
Research and Applications

Constructing co-occurrence network embeddings to assist association extraction for COVID-19 and other coronavirus infectious diseases

David Oniani¹, Guoqian Jiang², Hongfang Liu², and Feichen Shen²

¹Kern Center for the Science of Health Care Delivery, Mayo Clinic, Rochester, Minnesota, USA, and ²Division of Digital Health Sciences, Mayo Clinic, Rochester, Minnesota, USA

Corresponding Author: Feichen Shen, PhD, Division of Digital Health Sciences, Mayo Clinic, Rochester, MN 55901, USA; shen.feichen@mayo.edu

Received 12 May 2020; Revised 19 May 2020; Editorial Decision 20 May 2020; Accepted 22 May 2020

ABSTRACT

Objective: As coronavirus disease 2019 (COVID-19) started its rapid emergence and gradually transformed into an unprecedented pandemic, the need for having a knowledge repository for the disease became crucial. To address this issue, a new COVID-19 machine-readable dataset known as the COVID-19 Open Research Dataset (CORD-19) has been released. Based on this, our objective was to build a computable co-occurrence network embeddings to assist association detection among COVID-19–related biomedical entities.

Materials and Methods: Leveraging a Linked Data version of CORD-19 (ie, CORD-19-on-FHIR), we first utilized SPARQL to extract co-occurrences among chemicals, diseases, genes, and mutations and build a co-occurrence network. We then trained the representation of the derived co-occurrence network using node2vec with 4 edge embeddings operations (L1, L2, Average, and Hadamard). Six algorithms (decision tree, logistic regression, support vector machine, random forest, naïve Bayes, and multilayer perceptron) were applied to evaluate performance on link prediction. An unsupervised learning strategy was also developed incorporating the t-SNE (t-distributed stochastic neighbor embedding) and DBSCAN (density-based spatial clustering of applications with noise) algorithms for case studies.

Results: The random forest classifier showed the best performance on link prediction across different network embeddings. For edge embeddings generated using the Average operation, random forest achieved the optimal average precision of 0.97 along with a F1 score of 0.90. For unsupervised learning, 63 clusters were formed with silhouette score of 0.128. Significant associations were detected for 5 coronavirus infectious diseases in their corresponding subgroups.

Conclusions: In this study, we constructed COVID-19–centered co-occurrence network embeddings. Results indicated that the generated embeddings were able to extract significant associations for COVID-19 and coronavirus infectious diseases.

Key words: COVID-19, coronavirus INFECTIOUS diseases, co-occurrence network embeddings, association extraction

INTRODUCTION

Having now affected millions of people worldwide, coronavirus disease 2019 (COVID-19)/the novel coronavirus has become a major pandemic of the century. Most countries have declared a state of national emergency and took actions effective immediately to slow the spread. Researchers and medical personnel around the world have published and released thousands of articles over a short period of time, covering a vast scientific ground and exploring medical treatments and possible vaccines for the virus.¹ With all this information, it is important to assemble all the available heterogeneous information and be aware of the explicit or implicit associations among subjects related to COVID-19 (eg, certain genes could be linked to other genes and/or mutations related to COVID-19 and other coronavirus infectious diseases). Figuring out which subjects appear together is one of the approaches for identifying these associations and linking them together. Traditionally, text semantic similarity^{2,3} is one of the approaches for detecting links between words or sentences from unstructured data. One limitation is that it is inefficient to apply this approach over a large collection of free-text data, hampering the global view to detect significant associations across literature from heterogeneous domains. Normalized data stored in semi-structured graph format is more suitable for global link detection, as linked data⁴ by nature provide efficient query scheme over triplets to interpolate between “macroscopic” and “microscopic” search.

Several efforts were made in the graph-based analysis of COVID-19. For example, Ahamed and Samad⁵ developed a graph-based model using abstracts of 10 683 COVID-19-related scientific articles and applying betweenness-centrality to rank order the importance of keywords related to drugs, diseases, pathogens, hosts of pathogens, and biomolecules. Bellomarini et al⁶ presented a report on ongoing work about the application of automated reasoning and knowledge graph technology to address the impact of the COVID-19 outbreak on the network of Italian companies. Tsiotas and Magafas⁷ used visibility graphs to study Greek COVID-19 infection curve as a complex network. Per request of the White House Office of Science and Technology Policy, a new COVID-19 machine readable dataset (CORD-19)⁸ has been released, and several studies have featured it dataset to investigate COVID-19-related topics. For example, Wolinski⁹ used CORD-19 to extract diseases at risk and calculate relevant indicators as well as created VIDAR-19 (Visualization of Diseases At Risk in CORD-19). Wang et al¹⁰ conducted CORD-19 named entity recognition leveraging the distant supervision strategy. CORD-19-on-FHIR is a Linked Data version of CORD-19.¹¹ It is represented in FHIR RDF^{12,13} and was produced by data mining the CORD-19 dataset and adding semantic annotations. In addition, Groza¹⁴ featured CORD-19-on-FHIR in the analysis of how semantically annotated dataset can be applied for detecting and preventing the potential spread of deceptive information regarding COVID-19.

A vast co-occurrence information contained in CORD-19 datasets allows for detection of novel associations across findings from various research articles. However, such information has been largely unexplored for association extraction. Moreover, the lack of measurable association among heterogeneous biomedical entities hampers the capability for a quantitative analysis. Inspired by the success of word embeddings¹⁵ in building distributed semantic representations for each word given a corpus, network embeddings provide a solution to map graph nodes to distributional representations and translate nodes' relationships from graph space to embedding space, which makes the association between the nodes measur-

able.¹⁶⁻¹⁸ In this study, we filled this gap by constructing network embeddings for the CORD-19 co-occurrence network. Specifically, we first derived a co-occurrence network by querying the CORD-19-on-FHIR and focused on the extraction of biomedical entities falling in 4 categories: chemical, disease, gene, and mutation. We then applied the node2vec model over the generated network and constructed network embeddings. We conducted the evaluation quantitatively and qualitatively. For the quantitative evaluation, we generated different embeddings with 4 embeddings generation operations using a downstream application on graph link prediction and measured the performance with 6 machine learning algorithms. For the qualitative evaluation, we visualized clusters generated by the optimal COVID-19 network embeddings and analyzed associations of heterogeneous biomedical entities related to COVID-19 and other coronavirus infectious diseases.

MATERIALS AND METHODS

CORD-19-on-FHIR

The purpose of building CORD-19-on-FHIR is to represent linkage with other biomedical datasets and enable answering research questions. In this study, we used a subset of CORD-19-on-FHIR datasets annotated by PubTator¹⁹ and LitCovid,¹ including 3207 COVID-19-related articles in total. Each article was stored in one specific annotation file. For each file and for each paragraph in the file, CORD-19-on-FHIR provides a way to capture all the annotated biomedical entities. A high level example of data stored in the Terse RDF Triple Language (Turtle) format is shown:

```

• pmc: annotations [
•   pmc: id "1";
•   pmc: infons [pmc: identifier "MESH: D003371";
•     pmc: type "Disease"];
•   pmc: locations [pmc: length "5"^^xsd: int;
•     pmc: offset "20312"^^xsd: int];
•   pmc: text "cough";
•   pmc: id "2";
•   pmc: infons [pmc: identifier "MESH:
•     C000657245"; pmc: type "Disease"];
•   pmc: locations [pmc: length "19"^^xsd: int;
•     pmc: offset "14766"^^xsd: int];
•   pmc: text "2019-nCoV infection" ] ,
• pmc: annotations [
•   pmc: id "5";
•   pmc: infons [pmc: identifier "59272"; pmc:
•     ncbi_homologene "41448"; pmc: type "Gene"];
•   pmc: locations [pmc: length "31"^^xsd: int;
•     pmc: offset "1986"^^xsd: int];
•   pmc: text "angiotensin-converting enzyme
•     2" ] ,
•   pmc: id "7";
•   pmc: infons [pmc: identifier "MESH:
•     C000657245"; pmc: type "Disease"];
•   pmc: locations [pmc: length "19"^^xsd: int;
•     pmc: offset "14766"^^xsd: int];
•   pmc: text "2019-nCoV infection" ] ,

```

where “pmc: annotations” was used to differentiate different paragraphs within a same article, “pmc: id” was used to indicate different biomedical entities along with entity type (“pmc: type”), location and offset (“pmc: location”), and the original text from lit-

erature (“pmc: text”). Such encoding of the data made it possible to easily detect co-occurrence of biomedical entities within a single paragraph for building the network across the literature.

Node2Vec

The node2vec model used a random walk-based sampling strategy to balance the graph homophily²⁰ and structural equivalence.²¹ The reason we chose to use node2vec is its ability to learn node representations with a balance between the breadth-first search and depth-first search, which is essential for learning associations in a graph with both local and global views.

METHODS

The workflow of this study is made of 3 modules, including a COVID-19-on-FHIR-based co-occurrence network generation module, a network embeddings construction module, and an unsupervised learning module (Figure 1).

Co-occurrence network generation

For each literature, we treated paragraph-level co-occurrence in this study. We first designed a SPARQL query statement to extract paragraph-level co-occurrence of biomedical entities from COVID-19-on-FHIR. Particularly, in order to largely collect coronavirus related diseases and comorbidities, we built a list of keywords for diseases and symptoms to constrain the searching space, which includes COVID-19, SARS, pneumonia, fever, fibrosis, diarrhea, coronavirus, bronchitis, Ebola, influenza, and ZIKA. We extracted co-occurrences between gene-disease, mutation-disease, and chemical-disease using the following SPARQL query by replacing “Biomedical_Entity” with “Gene,” “Mutation,” and “Chemical,” respectively:

- PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
- PREFIX fhir: <http://hl7.org/fhir/>
- PREFIX pmc: <https://www.ncbi.nlm.nih.gov/pmc/articles#>
- SELECT distinct? pmc_id0? text0? pmc_id1? text1 (count(? text1) as? count) WHERE {
- ? pmc pmc: annotations
- [pmc: id? id0; pmc: text? text0; pmc: infons
- [pmc: type? type0; pmc: identifier? pmc_id0]] .
- FILTER ((? type0 = Biomedical_Entity)) .
- {select * where{
- ? pmc pmc: annotations
- [pmc: id? id1; pmc: text? text1; pmc: infons
- [pmc: type? type1; pmc: identifier? pmc_id1]] .

- FILTER ((? type1 = "Disease") && (contains (lcase(str(? text1)), "coronavirus") || contains (lcase(str(? text1)), "sars") || contains (lcase(str(? text1)), "covid-19") || contains (lcase(str(? text1)), "pneumonia") || contains (lcase(str(? text1)), "fever") || contains (lcase(str(? text1)), "fibrosis") || contains (lcase(str(? text1)), "diarrhea") || contains (lcase(str(? text1)), "bronchitis") || contains (lcase(str(? text1)), "ebola") || contains (lcase(str(? text1)), "influenza") || contains (lcase(str(? text1)), "zika"))).
- }
- }
- }Group by? pmc_id0? text0? pmc_id1? text1 Order by DESC(? count)

The outputs of the query were composed of a list of pairwise biomedical entities with co-occurrence frequency. We then built a network based on this list by adding a link between any 2 biomedical entities if they have co-occurred at least once. As shown in Figure 1, the co-occurrence network was represented by source-target pairs, which were then used as input data for training node representations.

Network embeddings representation learning

We applied the node2vec model in this module. Node2vec implements a second-order random walk over the graph topological structure, denoting that 3 types of node are involved in a specific walk, namely source entity, intermediate entity, and target entity. Given any source entity as E_s , target entity as E_t , intermediate entity that exists on the path between E_s and E_t as E_i , normalization constant as Z , the distribution of entity E_t with a fixed length of random walk can be represented as:

$$P(E_t|E_i) = \begin{cases} \frac{\pi(E_i, E_t)}{Z} & \text{if } (E_i, E_t) \text{ is an edge} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $\pi(E_i, E_t)$ is a transition probability between entities E_i and E_t . Given the weight over edge (E_i, E_t) as $w(E_i, E_t) = 1$, $\pi(E_i, E_t)$ could be calculated as:

$$\pi(E_i, E_t) = \alpha(E_s, E_t) \cdot w(E_i, E_t) \quad (2)$$

$\alpha(E_s, E_t)$ is a searching bias term developed in node2vec. Specifically, node2vec introduced 2 hyperparameters p and q to balance between the breadth-first search and depth-first search searching strategies for both local and global optimization. Given the shortest distance between E_s and E_t as $sd(E_s, E_t)$, α for entities E_s and E_t is computed based on p and q :




Figure 1. Study workflow.

$$\alpha(E_s, E_t) = \begin{cases} \frac{1}{p} & \text{if } sd(E_s, E_t) = 0 \\ 1 & \text{if } sd(E_s, E_t) = 1 \\ \frac{1}{q} & \text{if } sd(E_s, E_t) = 2 \end{cases}$$

After learning the sampled network data using random walk, we then leveraged the Skip-Gram model to train entity representations on the sampled data. For each entity node $E_s \in E$ and all its sampled neighbors $N(E_s)$, the loss function for entity representation learning could be described as:

$$\max_f \sum_{E_s \in E} \log P(N(E_s) | f(E_s)) \quad (4)$$

In the end, we normalized the prediction distribution by using a nonlinearity (eg, softmax) and optimize this loss function using Stochastic gradient descent.

Unsupervised clustering of network embeddings

To render the relatively high-dimensional embedding representations of network embeddings into a lower-dimensional space, we utilized the t-distributed stochastic neighbor embedding (t-SNE) algorithm²² to render the embeddings for all entity nodes into a 2-dimensional space. t-SNE does not perform clustering in and of itself, but instead renders each node embedding into a (x, y) coordinate. As such, additional postprocessing is needed to regroup these points into discrete clusters. The density-based spatial clustering of applications with noise (DBSCAN) algorithm²³ was therefore used over output generated by the t-SNE to further partition different entity groups into distinct clusters. Given a parameter ϵ that denotes how close points should be to each other and another parameter k that indicates the minimum number of points, the DBSCAN clustered similar entity nodes together based on density according to the pre-defined 2 parameters.

EXPERIMENTS

From CORON-19-on-FHIR, we extracted 49 696 co-occurred biomedical entities for 3626 coronavirus related diseases, 5741 genes, 524 mutations, and 6878 chemicals. Thus the derived co-occurrence network contains 16 769 nodes and 49 696 edges in total.

For quantitative evaluation, we generated the optimal network embeddings by performing a downstream link prediction task. Link prediction is a procedure where the goal is to predict the relationship between any 2 nodes and use the performance of a prediction to evaluate the quality of the generated network embeddings. Edge embeddings were used in this task in order to investigate the relationships between nodes leveraging distributional representations provided by entity embeddings. For any given nodes E_s, E_t and their corresponding entity representations $f(E_s)$ and $f(E_t)$, edge embeddings were calculated using 4 operations, namely Hadamard, Average, L1 and L2 as shown in equations 5-8, respectively:

$$\text{Hadamard}(E_s, E_t) = f(E_s) * f(E_t) \quad (5)$$

$$\text{Average}(E_s, E_t) = \frac{f(E_s) + f(E_t)}{2} \quad (6)$$

$$\text{L1}(E_s, E_t) = |f(E_s) - f(E_t)| \quad (7)$$

$$\text{L2}(E_s, E_t) = \sqrt{|f(E_s) - f(E_t)|^2} \quad (8)$$

We used 6 conventional classification algorithms to evaluate the performance of different edge embeddings on link prediction task, including decision tree (DT),²⁴ logistic regression (LR),²⁵ support

vector machine (SVM),²⁶ random forest (RF),²⁷ naïve Bayes (NB),²⁸ and multilayer perceptron (MLP).²⁵ Specifically, The Boolean function $L(E_s, E_t)$ was used to determine the existence of edge(s) between nodes E_s and E_t , where $L(E_s, E_t) = 1$ indicates positive links and $L(E_s, E_t) = 0$ represents negative links. We fit features of edge embeddings with labels provided by $L(E_s, E_t)$ to train the model. For positive examples, for each of the 4 networks, 60%, 10%, and 30% of all their edges were used for training, validation, and testing purposes, respectively. For negative examples, an equal number of node pairs were randomly sampled (with the same ratio among training, validation, and testing sets as 60%, 10%, and 30%, respectively).

For each classifier, we plotted the receiver-operating characteristic (ROC) curve and computed the area under the ROC curve to report link prediction performance. Moreover, as shown in equations 9-12, we used precision, recall, F1 score, and average precision (AP) to quantify the link prediction performance among 4 edge embeddings.

$$\text{Precision} = \frac{|\{\text{True Relations}\} \cap \{\text{Predicted Relations}\}|}{|\{\text{Predicted Relations}\}|} \quad (9)$$

$$\text{Recall} = \frac{|\{\text{True Relations}\} \cap \{\text{Predicted Relations}\}|}{|\{\text{True Relations}\}|} \quad (10)$$

$$\text{F1 score} = \frac{2 * \text{precision} * \text{recall}}{\text{precision} + \text{recall}} \quad (11)$$

$$\text{AP} = \sum_n (\text{Recall}_n - \text{Recall}_{n-1}) \text{Precision}_n \quad (12)$$

For qualitative evaluation, we first visualized the network embeddings clustering output and used the silhouette score to evaluate clustering outputs. Silhouette score is adopted to calculate the average distance to entities in the same cluster with the average distance to entities in other clusters. Given any entity node e in cluster C_e , the internal mean distance is defined as:

$$m(e) = \frac{1}{C_e - 1} \sum_{t \in C_e, e \neq t} d(e, t) \quad (13)$$

where $d(e, t)$ is the distance between node e and t in C_e . Similarly, external mean distance is described as:

$$n(e) = \min_{k \neq e} \frac{1}{C_k} \sum_{t \in C_k} d(e, t) \quad (14)$$

Overall, the silhouette score is calculated incorporating both internal and external mean distances:

$$s(e) = \begin{cases} 1 - \frac{m(e)}{n(e)}, & \text{if } m(e) < n(e) \\ 0, & \text{if } m(e) = n(e) \\ \frac{n(e)}{m(e)} - 1, & \text{if } m(e) > n(e) \end{cases} \quad (15)$$

For some selected coronavirus infectious diseases, we also located the cluster they belonged to and checked the most similar entities within the same cluster using cosine similarity. Let E_s denote any given coronavirus infectious disease and E_t denote any biomedical entity inferred by network embeddings, and $f(E_s)$ and $f(E_t)$ represent the embeddings for E_s and E_t respectively, cosine similarity was calculated as shown in equation 16.

$$\text{cosine_similarity}(E_s, E_t) = \frac{f(E_s) \cdot f(E_t)}{\|f(E_s)\| \|f(E_t)\|} \quad (16)$$

Table 1. Evaluation results for the 4 edge embeddings operations along with 6 machine learning algorithms

Operations	Algorithms	Average Precision (AP)	ROC score	Precision	Recall	F1 score
Hadamard	DT	0.79	0.82	0.84	0.82	0.81
	LR	0.89	0.83	0.86	0.82	0.81
	SVM	0.80	0.81	0.85	0.81	0.81
	RF	0.92	0.89	0.87	0.86	0.86
	NB	0.82	0.84	0.86	0.84	0.84
	MLP	0.56	0.60	0.63	0.60	0.57
Average	DT	0.81	0.84	0.85	0.84	0.84
	LR	0.94	0.92	0.87	0.85	0.85
	SVM	0.83	0.86	0.87	0.86	0.85
	RF	0.97 ^a	0.96 ^a	0.91 ^a	0.91 ^a	0.90
	NB	0.88	0.91	0.91 ^a	0.91 ^a	0.91 ^a
	MLP	0.78	0.84	0.84	0.84	0.84
L1	DT	0.75	0.80	0.80	0.80	0.80
	LR	0.95	0.94	0.89	0.89	0.89
	SVM	0.87	0.89	0.90	0.89	0.89
	RF	0.96	0.95	0.89	0.88	0.88
	NB	0.85	0.88	0.89	0.88	0.88
	MLP	0.87	0.89	0.89	0.89	0.89
L2	DT	0.75	0.80	0.80	0.80	0.80
	LR	0.94	0.93	0.89	0.88	0.88
	SVM	0.87	0.88	0.90	0.88	0.88
	RF	0.96	0.95	0.89	0.88	0.88
	NB	0.85	0.87	0.88	0.87	0.87
	MLP	0.85	0.87	0.88	0.87	0.87

DT: decision tree; LR: logistic regression; MLP: multilayer perceptron; NB: naïve Bayes; RF: random forest; ROC: receiver-operating characteristic; SVM: support vector machine.

^aHighest value.

RESULTS

Different embeddings were generated by using neighbor size of 10, number of walks of 10, window size of 10, and dimensionality of 128. The optimal p and q were also tuned as 0.5 and 0.25 for each training process, respectively. Detailed network embeddings and the corresponding clustering information could be found online (<https://github.com/shenfc/COVID-19-network-embeddings>). We have also implemented a Web-based user-friendly tool for clustering visualization and entity similarity checking (<https://www.davidoniani.com/covid-19-network>).

Quantitative evaluation

As shown in Table 1, we presented the evaluation results of 4 different edge embedding operations along with 6 different classification algorithms. We found that, in general, the embeddings trained by the Average operation achieved the best performance across all the evaluation metrics. The optimal AP, ROC score, precision, and recall was reached when the RF was used, and the optimal F1 score was achieved by using NB. L1 and L2 had roughly the same performance, both peaking at ROC = 0.95 and AP = 0.96 with RF classifier. Among all 4 approaches, Hadamard yielded the worst performance, peaking at ROC = 0.89 and AP = 0.92 with RF classifier. Across all 6 classification algorithms, the worst performance, on the other hand, was shown by DT and MLP classifiers.

Regarding different classification algorithms, for the L1 embeddings embedding operation, RF and LR had similar performance, with ROC_{RF} = 0.96, AP_{RF} = 0.97, and ROC_{LR} = 0.96, AP_{LR} = 0.95. SVM, NB, and MLP were also not much different from each other. DT had the worst performance, with ROC = 0.80 and AP = 0.75. Similarly, for L2, RF and LR had roughly the same

performance. SVM, NB, and MLP were also equally performant. DT had the worst performance, with ROC = 0.80 and AP = 0.75. For the Average operation, RF has outperformed all the other classification methods, with ROC = 0.96 and AP = 0.97. LR and NB had similar performance. DT, SVM, and MLP were similar, yet all of them were behind RF, LR, and NB. For Hadamard, RF showed the best performance, with ROC = 0.89 and AP = 0.92. The rest of the methods showed roughly the same performance, except for MLP, which had the worst performance across all classification algorithms

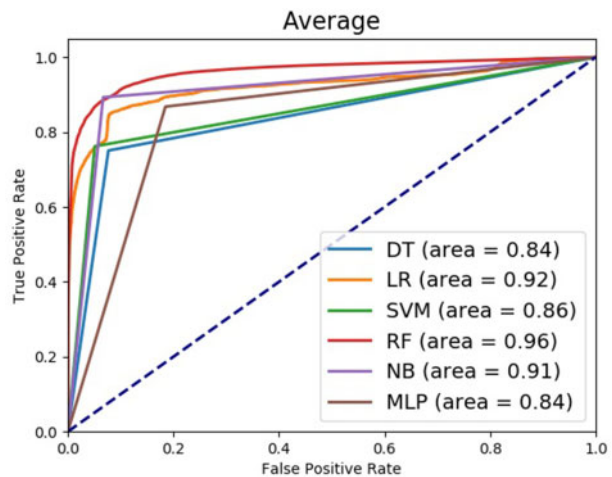


Figure 2. Receiver-operating characteristic scores for the average operation with 6 machine learning algorithms. DT: decision tree; LR: logistic regression; MLP: multilayer perceptron; NB: naïve Bayes; RF: random forest; SVM: support vector machine.

as well as across every embedding generation approach, with ROC = 0.60 and 0.56.

We finalized the network embeddings generated by the Average operation as the optimal one and an ROC curve for the performance across 6 classification algorithms is shown in Figure 2.

Qualitative evaluation

We clustered the network embeddings by selecting the optimal hyperparameters $\epsilon = 1.5$ and $k = 37$ for the DBSCAN algorithm. Sixty-three clusters were generated with a Silhouette score of 0.128.

We used the network embeddings generated through the optimal Average operation to conduct a further qualitative evaluation. We first visualized clusters for diseases and clusters for all the entities as shown in Figure 3. In Figure 3A, we found that *pneumonia*, *fever*, *fibrosis*, and *bronchitis* appeared in each cluster, indicating that they are common comorbidities among all types of coronavirus infectious diseases. We also observed that COVID-19 co-occurred with *coronavirus*, *SARS*, *Ebola*, and *ZIKA* in different clusters respectively, denoting that there exists tremendous overlap between these diseases regarding underlying mechanisms. Clusters shown in Figure 3B further illustrate how different genes, mutations, and chemicals can link diseases with similar mechanisms. For example, COVID-19 was grouped in cluster #6, which contains 606 biomedical entities in total, including *infection of SARS*, *Ebola viruses*, *rs180047* mutation, and *carbohydrates* chemical component. Based on literature search, we found that *rs180047* is strongly related to *TGF- β 1*, a master regulator for *pulmonary fibrosis*, which is a common comorbidity related to COVID-19, *infection of SARS*, and *Ebola viruses*.²⁹ In addition, *carbohydrates*-based diagnostic was recently reported to be a potential new approach for testing COVID-19,³⁰ was also detected from cluster #6. The comprehensive information for entities included in each cluster was illustrated online (<https://github.com/shenfc/COVID-19-network-embeddings>).

We then selected 5 coronavirus infectious diseases and listed top 10 closest entities using cosine similarity as shown in Table 2. COVID-19 was clustered in cluster #6, and the top 2 closest entities in cluster #6 are *VP35* and *HD11* (both being genes). *VP35* is a virus protein of the other highly infectious disease *Ebola*.³¹ *HD11*, also known as *Homeobox* protein, is known for regulating infectious diseases such as *Avian infectious bronchitis virus (IBV)*³² that

is in the coronavirus family. *Pulmonary coronavirus infection* was grouped in cluster #1 and has the closest associations with gene *PTP* (*protein tyrosine phosphatase*), which has been mentioned in SARS-CoV replication inhibition studies.³³ In case of SARS-CoV-infected *human airway epithelia cell cultures*, it is easy to notice that the entity is directly linked to the coronavirus infection.³⁴ As for SARS-CoV *infection damages lung* that listed in cluster #2, *IL-1-alpha* (*interleukin 1Alpha*) also known as *IL-1 α* , has the closest association with SARS-CoV, and evidence could be found in a research study.³⁵ In particular, *IL-1 α* is a *proinflammatory cytokine* that shows increase when infected by SARS-CoV. *Sucralfate* is a chemical compound that also holds close association with SARS-CoV, which has been studied as a potentially effective means against *early-onset ventilator-associated pneumonia*.³⁶ *Coronavirus upper respiratory infection* was found in cluster #23, and *pleuropneumoniae* (disease) and *plasmin* (gene) are most 2 similar entities. *Pleuropneumoniae* is a *pneumonia* complicated with *pleurisy*, which has been linked to *porcine upper respiratory tract*³⁷ and *plasminogen (PLG)* (the *zymogen of plasmin*) has also been proved to be related with *coronavirus upper respiratory infection* in research studies.³⁸ *Coronavirus-infected pneumonia* was detected in cluster #10. The top 2 closest diseases are *respiratory syncytial viral infection* and *pegylated interferon-alpha*, which could be proved by reviewing research studies in Zou and Zhu³⁹ and Haagmans et al.⁴⁰

DISCUSSION

In this study, we used 11 keywords to query COVID-19-on-FHIR for constructing the COVID-19-centered coronavirus co-occurrence network. As research studies of COVID-19 published on a daily basis, we will keep watching the new results and adding more significant diseases and comorbidities into the keyword list to provide timely update for the co-occurrence network.

PMID information was not incorporated as an attribute into the co-occurrence network, which creates difficulties on the capability to trace each biomedical entity back to the original literature. In the future, we will add PMID list for each entity. On the one hand, it will help our evaluation on detecting if the closely associated terms are from the same article or different publications. On the other hand, it will provide more evidence to scientists and clinical investi-

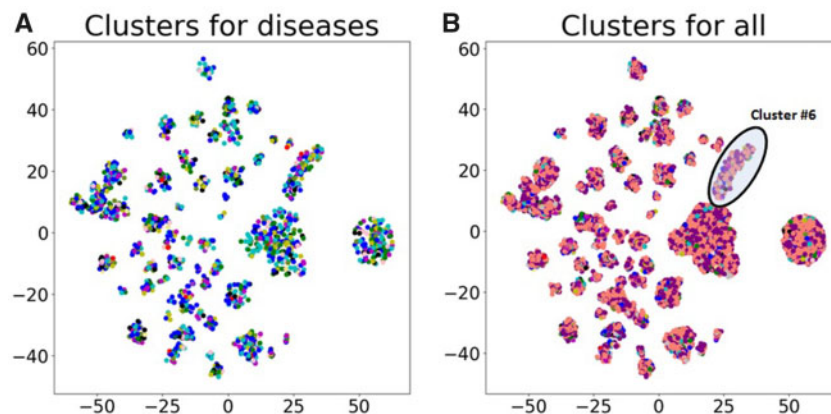


Figure 3. Clustering visualization for (A) diseases and (B) all the biomedical entities. COVID-19 (coronavirus disease 2019) is represented in red, SARS (severe acute respiratory syndrome) is represented in black, coronavirus is represented in green, pneumonia is represented in blue, fever is represented in cyan, fibrosis is represented in yellow, diarrhea is represented in magenta, bronchitis is represented in olive drab, Ebola is represented in pink, influenza is represented in dark orchid, ZIKA is represented in khaki, all the genes are represented in purple, all the mutations are represented in silver, and all the chemicals are represented in salmon.

Table 2. Top 10 intracluster closest biomedical entities for 5 selected coronavirus infectious diseases

Coronavirus infectious diseases	Top 10 closest entities	Cosine similarity score
COVID-19 (cluster #6)	VP35 (Gene)	0.9777
	HD11 (Gene)	0.9774
	Coronavirus infection process (Disease)	0.9700
	Fibroblast growth factor (FGF)-2 (Gene)	0.9655
	Acute respiratory infection illness (Disease)	0.9596
	PIGS (Gene)	0.9576
	TGF alpha (Gene)	0.9571
	SFPQ (Gene)	0.9561
	Tumor necrosis factor (TNF) (Gene)	0.9549
	Praziquantel (Chemical)	0.9537
Pulmonary coronavirus infection (cluster #1)	PTP (Gene)	0.9754
	SARS-CoV–infected human airway epithelia cell cultures (Disease)	0.9699
	“5′-tgg gat tca aca” (Chemical)	0.9672
	Trachea nasal respiratory epithelial cells and llamas (lama glama) (Disease)	0.9658
	Suppressor of cytokine signaling 3 (Gene)	0.9620
	KAT (Gene)	0.9604
	CD32 (Gene)	0.9573
	Maternal SARS infection (Disease)	0.9553
	Respiratory syndrome coronavirus (MERS-CoV) infections (Disease)	0.9547
	S27 (Gene)	0.9546
SARS-COV infection damages lung (cluster #2)	IL-1 α (Gene)	0.9560
	Sucralfate prn (Chemical)	0.9589
	Acute respiratory syndrome-cov infection (Disease)	0.9555
	IL-5– and IL-13–producing ilc-iis (Gene)	0.9487
	HAP1 (Gene)	0.9342
	FSK (Chemical)	0.9337
	Low fever (Disease)	0.9328
	HIV and Ebola virus infection (Disease)	0.9327
	YKL-40 (Gene)	0.9288
	ETF (Gene)	0.9280
Coronavirus upper respiratory infection (cluster #23)	Viruses <i>Actinobacillus pleuropneumoniae</i> (Disease)	0.9890
	Plasmin (Gene)	0.9719
	JAM-1 (Gene)	0.9654
	TNF receptor–associated factor 6 (Gene)	0.9648
	GPC3 (Gene)	0.9613
	Renin (Gene)	0.9582
	ZO-1 (Gene)	0.9563
	Cathepsin G (Gene)	0.9556
	rs5743313 (Mutation)	0.9547
	Alpha1 antitrypsin (Gene)	0.9544
Coronavirus-infected pneumonia (cluster #10)	Respiratory syncytial viral infection (Disease)	0.9923
	Pegylated interferon-alpha (Chemical)	0.9891
	IFITM6 (Gene)	0.9872
	Feline b (Chemical)	0.9858
	E119V (Mutation)	0.9854
	Epac2 (Gene)	0.9850
	GFTP2 (Gene)	0.9849
	Hepatitis coronavirus infection (Disease)	0.9843
	Ouabain (Chemical)	0.9797
	LY6G (Gene)	0.9786

Cluster ID and the type of entities are marked in parentheses.

gators for assisting their research studies on COVID-19 in an efficient manner.

We didn't consider the co-occurred frequency for the pairs of biomedical entities while training the network embeddings. Instead of treating each edge equally, in the future, we will add weights over edges using frequency in order to better represent associations in the network and provide more accurate edge embeddings to better quantify power of associations among biomedical entities.

In unsupervised learning approach, there is always a balance between the number of clusters and the silhouette score. In most cases, the silhouette score tends to be higher if the number of cluster is small.⁴¹ In this study, we used a heuristic way to determine silhouette score and the number of clusters for making clear separations over the biomedical entities. In the future, we sought to use our previously developed hierarchical clustering optimization algorithms to make dynamic balance between the optimal silhouette

score and suitable cluster density.^{42–44} Moreover, after checking top similar entities for 5 selected coronavirus infectious diseases, we observed that applying clustering over the network embeddings could detect both explicit and implicit associations. It is easy to check explicit links, as most of them might be documented in existing studies. But for implicit associations, although they might hold huge potential on new discoveries, in order to validate their correctness, we will invite a clinical investigator from Mayo Clinic Division of Pulmonary and Critical Care Medicine for manual evaluation.

CONCLUSION

This study has explored the construction of co-occurrence network embeddings for COVID-19 and related coronavirus infectious diseases. We have tested different edge embeddings operations along with different machine learning algorithms to optimize the final network embeddings and developed unsupervised clustering algorithms to deep dive into specific COVID-19-related associations. Results indicated that the co-occurrence network embeddings were able to perform link prediction task well and detect both explicit and implicit associations for COVID-19, demonstrating its potential usage for discovering new disease management and treatment plan for COVID-19. Detailed implementations and data sources could be found online (<https://github.com/shenfc/COVID-19-network-embeddings>). A Web-based user-friendly tool for clustering visualization and entity similarity check is available online (<https://www.davidoniani.com/covid-19-network>).

FUNDING

This work was supported by National Institute of Health grant U01TR0062-1 (PI: HL and Co-I: FS).

AUTHOR CONTRIBUTIONS

DO was responsible for algorithm development, tool implementation, and manuscript drafting. GJ was responsible for extracting needed information from CORD-19-on-FHIR. HL provided supports on system evaluations. FS conceived and supervised the study. FS was responsible for research topic formulation, algorithm development, experiment design, tool implementation, and manuscript writing and revision.

ACKNOWLEDGMENTS

We thank the FHIRCat team for building the CORD-19-on-FHIR datasets and providing support on query template design.

CONFLICT OF INTEREST STATEMENT

None declared.

REFERENCES

- Chen Q, Allot A, Lu Z, *et al*. Keep up with the latest coronavirus research. *Nature* 2020; 579 (7798): 193.
- Mihalcea R, Corley C, Strapparava C. Corpus-based and knowledge-based measures of text semantic similarity. In: AAAI'06: proceedings of the 21st National Conference on Artificial Intelligence; 2006: 775–80.
- Oliva J, Serrano JI, del Castillo MD, *et al*. SyMSS: a syntax-based measure for short-text semantic similarity. *Data Knowl Eng* 2011; 70 (4): 390–405.
- Heath T, Bizer C. Linked data: evolving the web into a global data space. In: Hendler J, ed. *Synthetic Lectures on the Semantic Web: Theory and Technology*. San Rafael, CA: Morgan & Claypool; 2011: 1–136.
- Ahamed S, Samad M. Information mining for COVID-19 research from a large volume of scientific literature. *arXiv:2004.02085*; 2020.
- Bellomarini L, Benedetti M, Gentili A, *et al*. COVID-19 and company knowledge graphs: assessing golden powers and economic impact of selective lockdown via AI reasoning. *arXiv:2004.10119*; 2020.
- Tsiotas D, Magafas L. The effect of anti-COVID-19 policies to the evolution of the disease: a complex network analysis to the successful case of Greece. *arXiv:2004.06536*; 2020.
- Allen Institute for AI. CORD-19: COVID-19 Open Research Dataset. <https://allenai.org/data/cord-19> Accessed May 20, 2020.
- Wolinski F. Visualization of diseases at risk in the COVID-19 Literature. *arXiv:2005.00848*; 2020.
- Wang X, Song X, Guan Y, Li B, Han J. Comprehensive named entity recognition on CORD-19 with distant or weak supervision. *arXiv:2003.12218*; 2020.
- CORD-19-on-FHIR. <https://github.com/fhircat/CORD-19-on-FHIR> Accessed April 6, 2020.
- Bender D, Sartipi K. HL7 FHIR: an Agile and RESTful approach to healthcare information exchange. In: proceedings of the 26th IEEE International Symposium on Computer-Based Medical Systems; 2013: 326–31.
- Miller E. An introduction to the resource description framework. *Bull Am Soc Inf Sci Technol* 2005; 25 (1): 15–9.
- Groza A. Detecting fake news for the new coronavirus by reasoning on the Covid-19 ontology. *arXiv:2004.12330*; 2020.
- Mikolov T, Wen-tau Yih S, Zweig G. Linguistic regularities in continuous space word representations. In: proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies; 2013: 746–51.
- Grover A, Leskovec J. node2vec: Scalable feature learning for networks. In: proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2016: 855–64.
- Shen F, Liu S, Wang Y, *et al*. Constructing node embeddings for human phenotype ontology to assist phenotypic similarity measurement. In: proceedings of the 2018 IEEE International Conference on Healthcare Informatics Workshop; 2018: 29–33.
- Shen F, Peng S, Fan Y, *et al*. HPO2Vec+: leveraging heterogeneous knowledge resources to enrich node embeddings for the human phenotype ontology. *J Biomed Inform* 2019; 96: 103246.
- Wei C-H, Kao H-Y, Lu Z, *et al*. PubTator: a web-based text mining tool for assisting biocuration. *Nucleic Acids Res* 2013; 41 (W1): W518–22.
- Tang L, Liu H. Leveraging social media networks for classification. *Data Min Knowl Disc* 2011; 23 (3): 447–78.
- Fortunato S. Community detection in graphs. *Phys Rep* 2010; 486 (3–5): 75–174.
- Maaten L, *et al*. Visualizing data using t-SNE. *J Mach Learn Res* 2008; 9: 2579–605.
- Sander J, Ester M, Kriegel H-P, *et al*. Density-based clustering in spatial databases: the algorithm GDBSCAN and its applications. *Data Min Knowl Discov* 1998; 2 (2): 169–94.
- Quinlan JR. Induction of decision trees. *Mach Learn* 1986; 1 (1): 81–106.
- Walker SH, Duncan DB. Estimation of the probability of an event as a function of several independent variables. *Biometrika* 1967; 54 (1–2): 167–79.
- Cortes C, Vapnik V. Support-vector networks. *Mach Learn* 1995; 20 (3): 273–97.
- Ho TK. Random decision forests. In: proceedings of the 3rd International Conference on Document Analysis and Recognition; 1995: 278–82.
- Rish I. An empirical study of the naive Bayes classifier. *IJCAI 2001 Workshop on Empirical Methods in Artificial Intelligence*; 2001: 41–6.
- Osadnik T, Strzelczyk JK, Lekston A, *et al*. The association of functional polymorphisms in genes encoding growth factors for endothelial cells and smooth muscle cells with the severity of coronary artery disease. *BMC Cardiovasc Disord* 2016; 16 (1): 218.

30. Carbohydrate-based diagnostics: a new approach to COVID-19 testing? <https://www.technologynetworks.com/diagnostics/news/carbohydrate-based-diagnostics-a-new-approach-to-covid-19-testing-332313> Accessed May 22, 2020.
31. Fajardo-Ortiz D, Ortega-Sánchez-de-Tagle J, Castaño VM, *et al.* Hege-monic structure of basic, clinical and patented knowledge on Ebola re-search: a US army reductionist initiative. *J Transl Med* 2015; 13 (1): 124.
32. Li H, Li J, Zhai Y, *et al.* Gga-miR-30d regulates infectious bronchitis virus infection by targeting USP47 in HD11 cells. *Microbial Pathog* 2020; 141: 103998.
33. Yap Y, Zhang XWu, Andonov A, *et al.* Structural analysis of inhibition mechanisms of aurintricarboxylic acid on SARS-CoV polymerase and other proteins. *Comput Biol Chem* 2005; 29 (3): 212–9.
34. Sheahan TP, Sims AC, Zhou S, *et al.* An orally bioavailable broad-spectrum antiviral inhibits SARS-CoV-2 in human airway epithelial cell cultures and multiple coronaviruses in mice. *Sci Transl Med* 2020; 12 (541): eabb5883.
35. Barnard DL, Day CW, Bailey K, *et al.* Enhancement of the infectivity of SARS-CoV in BALB/c mice by IMP dehydrogenase inhibitors, including ri-bavirin. *Antiviral Res* 2006; 71 (1): 53–63.
36. Bornstain C, Azoulay E, De Lassence A, *et al.* Sedation, sucalfate, and anti-biotic use are potential means for protection against early-onset ventila-tor-associated pneumonia. *Clin Infect Dis* 2004; 38 (10): 1401–8.
37. Sidibe M, Messier S, Lariviere S, Gottschalk M, Mittal KR. Detection of *Actinobacillus pleuropneumoniae* in the porcine upper respiratory tract as a complement to serological tests. *Can J Vet Res* 1993; 57: 204–8.
38. Wang J, Hajizadeh N, Moore EE, *et al.* Tissue plasminogen activator (TPA) treatment for COVID-19 associated acute respiratory distress syn-drome (ARDS): a case series. *J Thromb Haemost* 2020 Apr 8 [E-pub ahead of print].
39. Zou S, Zhu X. FDG PET/CT of COVID-19. *Radiology* 2020 Mar 6 [E-pub ahead of print].
40. Haagmans BL, Kuiken T, Martina BE, *et al.* Pegylated interferon- α pro-TECTS type 1 pneumocytes against SARS coronavirus infection in macaques. *Nat Med* 2004; 10 (3): 290–3.
41. Shen F, Lee Y. Knowledge discovery from biomedical ontologies in cross domains. *PLoS One* 2016; 11 (8): e0160005.
42. Shen F, Lee Y. Biobroker: Knowledge discovery framework for heteroge-neous biomedical ontologies and data. *J Intell Learn Syst Appl* 2018; 10 (1): 1–20.
43. Shen F, Liu H, Sohn S, Larson DW, Lee Y. BmQGen: Biomedical query generator for knowledge discovery. In: 2015 IEEE International Confer-ence on Bioinformatics and Biomedicine; 2015: 1092–7.
44. Shen F, Liu H, Sohn S, *et al.* Predicate oriented pattern analysis for biomedical knowledge discovery. *Intell Inf Manag* 2016; 8 (3): 66–85.