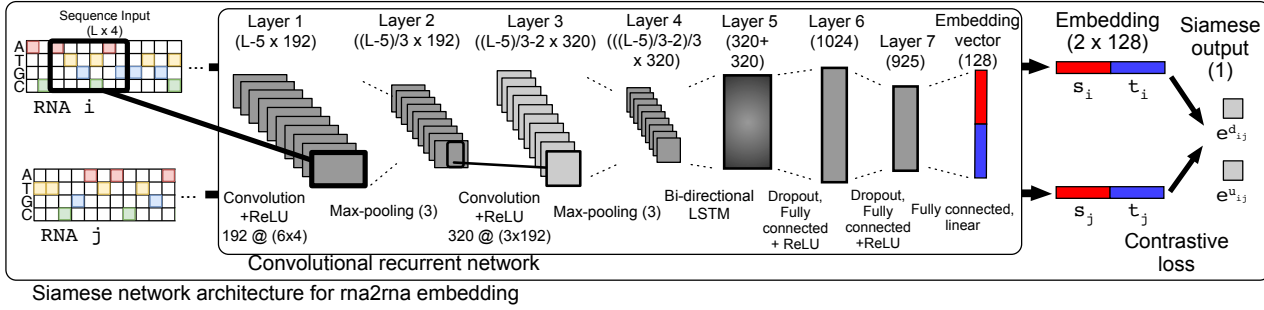


Fig. 1. The rna2na network embedding method utilizing Siamese architecture.

hard-threshold was arbitrarily chosen where the number of positive affinity edges covers no more than 0.1% sparsity of the entire affinity matrix. We also utilize the set of undirected affinity edges with 0.0 weight, indicating a negative edge that suggests functional dissimilarity.

3. Network Embedding with Source-Target Contexts

A network embedding is mapping from each RNA node to a low-dimensional representation, denoted as a mapping function $f : v_i \rightarrow y_i \in \mathbb{R}^d, \forall v_i \in V$, where d is the embedding dimensionality *s.t.* $d \ll n$. The embedding y_i associated with each node v_i is learned such that nodes preserve some meaningful proximities to each other according to the given topology in both networks G_d and G_u . We propose the embedding to have two components: source context and target context. That is, each embedding vector $y_i = [s_i, t_i]^T$ is represented as a concatenation of the source context, $s_i \in \mathbb{R}^{d/2}$, and target context, $t_i \in \mathbb{R}^{d/2}$. This embedding representation can simultaneously capture directed and undirected edges by our proposed first-order directed proximity and second-order undirected proximity. The first-order directed proximity represents a binding affinity from node i 's source context to node j 's target context, defined as

$$d_1(v_i, v_j) = \sqrt{(s_i - t_j)^2} \quad (1)$$

Note, $d_1(v_i, v_j)$ can take on a different value than $d_1(v_j, v_i)$. The second-order undirected proximity represents the functional similarity between node i and node j , defined as

$$d_2(v_i, v_j) = \sqrt{(y_i - y_j)^2} \quad (2)$$

The value of these proximities is an Euclidean distance, where the embeddings are selected so that if two nodes have a positive (directed or undirected) edge, its respective embeddings should be more similar, i.e., having a smaller distance. Otherwise, if two nodes have a negative or non-interactions, their embeddings should be more dissimilar, incurring a greater distance.

3.1. Representation Learning for RNA Sequences to Reconstruct the Interactions and Functional Topology

Aside from the interaction topology data, each RNA v_i also has an associated transcript sequence in one-hot vector representation denoted by $x_i \in \{1, 2, 3, 4\}^{l_i}$, where l_i is the variable length of an RNA sequence. We propose the network embedding function $f : x_i \rightarrow y_i \in \mathbb{R}^d$ to

be a deep Siamese neural network that maps an RNA sequence input x_i to an embedding y_i of a fixed dimension $d = 128$, outlined in Fig. 1. Initially proposed for signature verification,²⁹ the Siamese network is an architecture where a pair of objects can each be encoded such that its resulting embeddings can determine if the two objects are similar or dissimilar. More specifically, the network learns to output embeddings for a pair of RNA sequences, guided by edge weight e_{ij}^u as the label indicating whether the pair is functionally similar or dissimilar. Additionally, for an interacting pair of RNAs, the directed edge e_{ij}^d would indicate whether RNA i interacts with RNA j . For the output embeddings to preserve proximities across all edges in both G_d and G_u network topologies, we minimize the binary cross-entropy losses:³⁰

$$L_1(X, E^d, f) = \sum_{e_{ij}^d \in E^d} e_{ij}^d \log(d_1(f(x_i), f(x_j))) + (1 - e_{ij}^d) \log(1 - d_1(f(x_i), f(x_j))) \quad (3)$$

$$L_2(X, E^u, f) = \sum_{e_{ij}^u \in E^u} e_{ij}^u \log(d_2(f(x_i), f(x_j))) + (1 - e_{ij}^u) \log(1 - d_2(f(x_i), f(x_j))) \quad (4)$$

The network weights in $f(x)$ are trained with Stochastic Gradient Descent using the standard back-propagation algorithm. At each iteration, a batch of RNA nodes is sampled along with its associated sets of positive and negative, directed and undirected edges. This batch optimization strategy addresses the memory requirement to encode a large number of variable-length RNA sequences while efficiently sample the associated interaction edges. Additionally, it also circumvents the issue of size imbalance between the interaction sets of well-studied RNA classes and newly-emerged RNA classes. We implement a biased random node sampling strategy where we randomly select a set of nodes and fit the network on the set of edges induced by this sub-graph. The probability of selecting a node i is a function of its degree r_i , as a square-root sampling compression function³¹ $P(v_i) = \sqrt{r_i} / \sum_{v_j \in V} \sqrt{r_j}$, which aims to retain the power-law degree distribution while keeping the frequency ranking of each node. A batch of nodes S is sampled without replacement from this distribution, and each node has a set of positive edges. To obtain the negative directed edges, we then sample a number of nodes in S by adopting the negative sampling approach proposed in,³² where the ratio of negative edges to positive edges incident to each node is between 2.0 and 5.0. Given S and the set of directed E_S^d and undirected E_S^u edges containing both positive and negative interactions incident to S , we minimize the loss function $L(S, E_S^d, E_S^u, f) = L_1(S, E_S^d, f) + \lambda L_2(S, E_S^u, f)$ where λ is a coefficient parameter.

After training convergence, given any RNA full sequence inputs x_i and x_j , the learned model yields the embeddings y_i and y_j which is used to predict whether a relationship exists between them. We use the proximity $d_1(v_i, v_j)$ or $d_2(v_i, v_j)$ to either predict the existence of interaction or functional similarity, depending on the query. Finally, a Gaussian kernel function $P(v_i, v_j) = \exp(-\gamma * d(v_i, v_j)^2)$ is applied to yield a predicted affinity score.

4. Results

4.1. Large-scale Integration of RNA Sequences, Interactions, Expressions, and Annotations

We integrated various experimentally verified interaction databases to build a large-scale lncRNA-miRNA-mRNA interaction network. Additionally, various functional annotations,

Table 1. Overview of the data selection, harmonization, and integration for training and validation interaction databases in prospective evaluation.

Interaction database	Training Sets			
	Version	# interactions	# source nodes	# target nodes
miRTarBase	6.0	377,318	1,618 miRNAs	14,666 mRNAs
DIANA-lncBase	v2	53,926	631 miRNAs	2530 lncRNAs
NPInter	v2.0	85,335	12 lncRNAs	5023 mRNAs
lncRNA2Target	v1.0	1308	79 lncRNAs	471 mRNAs
BioGRID	v3.4	313,724	13,318 mRNAs	19,429 mRNAs
Interaction database	Validation Sets			
	Version	# interactions	# novel sources	# novel targets
miRTarBase	7.0	64,749	12 miRNAs	702 mRNAs
DIANA-lncBase	Predicted	337,031	0 miRNAs	0 lncRNAs
NPInter	v3.0	123,054	499 lncRNAs	2346 mRNAs
lncRNA2Target	v2.0	65,624	1037 lncRNAs	10,825 mRNAs
BioGRID	v3.5	33,522	178 mRNAs	178 mRNAs

expressions, sequences, and disease associations were integrated to enable the extraction of the undirected attribute affinity edges. In total, there are 12725 lncRNAs, 1870 microRNAs, and 20284 mRNAs considered in this study, comprised of comprehensive integration of the various databases outlined in Table 1. To accomplish the primary task of predicting novel interactions not seen at training time, we propose an experimental setup using prospective evaluation. All models were trained exclusively using a prior version of each interaction database. Then, we validate the link prediction model using the set of new interactions from the latest database version update. This type of evaluation, rarely done in the literature, is extremely important as it allows us to mimic a realistic scenario where the task is to discover novel RNA-RNA interactions based on the existing knowledge. The training sets are comprised of the interaction network of database versions released before 2015, while the validation sets are comprised of updates from the latest released database version. After integration of the training databases, self-interactions and redundant interaction edges are removed, and only interactions between RNAs with an associated transcript sequence is considered. In the validation set, we evaluated only the interactions that do not overlap with the training set.

After the integration of various multi-omics and annotation attributes, RNA-RNA pairwise functional affinities were computed, and undirected affinity edges were added to the undirected interactions training set. Given the affinities \mathbf{A} for all lncRNA, miRNA, and mRNA pairs and filtering second-order undirected affinities at a 0.8 threshold, 65864, 405, and 362362 undirected edges were added, respectively. With the negative sampling ratio set at 5.0 per positive edge, a total of 329320, 2025, and 724724 negative edges among the lncRNAs, miRNAs, and mRNAs were added to the undirected edges training set, respectively.

4.2. Novel Link Predictions

Our experiments included comparative analysis with existing state-of-the-art network embedding methods across various evaluation tasks. The methods considered are node2vec,³³

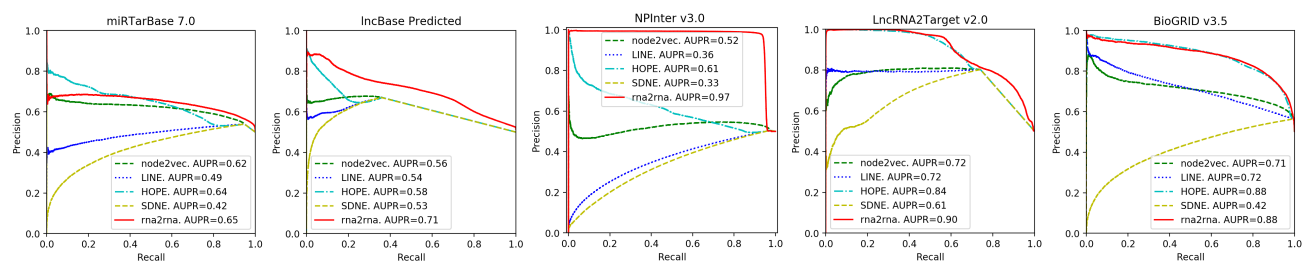


Fig. 2. Precision-recall curves for link prediction comparison analysis.

LINE,³⁴ HOPE,³⁵ and SDNE.³⁶ In the following experiments, we assessed each method by learning a 128-dimensions embedding representation from the training network. All other free parameters were set according to the default values in the proposed methods. We composed the training set for the link prediction task by a union of all ground-truth interactions set from the miRTarBase 6.0, lncBase v2, NPInter v2.0, lncRNA2Target v1.0, and BioGRID v3.4 databases. After the models have been trained, its estimated interaction adjacency matrix is computed and evaluated on the novel interactions from miRTarBase v7.0, lncBase predicted, NPInter v3.0, lncRNA2Target v2.0, and BioGRID v3.5 database separately. For a test to differentiate between positive interactions and random noise interactions, we also uniformly sample several interactions from the set of all possible pairwise interactions to consider as negative interaction using the random node sampling distribution specified in Riad *et al.*³¹ This set is denoted as E^n , and sampling size of negative interactions is such that the ratio of negative to positive interactions is 1.0. In the evaluation, the sets of ground truth validation edges E^d and random noise E^n edges are used to calculate the precision and recall rates. All methods were evaluated on the same set of positive and sampled negative interactions. The area under the precision-recall (AUPR) curves shown in Fig. 2 highlights the comparison analysis across five different interaction databases. Note that rna2rna can perform inductive link prediction to novel RNA sequences since the validation set contains many novel lncRNAs, miRNAs, and mRNAs not seen at training time. While other transductive methods were evaluated only on the interactions incident to the RNA nodes seen at training time, rna2rna were evaluated on a greater number of interactions, yet still achieved the top performance at the default parameters $\gamma = 2.0$ and $\lambda = 1.0$.

4.3. Inferring Functional Similarity From Embeddings

In comparison analysis, we first obtained the embeddings from each of the methods and performed K-Means clustering only on the nodes that have an associated functional annotation. The number of clusters in K-Means is the same as the total number of unique labels in a particular annotation. The clustering result of different methods is compared over the RNA family and transcript biotype annotations, shown in Table 2 and Table 3. The result shows that although there is a greater number of RNA nodes to assign to clusters, rna2rna embeddings can achieve the highest NMI score over the RNA functional family and biotype annotations.

Since rna2rna embeddings can preserve functional similarities, an essential next step is to assign putative biological functions to novel lncRNAs. To do this, we perform gene set

Table 2. Clustering analysis over 2343 true RNA functional family annotations.

Method	Homogeneity	Completeness	NMI	# nodes
node2vec	0.641	0.602	0.621	11735
LINE	0.689	0.614	0.650	11735
HOPE	0.525	0.571	0.570	11735
SDNE	0.613	0.588	0.600	11735
rna2rna*	0.508	0.530	0.519	14312
rna2rna	0.685	0.620	0.651	14312

Table 3. Clustering analysis over 24 true RNA transcript biotype annotations.

Method	Homogeneity	Completeness	NMI	# nodes
node2vec	0.147	0.089	0.111	23940
LINE	0.268	0.158	0.199	23940
HOPE	0.109	0.111	0.110	23940
SDNE	0.079	0.076	0.078	23940
rna2rna*	0.178	0.138	0.155	32530
rna2rna	0.355	0.235	0.283	32530

rna2rna* denotes the model trained on the directed interactions data alone, without the undirected functional affinity information.

enrichment analysis on K-mean clusters of the RNA node embeddings, select the clusters with the highest enriched functional term, then associate the lncRNAs belonging in this cluster with this term. In this experiment, K-means ($k = 2000$) is performed over 32,741 RNA embeddings trained from both training set and validation set. We then performed enrichment analysis on these 2000 clusters using Enrichr³⁷ over the KEGG Human 2019³⁸ terms, which includes both functional and disease pathways. Among the 2000 clusters, 559 have an adjusted P-value of less than 0.01, and 139 have less than 0.001. Interestingly, the highest-scoring gene sets usually contain some lncRNAs not previously associated with any other functional term, as shown in Table 4. It warrants additional experimental studies to verify the functional associations of these lncRNAs to the annotated genes within the same cluster.

4.4. Discussions and Conclusion

With the framework we have developed, heterogeneous functional attributes and interaction data are integrated to enable the characterization of RNA sequences using a network embedding representation. While the method of integrating various functional attributes is simple, its purpose is to characterize functional affinities for an extensive number of RNAs, even among sparsely annotated ones. Whereas very few lncRNAs have been fully annotated, especially in functional annotation or disease association, most have already been annotated with tissue-specific expressions, transcript biotype, and sequence. Thus, more training data was utilized so the model can recognize a greater number of RNA sequences to achieve better generalization.

Additionally, since our method was able to map the functional affinity between RNA nodes belonging in disconnected components in the interaction topology, we hypothesize rna2rna can effectively map any RNA transcript sequence to a functional manifold in the embedding representation. It is observed in the t-SNE visualization in Fig. 3 that the embeddings can preserve the local structure of the interactions and functional annotations, as well as exhibit good separation based on their transcript biotype classification. There is a clear separation between miRNAs to lncRNAs and mRNAs, and while the protein-coding mRNAs and lncRNAs may have some overlap, it is expected since the sequence structure of these two RNA classes is similar. Note that although no negative undirected edges between RNAs of different types (e.g., lncRNAs v. miRNAs) were sampled to explicitly designate different RNA types to have dissimilar embeddings, the network can still make a distinction between their functional roles.

Table 4. Gene set enrichment analysis over 2000 K-mean clusters. Each row indicates the highest enriched KEGG functional pathway for a given cluster gene set comprised of both mRNAs and lncRNAs.

Genes	lncRNAs	KEGG Term	Overlap	P-value
ZNF177,ZNF175,ZNF607,ZNF606,...	AC022150.4	Herpes simplex virus 1...	269/492	2.4e-323
OR7G2,OR8I2,OR7G1,OR9K2,OR11...	AC131571.1,...	Olfactory transduction	350/444	2.4e-323
GSK3B,HDAC2,PTGER3,PTEN,...	LINC00598	Pathways in cancer	25/530	1.13e-13
CHRND,PTGIR,EDNRB,MTNR1B,...	UCA1	Neuroactive ligand-recept...	12/338	9.67e-13
KIR2DS4,KIR2DL1,KIR3DL3,KIR2DL3	Z99756.1,...	Antigen processing...	4/77	2.62e-07
CHRNA,HTT1E,PTGER1,KISS1R,...	AL355297.4,...	Neuroactive ligand-recept...	6/338	2.35e-06
P2RY4,TAAR6,C5AR1,HTT5A,TRHR	AC008125.1	Neuroactive ligand-recept...	5/338	9.616e-06
AOC3,PAH	LINC01940,...	Phenylalanine metabolism	2/17	2.19e-04

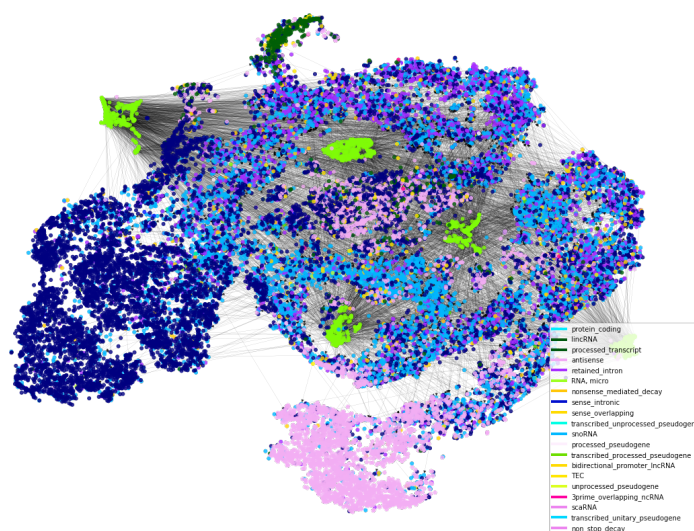


Fig. 3. Visualization of the 35,069 transcriptome-wide lncRNA-miRNA-mRNA nodes mapped to a 2-D t-SNE projection from 128-dimensional embeddings.

This shows that rna2rna's source-target embedding is an effective representation that can encode an RNA's biological function only by its given directed interactions.

Our main contribution proposes a highly versatile architecture aimed at predicting interactions between heterogeneous RNA transcripts while characterizing the functional landscape of various non-coding RNA species. The method currently has some limitations inherent in a deep learning method, such as requiring the tuning of free parameters and lacking direct biological interpretability of the embedding features. For future work, we can incorporate attention mechanisms in the encoder network to correlate individual features with specific regions in the sequence input. The notion of the source and target context embeddings would serve additional purposes to identify the binding regions of a sequence that would enable their regulatory functions. In foresight, we intend this method to be the groundwork for subsequent down-stream analysis, where the data inputs and model parameters can be fine-tuned for specific prediction tasks such as the inference of a functional regulatory pathway or the

discovery of unknown RNA families and novel disease biomarkers. Further development of this framework can provide an invaluable tool to support significant discoveries in systems biology, especially for newly identified non-coding RNAs.

References

1. J.-H. Yoon, K. Abdelmohsen and M. Gorospe, Functional interactions among micrnas and long noncoding rnas, in *Seminars in cell & developmental biology*, 2014.
2. P. P. Amaral, M. B. Clark, D. K. Gascoigne, M. E. Dinger and J. S. Mattick, Incrnadb: a reference database for long noncoding rnas, *Nucleic acids research* **39**, D146 (2010).
3. B. S. Gloss and M. E. Dinger, The specificity of long noncoding rna expression, *Biochimica et Biophysica Acta (BBA)-Gene Regulatory Mechanisms* **1859**, 16 (2016).
4. X. Chen, C. C. Yan, C. Luo, W. Ji, Y. Zhang and Q. Dai, Constructing lncrna functional similarity network based on lncrna-disease associations and disease semantic similarity, *Scientific reports* **5**, p. 11338 (2015).
5. H. N. Chua, W.-K. Sung and L. Wong, Exploiting indirect neighbours and topological weight to predict protein function from protein–protein interactions, *Bioinformatics* **22**, 1623 (2006).
6. K. Kishan, R. Li, F. Cui, Q. Yu and A. R. Haake, Gne: a deep learning framework for gene network inference by aggregating biological information, *BMC systems biology* **13**, p. 38 (2019).
7. E. Asgari and M. R. Mofrad, Continuous distributed representation of biological sequences for deep proteomics and genomics, *PLoS one* **10**, p. e0141287 (2015).
8. D. Quang and X. Xie, Danq: a hybrid convolutional and recurrent deep neural network for quantifying the function of dna sequences, *Nucleic acids research* **44**, e107 (2016).
9. X. Pan and H.-B. Shen, Learning distributed representations of rna sequences and its application for predicting rna-protein binding sites with a convolutional neural network, *Neurocomputing* **305**, 51 (2018).
10. C.-H. Chou, S. Shrestha, C.-D. Yang, N.-W. Chang, Y.-L. Lin, K.-W. Liao, W.-C. Huang, T.-H. Sun, S.-J. Tu, W.-H. Lee *et al.*, mirtarbase update 2018: a resource for experimentally validated microrna-target interactions, *Nucleic acids research* **46**, D296 (2017).
11. M. D. Paraskevopoulou, I. S. Vlachos, D. Karagkouni, G. Georgakilas, I. Kanellos, T. Vergoulis, K. Zagganas, P. Tsanakas, E. Floros, T. Dalamagas *et al.*, Diana-lncbase v2: indexing microrna targets on non-coding transcripts, *Nucleic acids research* **44**, D231 (2015).
12. Y. Hao, W. Wu, H. Li, J. Yuan, J. Luo, Y. Zhao and R. Chen, Npinter v3. 0: an upgraded database of noncoding rna-associated interactions, *Database* **2016** (2016).
13. L. Cheng, P. Wang, R. Tian, S. Wang, Q. Guo, M. Luo, W. Zhou, G. Liu, H. Jiang and Q. Jiang, Lncrna2target v2. 0: a comprehensive database for target genes of lncrnas in human and mouse, *Nucleic acids research* **47**, D140 (2018).
14. A. Chatr-Aryamontri, R. Oughtred, L. Boucher, J. Rust, C. Chang, N. K. Kolas, L. O'Donnell, S. Oster, C. Theesfeld, A. Sellam *et al.*, The biogrid interaction database: 2017 update, *Nucleic acids research* **45**, D369 (2017).
15. J. Lonsdale, J. Thomas, M. Salvatore, R. Phillips, E. Lo, S. Shad, R. Hasz, G. Walters, F. Garcia, N. Young *et al.*, The genotype-tissue expression (gtex) project, *Nature genetics* **45**, p. 580 (2013).
16. J. Harrow, A. Frankish, J. M. Gonzalez, E. Tapanari, M. Diekhans, F. Kokocinski, B. L. Aken, D. Barrell, A. Zadissa, S. Searle *et al.*, Gencode: the reference human genome annotation for the encode project, *Genome research* **22**, 1760 (2012).
17. S. Griffiths-Jones, R. J. Grocock, S. Van Dongen, A. Bateman and A. J. Enright, mirbase: microrna sequences, targets and gene nomenclature, *Nucleic acids research* **34**, D140 (2006).
18. T. A. Eyre, F. Ducluzeau, T. P. Sneddon, S. Povey, E. A. Bruford and M. J. Lush, The hugo gene nomenclature database, 2006 updates, *Nucleic acids research* **34**, D319 (2006).

19. R. Consortium, Rnacentral: a comprehensive database of non-coding rna sequences, *Nucleic acids research*, p. gkw1008 (2016).
20. D. Bu, K. Yu, S. Sun, C. Xie, G. Skogerbø, R. Miao, H. Xiao, Q. Liao, H. Luo, G. Zhao *et al.*, Noncode v3. 0: integrative annotation of long noncoding rnas, *Nucleic acids research* **40**, D210 (2011).
21. P.-J. Volders, K. Helsens, X. Wang, B. Menten, L. Martens, K. Gevaert, J. Vandesompele and P. Mestdagh, Lncipedia: a database for annotated human lncrna transcript sequences and structures, *Nucleic acids research* **41**, D246 (2012).
22. V. Agarwal, G. W. Bell, J.-W. Nam and D. P. Bartel, Predicting effective microRNA target sites in mammalian mrnas, *elife* **4**, p. e05005 (2015).
23. I. Kalvari, J. Argasinska, N. Quinones-Olvera, E. P. Nawrocki, E. Rivas, S. R. Eddy, A. Bateman, R. D. Finn and A. I. Petrov, Rfam 13.0: shifting to a genome-centric resource for non-coding rna families, *Nucleic acids research* **46**, D335 (2017).
24. J. Piñero, À. Bravo, N. Queralt-Rosinach, A. Gutiérrez-Sacristán, J. Deu-Pons, E. Centeno, J. García-García, F. Sanz and L. I. Furlong, Disgenet: a comprehensive platform integrating information on human disease-associated genes and variants, *Nucleic acids research*, p. gkw943 (2016).
25. Z. Huang, J. Shi, Y. Gao, C. Cui, S. Zhang, J. Li, Y. Zhou and Q. Cui, Hmdd v3. 0: a database for experimentally supported human microRNA–disease associations, *Nucleic acids research* (2018).
26. G. Chen, Z. Wang, D. Wang, C. Qiu, M. Liu, X. Chen, Q. Zhang, G. Yan and Q. Cui, Lncrnadisease: a database for long-non-coding rna-associated diseases, *Nucleic acids research* **41**, D983 (2012).
27. L. R. Dice, Measures of the amount of ecologic association between species, *Ecology* **26**, 297 (1945).
28. J. C. Gower, A general coefficient of similarity and some of its properties, *Biometrics*, 857 (1971).
29. J. Bromley, I. Guyon, Y. LeCun, E. Säkingner and R. Shah, Signature verification using a” siamese” time delay neural network, in *Advances in neural information processing systems*, 1994.
30. G. Koch, R. Zemel and R. Salakhutdinov, Siamese neural networks for one-shot image recognition, in *ICML Deep Learning Workshop*, 2015.
31. R. Riad, C. Dancette, J. Karadayi, N. Zeghidour, T. Schatz and E. Dupoux, Sampling strategies in siamese networks for unsupervised speech representation learning, *arXiv preprint arXiv:1804.11297* (2018).
32. T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado and J. Dean, Distributed representations of words and phrases and their compositionality, in *Advances in neural information processing systems*, 2013.
33. A. Grover and J. Leskovec, node2vec: Scalable feature learning for networks, in *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016.
34. J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan and Q. Mei, Line: Large-scale information network embedding, in *Proceedings of the 24th International Conference on World Wide Web*, 2015.
35. M. Ou, P. Cui, J. Pei, Z. Zhang and W. Zhu, Asymmetric transitivity preserving graph embedding, in *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016.
36. D. Wang, P. Cui and W. Zhu, Structural deep network embedding, in *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016.
37. M. V. Kuleshov, M. R. Jones, A. D. Rouillard, N. F. Fernandez, Q. Duan, Z. Wang, S. Koplev, S. L. Jenkins, K. M. Jagodnik, A. Lachmann *et al.*, Enrichr: a comprehensive gene set enrichment analysis web server 2016 update, *Nucleic acids research* **44**, W90 (2016).
38. M. Kanehisa and S. Goto, Kegg: kyoto encyclopedia of genes and genomes, *Nucleic acids research* **28**, 27 (2000).