Gene2Vec: Distributed Representation of Genes Based on Co-Expression

Jingcheng Du[*], Peilin Jia[*], Yulin Dai, Cui Tao, Zhongming Zhao[#], Degui Zhi[#]

School of Biomedical Informatics, University of Texas Health Science Center at Houston, Houston, TX, USA

[*]equal contribution
[#]equal contribution

Abstract

Existing functional description of genes are categorical, discrete, and mostly through manual process. In this work, we explore the idea of gene embedding, distributed representation of genes, in the spirit of word embedding. From a pure data-driven fashion, we trained a 300 dimension vector representation of all human genes, using gene co-expression patterns in 984 data sets from the GEO databases. These vectors capture functional relatedness of genes in terms of recovering known pathways - the average inner product (similarity) of genes within a pathway is 1.68X greater than that of random genes. Using t-SNE, we produced a gene co-expression map that shows local concentrations of tissue specific genes. We also illustrated the usefulness of the embedded gene vectors, laden with rich information on gene co-expression patterns, in tasks such as gene-gene interaction prediction. Overall, we believe that this distributed representation of genes may be useful for more bioinformatics applications.

Introduction

Genes, discrete segments of the genome that are transcribed, are basic building blocks of molecular biological systems. Although almost all transcripts in the human genome have been identified, functional annotation of genes is still a challenging task. Most existing annotation efforts organize genes into functional categories, e.g., pathways, or represent their relationship into networks. Pathways and networks crystallize biological knowledges and are convenient qualitative conceptualization of gene functions. Yet the exact functions of a gene are often more subtle and elusive to be expressed in qualitative terms.

The challenge of creating a quantitative semantic representation of discrete units of a complex system is not unique to gene systems. For a long time, creating a quantitative representation of words had been challenging for linguistic modeling. Hinton proposed the pioneering idea of 'learning distributed representations of words' (Hinton 1986), i.e., representing the semantics of a word by mapping them to vectors in a high-dimension space. However Hinton's idea did not lead to real implementation in mainstream natural language processing (NLP) research, until recently. The word2vec model achieved success in NLP modeling (Mikolov et al. 2013). This process of distributed representation is often called neural embedding because the embedding

function is often expressed by a neural network with a large number of parameters. This success of word2vec inspires us to investigate the possibility to represent gene functions via neural embedding.

In this study, we aim to represent genes as vectors in a high-dimension space, i.e., a gene embedding. In the word2vec model (Mikolov et al. 2013), a word embedding is trained by maximizing the probability of word co-occurrences in context, i.e., only a few words apart in a same sentence. Analogously, we defined the context of a gene by the other genes that co-expressed with it. We derive an embedding such that the probability of the context of a gene is maximized. While it is possible to training a gene embedding using the standard NLP word embedding by using a biomedical corpus, such as the PubMed abstracts, published literature is incomplete and biased towards genes that are well-studied. Therefore, we intended to adopt a purely data-driven fashion.

Using co-expression patterns of all human genes in 984 whole transcriptome human gene expression data sets from Gene Expression Omnibus (GEO), we learned a gene embedding using a neural network. We show that our embedding grouped related genes in clusters. Moreover, we demonstrate the usefulness of the learned gene embedding to downstream tasks in the problem of prediction of gene-gene interaction.

Methods

Data collection

Overview. We chose to use GEO data with rationale from both biological and technical aspects. In cellular systems, the mRNA expression levels represent activities of genes with fine resolutions. Over the past 10-20 years, GEO deposits the majority of microarray-based gene expression data in various conditions. Although the recent development of RNA-sequencing has generated transcriptomic data with advantages in both accuracy and scales than array-based data, the large cohort of GEO data provides features that are more suitable for our work. GEO data have been curated for over ten years and hence, the measurement covers a wide range of cell and tissue types, cellular conditions, disease status, and developmental stages. As our ultimate goal is to build a gene co-expression map that could be used for inferences in various conditions, we collected GEO data for our task. In addition, we chose one single platform to reduce technical variability and required the organism to be Homo sapiens.

Gene expression. We used the keywords “expression and human” to search in GEO on 12/24/2017 and retrieved all GSE sets that were conducted using the platform Affymetrix Human Genome U133 Plus 2.0 Array (GPL570). We required each dataset to have >= 30 samples. The downloaded gene expression intensity data were log transformed and quantile-normalized. For genes with multiple probe sets, we chose the probe set with the largest variance across all samples. Gene co-expression was measured using Pearson Correlation Coefficient (PCC) for each data set. In each data set, gene pairs with the absolute PCC >= 0.9 were selected for following analysis. Selected gene pairs from all data sets were merged and serve as training data. We did not distinguish biological conditions.

Gene types on chip. The U133 array is one of the most widely utilized platform to measure human gene expression. The chip has 54,675 probe sets for 24,442 genes. The number of probe sets per gene ranged between 1 and 15, with more than half of genes (52.08%) have one probe sets. Among these genes, 21,960 (89.85%) genes could be mapped to the current version of NCBI Entrez gene annotation. These mappable genes include 18,055 protein-coding genes, 2,660 ncRNA, 730 pseudo genes, 132 snoRNA, and 383 other types of genes. Particularly for ncRNAs, there are 202 microRNA genes. Gene set enrichment analysis was conducted using Fisher's Exact Test.

Gene-gene interaction dataset. We followed previous work (Lee et al. 2011) to build datasets for gene-gene interaction based on shared Gene Ontology (GO) annotations. GO annotation was obtained using the R (x64 3.4.3) package "org.Hs.eg.db" (version 3.5.0). GO structure file in the obo format was downloaded from http://geneontology.org/ontology/go.obo (2018-02-14). All genes were mapped to NCBI Entrez Gene (downloaded 11/6/2017). We defined gene pairs that shared GO annotations as the positive set of functional association. To this end, we chose the GO category "Biological Process" with experimental evidence: IDA (inferred from direct assay), IMP (inferred from mutant phenotype), IPI (inferred from protein interaction), IGI (inferred from genetic interaction), and TAS (traceable author statement). To minimize generalized annotation, we excluded the highly over-represented GO terms including (1) "signal transduction" (GO:0007165); (2) three phosphorylation terms: "protein amino acid phosphorylation" (GO:0006468), "protein amino acid autophosphorylation" (GO:0046777), and "protein amino acid dephosphorylation" (GO:0006470); and (3) all terms at the first three levels of GO hierarchy (assuming the root term of biological process, "GO:0008150", is level 0). This lead to a total of 270,704 pairs involving 5369 genes. To build the negative data set, we obtained all gene-pairs that did not share any GO term or their children GO terms. This resulted a total of 40,879,714 gene pairs involved in 12,521 (64.85% of 19307) human genes, serving as the set in which pairs of genes are not functionally associated.

Tissue-specific genes. GTEx data (version 6) (The GTEx Consortium 2015) was used to estimate the tissue-specific expression pattern of genes in 27 tissues, each with ≥ 30 samples. For each gene, a *z*-score was calculated to measure its tissue specificity by comparing the average gene expression of the gene across all tissues (Zhao et al. 2018).

Functional gene sets. We use clusteredness of MSigDB pathways (v5.1)(Subramanian et al. 2005) as the target function for hyper-parameter tuning for gene embedding training. Specifically, we used the the category c2 including curated pathways from various online resources such as KEGG[1], Biocarta[2], and Reactome[3] . A total of 4,726 pathways were downloaded.

[1] http://www.pathway.jp
[2] http://cgap.nci.nih.gov/Pathways/BioCarta_Pathways
[3] http://www.reactome.org

Concept embedding of genes

Distributed representation of word, or neural word embedding, was a recent breakthrough in NLP research based on deep learning. The goal of word embedding is to derive a linear mapping, i.e., embedding, from the discrete space of individual words to a continuous euclidean space such that similar words will be mapped to points in close vicinity in the embedding space. The direct benefit of word embedding is that such representation of individual words, vectors in continuous space, becomes differentiable and thus amenable for back-propagation-based neural network modeling. Meanwhile, a nice surprising result is that embedded space admits basic geometry. E.g., the KING - QUEEN $\approx$ MAN - WOMAN.

Inspired by the success of word embedding, we intend to produce an embedding of genes, also a discrete conceptual unit, such that similar genes are mapped to similar vectors. While for genes we do not have a natural equivalent concept of sentence in natural languages, we will use the notion of co-expression. This is analogy of the concept of co-occurrence in natural languages.

For neural embedding, a neural network is designed that maximizes an objective function, often in a form of likelihood, such as the probability of a word given its context. The most commonly used architectures are skip-gram and continuous bag-of-words (CBOW) that discussed in the word2vec (Mikolov et al. 2013) approach. In both architectures, a two-layer neural network is constructed to predict word co-occurrence, or the co-occurrence of a word and its surrounding words, or context. In CBOW, the input is the context and the output is the word; in skip-gram, the input is the the word and the output is the context. For both architectures, input and output are connected through a middle projection layer. Note that neither neural network would offer satisfactory predictions for most of the words. But the real goal of word embedding is to learn a distributional representation, i.e., the parameters of the embedding mapping from the input to the middle projection layer. A simple fully connected linear layer was used for the embedding mapping. For CBOW, the embedded vectors of all words in the context are averaged and thus provide a uniform size vector for the next layer. The second layer for both architecture is a linear layer with a soft-max. A cross-entropy loss is minimized.

In gene embedding, we are using the genes who are co-expressed with the gene of interest as its context. Since the number of co-expressed genes may vary, the size of the context may vary as well. For simplicity, in this work, we extract all pairs of co-expressed genes and maximize the probability of one given the other for each pair. This is equivalent to the skip-gram model. Since we are optimizing the total probability of all edges in a co-expression network, our approach can also be viewed as a graph embedding (Cai, Zheng, and Chang 2017).

More formally, the input of the training problem is a list of gene pairs that are highly co-expressed, $T = \{(g_{i1}, g_{i2})\}$, we will train an embedding network. The input of the network is a one-hot encoded vector for gene $g_i \in R^d$, where d is the number of genes and the elements of g_i are all 0 except $g_i[x_i] = 1$, where x_i is the dimension corresponds to the gene g_i. The out of

the network is a vector of dimension $v_i \in R^k$, the embedding dimension. The parameter of the network is a matrix $W \in R^{d \times k}$ such that $v_i = W g_i$. If the we define the probability

$$Pr(g_i|g_j) = \frac{exp(v_i^T v_j)}{\sum_{j'} exp(v_i^T v_{j'})}$$

The loss function that is to be minimized is the negative likelihood $-\sum_{(g_{i1}, g_{i2}) \in T} log(Pr(g_{i1}|g_{i2}))$.

It can be shown that this complex loss function for this single layer network is equivalent to a two-layer network with shared weight matrices of W and W^T, and the loss function as the standard cross-entropy after softmax (**Figure 1**).

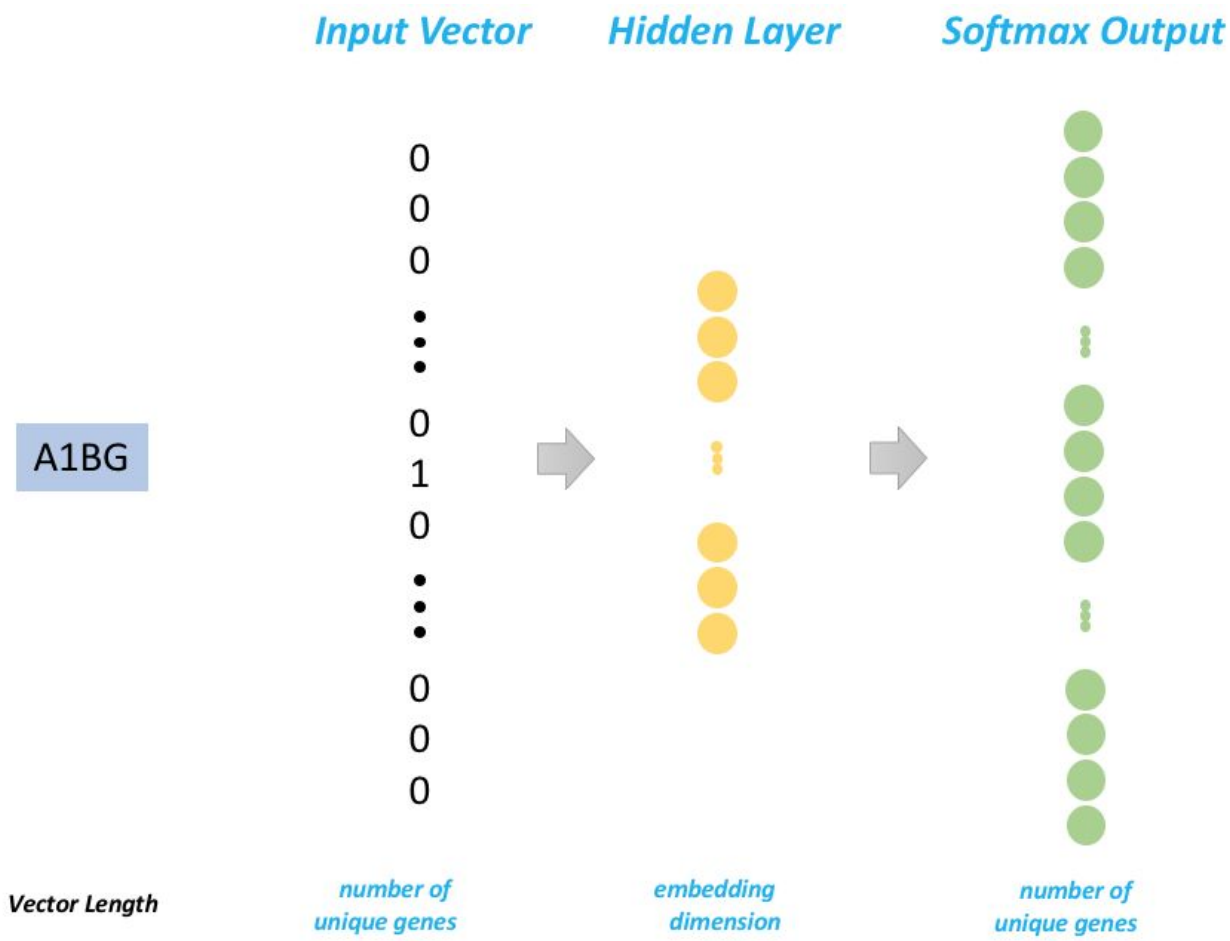

Figure 1: The Skip-Gram architecture was used for training for gene embedding. This is the modified architecture which is equivalent to the original word2vec, adopted from this blog[4].

[4] http://mccormickml.com/2016/04/19/word2vec-tutorial-the-skip-gram-model/

Training of embedding

We took all the gene pairs that have a Pearson Correlation Coefficient equal or larger than 0.9 as the input. This is a choice due to limited computational resources. We shuffled the gene pairs in each dataset on every iteration to avoid the impact caused by the order of gene pairs in the datasets. The embedding was trained on all genes with a minimum frequency at 5. As number of iterations and dimensionality of the embedding are considered as two major hyper-parameters parameters for word embedding (Lai et al. 2016), in order to generate "best" gene embedding, we did a preliminary parameters tuning and performed a grid search to find best parameters. The search ranges for number of iterations and embedding dimension are set at 1 to 10 and 50, 100, 200, and 300 respectively. We used the word2vec function implemented in the gensim library (Řehůřek and Sojka 2010) to generate gene embedding. Other parameters were set as default.

Since our goal is to obtain a gene embedding that reflects the functional relationships among genes, we selected the set of hyper-parameters that maximizes the clusteredness of genes within functional pathways. We optimized the following target function:

$$clusteredness = \frac{\frac{1}{|Q|}\sum_{P\in Q}\frac{1}{\#\,gene\ pairs\ in\ P}\sum_{g_i,g_j\in P}(v_i{}^T v_j)}{\frac{1}{\#\,gene\ pairs\ in\ Q'}\sum_{g_1,g_2\in Q'}(v_1{}^T v_2)},$$

where Q is the set of pathways in MSigDB, and Q' is a set of random gene pairs. Due to the limitation of computation power, we selected all the pathways from the MSigDB with the number of genes equal or fewer than 50. In total, 6,729 pathways were selected as Q. We randomly selected 1,000 genes from gene embedding and generated all possible unique gene pairs (499,500 in total) as Q'.

Visualization by t-SNE

A common way to visualize high-dimensional datasets is to map the datasets into 2D or 3D array. t-Distributed Stochastic Neighbor Embedding (t-SNE) is a machine learning algorithm for dimensionality reduction, which optimizes for neighborhood preserving and thus particularly well suited for the visualization of high-dimensional datasets(Maaten and Hinton 2008). Visualizations produced by t-SNE have been found significantly better than those produced by the other techniques (Maaten and Hinton 2008).

In order to speed up the t-SNE on the high-dimensional gene embedding, we first reduced the dimension to 50 using PCA and then applied a multicore modification of Barnes-Hut t-SNE by L. Van der Maaten(Van Der Maaten 2014; Ulyanov 2016). The perplexity was set at 30 and the learning rate was set at 200. To get a stable t-SNE results, we set the number of iterations at 100,000.

Prediction of gene-gene interaction

To investigate the usefulness of the trained gene embedding for downstream tasks, we applied the embedding to the problem of gene-gene interaction prediction. The goal is, given a pair of genes, we design a gene-gene interaction predictor neural network (GGIPNN) to predict if they will be together in any of the annotated pathway (Lee et al. 2011).

The architecture of GGIPNN can be seen in **Figure 2**. We first convert the genes in each gene pair to one-hot vectors and then map the one-hot vectors to gene embedding vectors using a shared embedding matrix. Then, the two gene embedding vectors will be concatenated together and be fed to a fully connected layer with a dimension at 100. The output of the first fully connected layer will then be fed to a second fully connected layer with a dimension at 10. The output will be then fed to a softmax (same as sigmoid as this is a binary classification) layer. We compute the cross entropy of the softmax function output and then compute the mean of elements across results as the loss. To avoid overfitting, we apply dropout on both the first and second fully connected layers. The dropout out rates are set at 0.1. Area Under Curve (AUC) will be computed to measure the performance of the prediction. We compared the AUC score of our pre-trained gene embedding and embedding which is randomly initialized. We also investigated the impact of trainable embedding layer (fine tuning during the training) versus non-trainable embedding layer (fixed during the training) on the prediction. This model was implemented in TensorFlow[5].

We took all the pairs from HumanNet.v1.benchmark (Lee et al. 2011) as the positive pairs. This led to a total of 270,704 pairs involving 5,369 genes. To build the negative data set, we obtained all gene-pairs that did not share any GO term or their children GO terms. This resulted a total of 40,879,714 gene pairs involved in 12,521 (64.85% of 19,307) human genes, serving as the set in which pairs of genes are not functionally associated. To avoid the impact of the imbalanced labels distribution, we randomly selected negative pairs with the equal number of the positive pairs. We then split all the unique genes into training, validation and testing sets with an proportion of 7 : 1 : 2. The pairs that the both two genes belong to training set are used as training; the pairs that the both two genes belong to validation set are used as validation; the pairs that the both two genes belong to testing set are used as testing. By doing so, we avoid the possibility that the neural network "memorizes" the likelihood of genes to be interacting with any other genes. In total, the training dataset has 263,016 pairs (involving 8,832 genes), while the validation and testing dataset have 5,568 pairs (1,173 genes) and 21,448 pairs (2,467 genes) respectively.

[5] https://www.tensorflow.org/

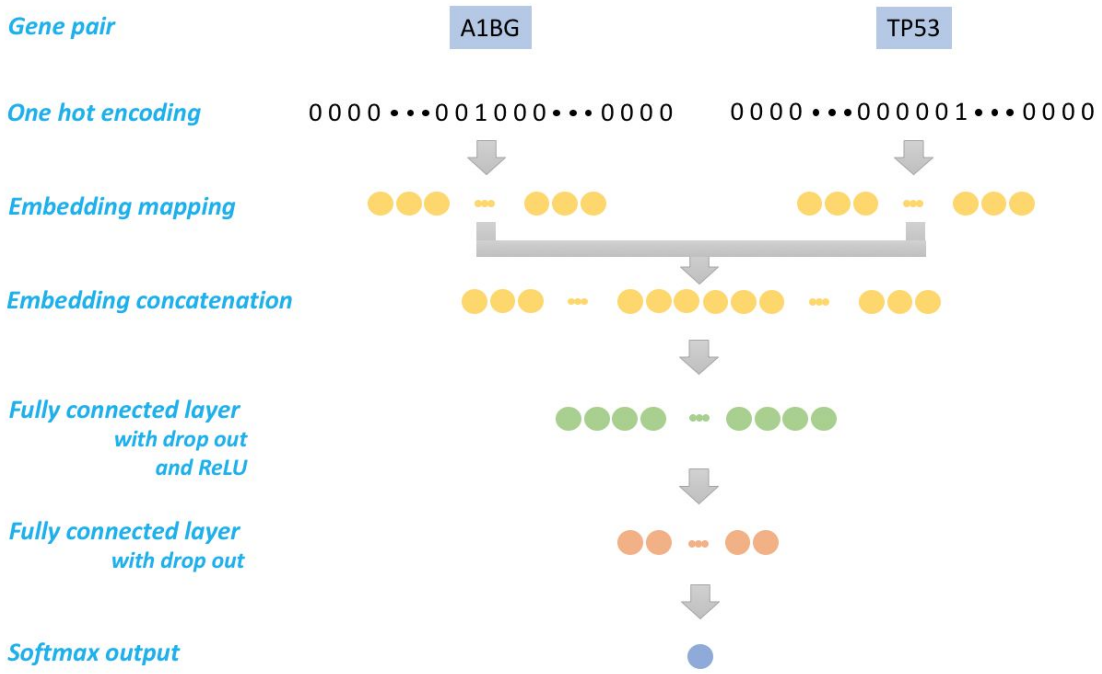

Figure 2: The architecture of gene-gene interaction predictor neural network (GGIPNN).

Results

Parameter tuning results by clusteredness:

The parameter tuning results can be seen in **Table 1**. As we can observe, the dimension of 300 at iteration 9 produced best gene embedding using clusteredness as the target function. As result, we chose this embedding for all following analyses.

Dimension	**Number of Iterations**									
	1	2	3	4	5	6	7	8	9	10
50	1.48	1.52	1.56	1.60	1.61	1.58	1.59	1.59	**1.61**	1.57
100	1.47	1.57	1.57	1.60	1.61	1.64	**1.67**	1.65	1.64	1.63
200	1.50	1.58	1.63	1.63	1.67	1.63	1.66	1.66	**1.67**	1.64
300	1.51	1.55	1.65	1.64	1.65	1.63	1.65	1.67	**1.68**	1.66

Table 1: Hyperparameter tuning using clusteredness as target function.

Gene embedding groups similar genes into spatial clusters

Using the first and second components from the t-SNE representation, we produced a *gene co-expression map*, based on which we explored the distribution of all human genes from our results (**Figure 3**). A direct visualization of the gene distribution revealed that the majority of genes formed one single cloud while several isolated groups of genes scattered around. We extracted these gene islands and found they were mainly non-protein-coding genes. The island on top of the figure (island 1) was significantly populated with the snoRNA genes (p=6.44e-71, Fisher's Exact Test). Island 2, located to the very right of the plot, mainly contains human cDNA/PAC clone genes. microRNA genes were mainly distributed in island 1 (p=2.46e-19), island 3 (p=2.22e-40), and island 5 (p=2.25e-69). A group of ncRNAs which start with "LOC" and are often uncharacterized split the whole distribution into the left panel and the right panel. In the right panel, we observed a cluster of open reading frames in the human genome.

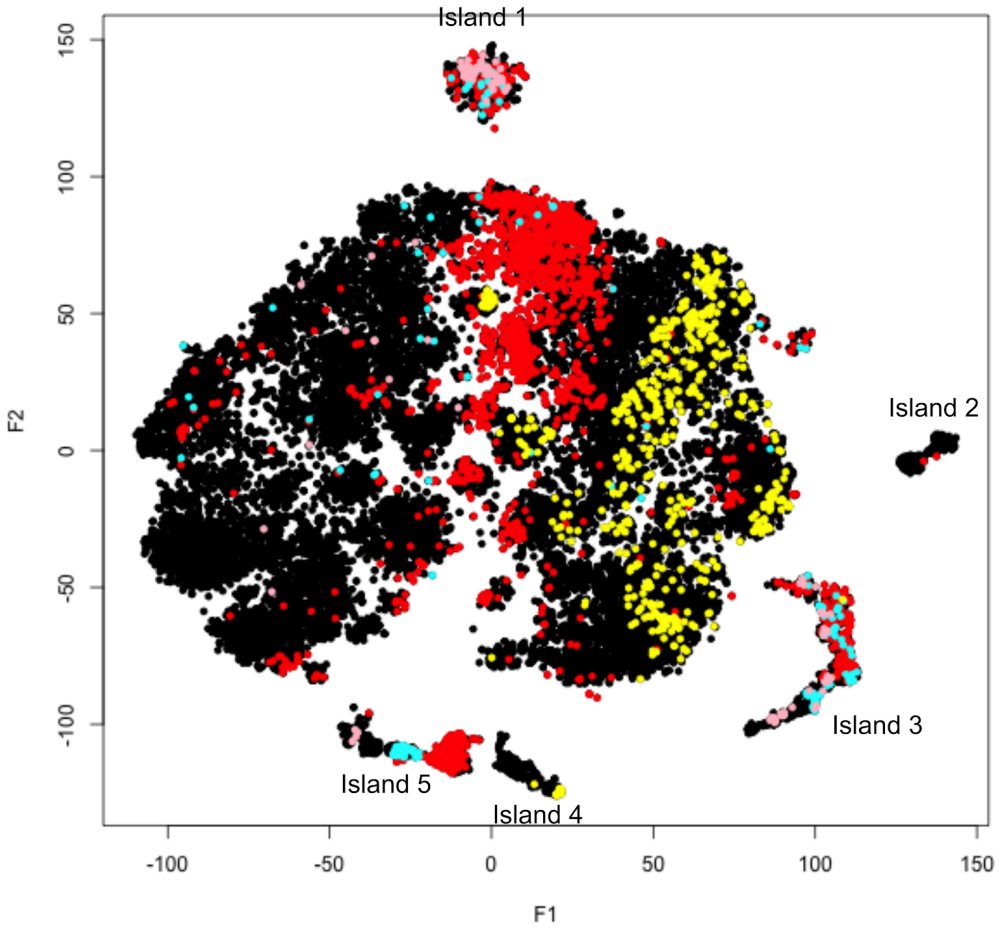

Figure 3: Gene co-expression map generated from embedding reveals clusters of functionally related genes. F1 and F2 are the first and the second dimensions of tSNE. Red: LOC non-coding genes; cyan: microRNA; pink: small nucleolar RNA (snoRNA); yellow: undercharacterized ORFs.

Tissue specific genes form spatial patterns in gene embedding

We mapped genes with z-scores representing their tissue-specific expression onto the gene co-expression map. We observed clear clusters in several tissues such as blood, skin, spleen, and lung (**Figure 4**, **supplementary figures**). Genes with high tissue specificity in blood formed two main clusters that are closely located to each other in the middle. This is likely because that blood samples are relatively more widely used in gene expression studies and blood-specific genes and their relationships are thus better represented in our map. Tissues that are biologically relevant showed similar patterns. For example, tissues of female reproductive systems presented graded and similar patterns, including breast, cervix uteri, fallopian tube, ovary, and uterus. In these tissues, genes located in the bottom part of the map in general showed increased tissue specificity, compared to genes located on the top part of the map (**Figure 4**, **supplementary figures**). Interestingly, we found a group of ribosomal genes (~50) that were highly expressed in ovary and formed a small cluster in our map. In addition, cognition and neurology related tissues, such as brain, nerve, and pituitary, presented quite diverse patterns. Nerve and pituitary are more similar to each other, with a wide range of genes showing moderate tissue-specificity distributed across the whole map. In contrast, active genes in brain, which are mainly distributed on the top part of the map, are much smaller in numbers but showed much stronger tissue-specificity (red dots). Notably, all tissues except the blood are expected to be under-represented in the GEO data we used because tissue samples are difficult to obtain for human.

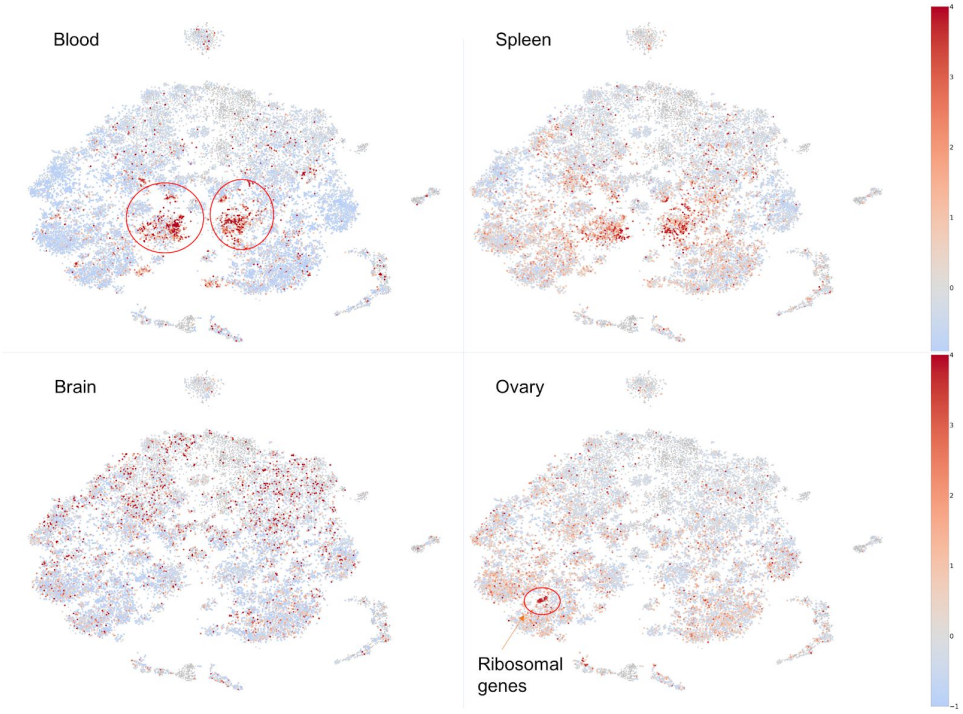

Figure 4: Embedding reveals clusters of genes with tissue-specificity. Blood and spleen have clear patterns of tissue-specific genes. Reproductive system (e.g., ovary) also showed distinguished genes. Genes not available in GTEx data were colored grey.

Prediction of gene-gene interaction using embedded vectors

The performances of GGIPNN with embedding matrix are presented in **Figure 5**. Using gene embedding matrix derived from GEO but do not make them trainable, we achieved an AUC of 0.698 over the test set, in which there are no gene overlapping with the training set nor the validation set. The AUC scores is lower, 0.649, for the GGIPNN with gene embedding matrix derived from GEO as initial weights but trainable. This is understandable as the gene embedding matrix for the genes in the training set was updated and leaving the gene embedding matrix in the test set "out of sync" with that for the training set, i.e., overfitting. As expected, the GGIPNN with random untrainable embedding matrix has an AUC score (0.493) close to random. This can be slightly improved if the embedding matrix is trainable (0.525), but still far from the performances of the informative gene embedding matrices.

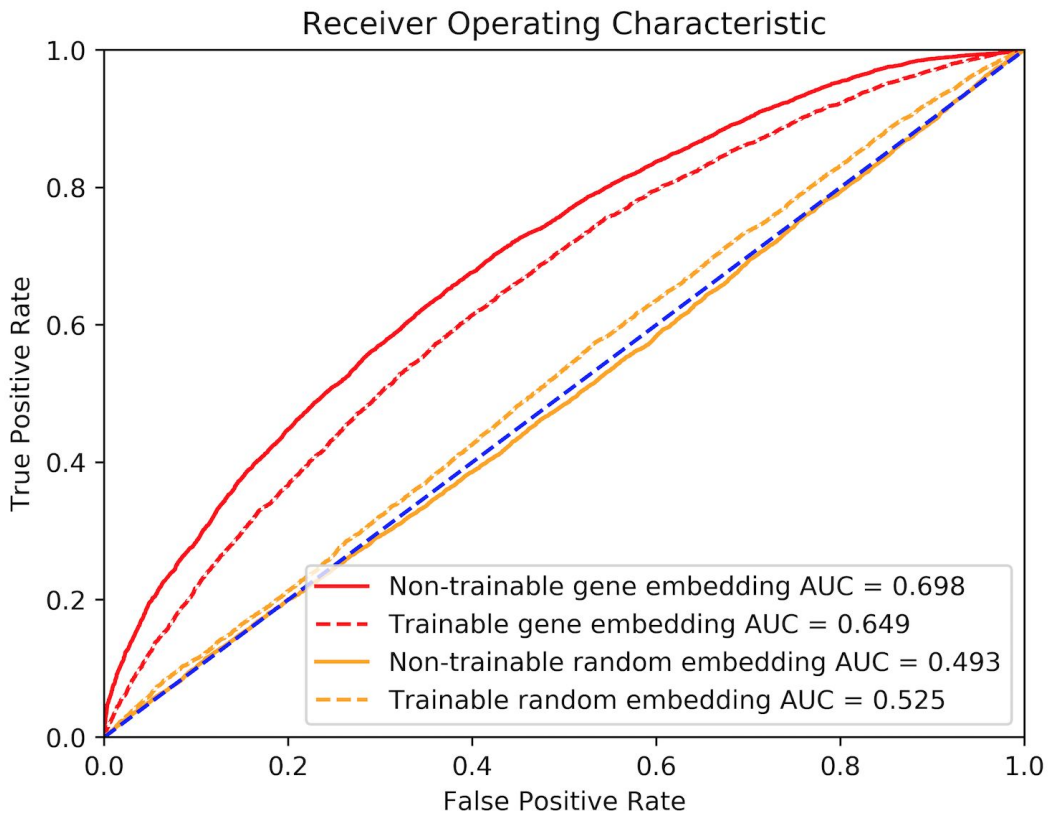

Figure 5: ROC curves for gene-gene interaction predictor neural networks.

Discussions

In this work we explored the idea of distributed representation of genes using their co-expression. Purely trained from their co-expression patterns in GEO, except using MSigDB as hyper-parameter tuning, the trained embedding matrix captures functional relationships among genes. In the t-SNE generated gene co-expression map of the embedding matrix, tight clusters of

non-coding genes are formed, while broader clusters corresponding to tissue specific genes are also visible.

The usefulness of gene embedding is beyond simply a nice visualization. Using the gene embedding as the basic layer for a multi-layer neural network, we can predict the gene-gene interaction with an AUC of 0.696. This is an intriguing result because the only input to the predictor is the names of the two genes. Therefore, the distributed representation of the genes, i.e., their embeddings, are laden with rich semantic information about their function.

The concept of concept embedding is not new to molecular biology. Works had been done to geometrical embedding gene co-expression networks into 2-D planar networks (Song and Zhang 2015). Recently, in the spirit of embedding everything, the work of bioVectors have been developed to embedding kmers in biological sequences into distributed representation (Asgari and Mofrad 2015). Similarly, a project named 'Gene2vec' is available embedding gene sequences[6]. However, to the best of our knowledge, our work is the first to directly embed genes into distributed representations based on their natural context - their expression and co-expression.

In this work, we are using the gene co-expression as the definition of "context" for gene embedding. However, it is possible to extend the current work to include other definitions of context for genes. For example, co-occurrence of genes across species, gene-gene and protein-protein interactions from experiments, and co-occurrences of genes in literature, all can be a source of information to define context.

The distributed representation of genes can enable new applications. E.g., as illustrated in the Results, with continuous representation of genes, it is possible to direct feed gene as inputs to neural networks, and can be useful for any prediction tasks with gene names as input.

A limitation of current approach is the lack of higher order semantics. In word embedding a surprising result was that the direction of the embedding space can be interpreted. For example, the vector representations of the words King, Queen, Man, and Woman formed a parallelogon. This higher order of semantics from NLP modeling may be due to that the concepts between these words were connected by certain relationships, which is reflected by the occurrences of these words being connected by certain verbs. To achieve this level of semantic embedding of genes, future works modeling more information about genes are warranted.

Acknowledgments

The authors acknowledge the support from the UTHealth Innovation for Cancer Prevention Research Training Program Pre-doctoral Fellowship (Cancer Prevention and Research Institute of Texas grant # RP160015).

[6] https://nbviewer.ipython.org/github/davidcox143/Gene2vec/blob/master/report/Gene2vec.ipynb

Disclaimer

The content is solely the responsibility of the authors and does not necessarily represent the official views of the Cancer Prevention and Research Institute of Texas.

References

Asgari, Ehsaneddin, and Mohammad R. K. Mofrad. 2015. "Continuous Distributed Representation of Biological Sequences for Deep Proteomics and Genomics." *PloS One* 10 (11). journals.plos.org:e0141287.

Cai, Hongyun, Vincent W. Zheng, and Kevin Chen-Chuan Chang. 2017. "A Comprehensive Survey of Graph Embedding: Problems, Techniques and Applications." *arXiv [cs.AI]*. arXiv. http://arxiv.org/abs/1709.07604.

Hinton, Geoffrey E. 1986. "Learning Distributed Representations of Concepts." In *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, 1:12. Amherst, MA.

Lai, S., K. Liu, S. He, and J. Zhao. 2016. "How to Generate a Good Word Embedding." *IEEE Intelligent Systems* 31 (6). ieeexplore.ieee.org:5–14.

Lee, Insuk, U. Martin Blom, Peggy I. Wang, Jung Eun Shim, and Edward M. Marcotte. 2011. "Prioritizing Candidate Disease Genes by Network-Based Boosting of Genome-Wide Association Data." *Genome Research* 21 (7):1109–21.

Maaten, Laurens van der, and Geoffrey Hinton. 2008. "Visualizing Data Using T-SNE." *Journal of Machine Learning Research: JMLR* 9 (Nov):2579–2605.

Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. "Efficient Estimation of Word Representations in Vector Space." *arXiv [cs.CL]*. arXiv. http://arxiv.org/abs/1301.3781.

Řehůřek, Radim, and Petr Sojka. 2010. "Software Framework for Topic Modelling with Large Corpora." In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, 45–50. Valletta, Malta: ELRA.

Song, Won-Min, and Bin Zhang. 2015. "Multiscale Embedded Gene Co-Expression Network Analysis." *PLoS Computational Biology* 11 (11). Public Library of Science:e1004574.

Subramanian, Aravind, Pablo Tamayo, Vamsi K. Mootha, Sayan Mukherjee, Benjamin L. Ebert, Michael A. Gillette, Amanda Paulovich, et al. 2005. "Gene Set Enrichment Analysis: A Knowledge-Based Approach for Interpreting Genome-Wide Expression Profiles." *Proceedings of the National Academy of Sciences of the United States of America* 102 (43). National Acad Sciences:15545–50.

The GTEx Consortium. 2015. "The Genotype-Tissue Expression (GTEx) Pilot Analysis: Multitissue Gene Regulation in Humans." *Science* 348 (6235). American Association for the Advancement of Science:648–60.

Ulyanov, Dmitry. 2016. "Multicore-TSNE." *GitHub Repository*. GitHub. https://github.com/DmitryUlyanov/Multicore-TSNE.

Van Der Maaten, Laurens. 2014. "Accelerating T-SNE Using Tree-Based Algorithms." *Journal of Machine Learning Research: JMLR* 15 (1). jmlr.org:3221–45.

Zhao, Junfei, Feixiong Cheng, Peilin Jia, Nancy Cox, Joshua C. Denny, and Zhongming Zhao. 2018. "An Integrative Functional Genomics Framework for Effective Identification of Novel Regulatory Variants in Genome--Phenome Studies." *Genome Medicine* 10 (1). BioMed Central:7.