

MIT Open Access Articles

A Graph Neural Network Approach for Product Relationship Prediction

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Ahmed, Faez, Cui, Yaxin, Fu, Yan and Chen, Wei. 2021. "A Graph Neural Network Approach for Product Relationship Prediction." Volume 3A: 47th Design Automation Conference (DAC).

As Published: 10.1115/DETC2021-69462

Publisher: ASME International

Persistent URL: <https://hdl.handle.net/1721.1/150663>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



A Graph Neural Network Approach for Product Relationship Prediction

Faez Ahmed¹, Yaxin Cui², Yan Fu³, and Wei Chen²

¹Dept. of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA

²Dept. of Mechanical Engineering, Northwestern University, Evanston, IL

³Insight and Analytics, Ford Motor Company, Dearborn, MI

Paper accepted in ASME IDETC 2021

Abstract

Graph Neural Networks have revolutionized many machine learning tasks in recent years, ranging from drug discovery, recommendation systems, image classification, social network analysis to natural language understanding. This paper shows their efficacy in modeling relationships between products and making predictions for unseen product networks. By representing products as nodes and their relationships as edges of a graph, we show how an inductive graph neural network approach, named GraphSAGE, can efficiently learn continuous representations for nodes and edges. These representations also capture product feature information such as price, brand, or engineering attributes. They are combined with a classification model for predicting the existence of the relationship between products. Using a case study of the Chinese car market, we find that our method yields double the prediction performance compared to an Exponential Random Graph Model-based method for predicting the co-consideration relationship between cars. While a vanilla GraphSAGE requires a partial network to make predictions, we introduce an ‘adjacency prediction model’ to circumvent this limitation. This enables us to predict product relationships when no neighborhood information is known. Finally, we demonstrate how a permutation-based interpretability analysis can provide insights on how design attributes

impact the predictions of relationships between products. This work provides a systematic method to predict the relationships between products in many different markets.

1 Introduction

Complex engineering systems contain multiple types of stakeholders and many individual entities, which exhibit complex interactions and interconnections. An example of a complex engineering system is the car market, where there are many interactions between stakeholders. The success of a new car depends not only on its engineering performance but also on the car’s competitiveness relative to similar cars and factors such as perceived market position. Customers from different geographies may prefer different types of cars, while a design intervention in the car market, either by introducing changes in existing cars or by launching a new design of car may encourage customers to change their driving behavior. To solve this complexity, it is necessary to consider the complex relationship between customers and products, such as the social network between customers and the competitive relationship between products.

Network analysis has emerged as a key method for statistical analysis of engineering systems in a wide variety of scientific, social, and engineering do-

mains [1, 2, 3, 4, 5, 6]. A few studies have begun exploring the capability of statistical network models in modeling complex customer-product relationships [7, 8, 9]. The premise underlying the network-based approach is that, similar to other engineering systems exhibiting dynamic, uncertain, and emerging behaviors, the relationship between customers and products can be viewed as a complex socio-technical system and analyzed using social network theories and techniques. The structural and topological characteristics identified in customer-product networks can reveal emerging patterns of the customer-product relations while taking into account the heterogeneity among customers and products.

Exponential random graph models (ERGMs) have been employed in literature as a statistical inference framework to interpret complex customer-product relations. ERGMs were used to study customers' consideration behaviors using a unidimensional network at the aggregated market level [10] and multidimensional network at the disaggregated customer level [11], respectively. In unidimensional models, the product competition was established based on the customers consideration behaviour. The estimated unidimensional model was used to forecast the impact of technological changes (*e.g.* turbo engines) on market competition [12], which illustrated the benefits of using the network-based preference model for predicting the outcome of design decisions.

However, ERGMs have a few limitations. First, they are typically appropriate for small- to medium-sized networks with a few attributes. For large datasets, the MCMC approach to estimate ERGM parameters does not converge [13]. This leads to an important limitation for product manufacturers, who now want to make the most of huge datasets but still want statistical models that can help them understand what is happening under the hood. In addition, previously published research shows that future market forecasts based on ERGMs are not sufficiently accurate at capturing the true network [14]. In this paper, we provide an alternative approach of modeling networks using neural networks, which does not face these issues.

Graph neural network (GNN) are increasingly gaining popularity, given their expressive power and

explicit representation of graphical data. Hence, they have a wide range of applications in domains that can harness graph structures out of their data. They offer fundamental advantages over more traditional unstructured methods in supporting interpretability, causality, and inductive generalization. Learning graph representations and performing reasoning and prediction has achieved impressive progress in applications ranging from drug discovery [15], image classification, natural language processing and social network analysis [16]. A few of the well-known applications of GNNs are Uber Eats [17], who used them to recommend food items and restaurants and Alibaba using them to model millions of nodes for product recommendation [18]. These successes motivated us to use them for studying product relationships.

We demonstrate a GNN approach for predicting product relationships. In our approach, the products are viewed as nodes, and the relationship among them (product association, market competition) is viewed as links. Hence, the problem of predicting relationships between products is posed as a graph link (or edge) prediction problem. The new approach we develop in this study is based on GraphSAGE, a type of GNN method, which allows modeling of design attributes. GraphSAGE first represents a graph (network) structure in lower-dimension vectors and utilizes the vectors as the downstream classification input. Meanwhile, we develop a permutation-based method to examine the feature importance to assist design decisions. In summary, the contributions of this study are:

1. Propose a GNN-based method for modeling a product relationship network and enabling a systematic way to predict the relationship links between unseen products for future years.
2. Show that the link prediction performance of GNNs is better than existing network modeling methods.
3. Demonstrate the scalability of the GNN method by modeling the effect of a large number of continuous and categorical attributes on link prediction.

4. Uncover the importance of attributes to help make design decisions using permutation-based methods.

2 Related Work

This paper applies GNNs to product relationship networks for link prediction and uncovers the importance of engineering design attributes using permutation-based analysis. In this work, we focus on the product co-consideration relation as a demonstration, but the method can be generalized to other product relationships, such as product association relationship. Below, we discuss related work on product co-consideration networks, GNNs, and interpretable machine learning.

2.1 Product Co-consideration Networks

Co-consideration of products describes the situation where customers consider multiple products at the same time prior to making a purchase [19]. The consideration behavior involves the comparison and evaluation of product alternatives and is accordingly a key step in the customer’s decision-making process [20]. At the same time, product co-consideration also means a market competition relationship between products, which is crucial to the company’s product positioning plans and market strategies. As a single product may be chosen by a customer considering two or more products, those products can increase their market share by understanding competition relationships and introducing interventions for them to be preferred over their competitors. Therefore, the successful modeling of the product co-consideration relationship helps enterprises understand the embedded market competition and provides new opportunities for enterprises to formulate design solutions to meet customer needs.

In order to understand the underlying patterns of customer consideration behaviors, researchers have developed multiple methods and models of customer considerations. Some models of customer consideration set composition are based on the marginal ben-

efits of considering an additional product [21, 22]. Other pioneering works have built models for investigating the impact of the consideration stage on the customer decision-making process [23, 24]. Also, in the Big Data era, many works use both online and offline customer activity data to infer the product co-consideration behavior [25]. In recent years, the network-based approach has emerged to understand the product competition by describing the product co-consideration relation based on customer cross-shopping data [19, 14, 13]. Depicted in a simple network graph, where nodes represent individual products and edges represent their co-consideration relation based on aggregated customer preference, network-based analysis views co-consideration relations in terms of network theories, and the links in the observed network are explained by the underlying social processes.

Several works that investigate the product co-consideration network are based on a dataset of car purchases. Wang *et al.* [26] have applied the correspondence analysis methods and network regression methods to investigate the formation of car co-consideration relationship. Sha *et al.* [14] have applied ERGMs in understanding the underlying customer preference in car co-consideration networks. However, the previous explorations are restricted to using the traditional network-based statistical methods, which leads to a low computation efficiency, low prediction accuracy for the future market competition as well as inability to model many design attributes. To overcome the limitations of the ERGMs, we have developed a new method to investigate the underlying effect of customers’ consideration behavior by using GNN methods. Applied to the same dataset, a comparison of the ERGMs and this work is summarized in Table 1.

2.2 Graph Neural Networks

Network data can be naturally represented by a graph structure that consists of nodes and links. Recently, research on analyzing graphs with machine learning has grown rapidly. The graph-based machine learning tasks in networks include **node classification** (predict a type of given node), **link prediction** (pre-

Table 1: Comparison of this work with prior studies on modeling car relationship using ERGM models

Topic	Past work using ERGM model	This work using GNN model
Train nodes	296 cars (common cars between 2013 and 2014)	388 cars (all cars from 2013)
Test nodes	Tested on 296 cars from 2014	Tested on 403 cars from 2014 and 422 cars from 2015
Unseen data	Predictions restricted to cars in the training data	Predictions for completely new cars too (107 unseen cars in 2014)
Attributes	6 numerical design attributes	29 design attributes, including categorical attributes
Interpretability	Coefficient-based	Permutation-analysis based

dict whether two nodes are linked), **community detection** (identify densely linked clusters of nodes), **network similarity** (determine how similar two networks are), **anomaly detection** (find the outlier nodes), and **attribute prediction** (predict the features of a node) [27].

In a graph, each node is naturally defined by its features and the neighborhood of connected nodes. Therefore, learning the representation of nodes in a graph, called **node embedding**, is an essential part of downstream tasks such as classification and regression. Most node embedding models are based on spectral decomposition [28, 29] or matrix factorization methods [30, 31]. However, most embedding frameworks are inherently transductive and can only generate embeddings for a single fixed graph. These transductive approaches do not efficiently generalize to unseen nodes (e.g., in evolving graphs), and these approaches cannot learn to generalize across different graphs. In contrast, **GraphSAGE** is an inductive framework that leverages node attribute information to efficiently generate representations on previously unseen data. GraphSAGE samples and aggregates features from a node’s local neighborhood [32]. By training a GraphSAGE model on an example graph, one can generate node embeddings for previously unseen nodes as long as they have the same attribute schema as the training data. It is especially useful for graphs that have rich node attribute information, which is often the case for product networks.

2.3 Interpretable Machine Learning

In addition to using machine learning models for prediction, there is growing attention on the capability to interpret what a model has learned. Interpretable machine learning would be an effective tool to explain or present the model results in terms understandable to humans [33, 34].

As a traditional machine learning explanation method, feature importance indicates the statistical contribution of each feature to the underlying model [35]. Among the techniques to unravel the feature importance, model-agnostic interpretation methods [36] treat a model as a black-box and do not inspect internal model parameters, which has the advantage that the interpretation method can work with any machine learning model. A representative approach is the **permutation feature importance measurement**, which was introduced by Breiman [37] for random forests. Based on this idea, Fisher *et al.* [38] developed a model-agnostic version of the feature importance and called it model reliance. The key idea is that the importance of a specific feature to the overall performance of a model can be determined by calculating how the model prediction accuracy deviates after permuting the values of that feature [39]. The permutation-based feature importance method has been applied to bioinformatics [40], engineering [41], and political science [42] to provide a highly compressed, global insight into machine learning models. In our study, we use the permutation-based methods to examine important

product attributes that impact the link prediction between cars.

3 Methodology

We establish a product co-consideration network to model product competition behavior and use a GNN approach to predict future product competition. The methodology of the training and prediction process for the link existence is shown in Fig.1.

Our methodology comprises five main components, which include representing products and their relationships as a graph, training the GNN to learn the graph structure, training classification models to make predictions, creating an adjacency prediction model to augment the GNN for unseen data, and, finally, interpreting the importance of design attributes. These components are described next:

3.1 Network Construction

Networks present a natural way to simultaneously model products as nodes and relationships between them as edges. Before purchasing a product, customers often consider multiple products and then select one or more products among them. When two products are simultaneously considered by the same customer in their decision-making process, we define this relationship as a co-consideration relationship. Assuming the customer only buys one product in the end, products that are co-considered are assumed to be in competition in this paper. Note that there are many different methods to measure competition between any two products, and the methods we describe next generalize to any measure of choice. Next, we discuss how a graph is created for co-considered products. Readers who already have a predefined network using some other method can skip this section.

We studied a unidimensional product network that can reveal product market competition by describing products’ co-consideration relationship. Each product corresponds to a unique node. Each node is associated with a set of attributes such as price, fuel consumption, and engine power. The product co-consideration network is constructed using data from

customers’ consideration sets. The presence of a co-consideration binary link between two nodes (products) is determined by the number of customers who consider them together:

$$E_{i,j} = \begin{cases} 1, & n_{i,j} \geq cutoff \\ 0, & otherwise \end{cases} \quad (1)$$

where $E_{i,j}$ refers to the edge connected by node i and node j . $n_{i,j}$ is the number of customers who have considered products i and j together. *cutoff* is a domain-dependent threshold, which defines the strength of the relationship considered in the analysis. In other words, we define an undirected link between node i and node j , if there exists at least one customer who considers both products i and j together. Based on Equation 1, the network adjacency matrix is symmetric and binary. This study uses a cut-off value equal to 1.

3.2 Inductive Representation Learning on Networks

Finding a low-dimensional vector embedding of nodes and edges in graphs can enable many machine learning applications such as node classification, clustering, and link prediction. This section describes a method to learn such embeddings, named GraphSAGE. GraphSAGE is a representation learning technique for dynamic graphs. It can predict the embedding of a new node, without needing a re-training procedure. To do this, GraphSAGE uses inductive learning. It learns aggregator functions that can induce new node embedding, based on the features and neighborhood of the node.

As illustrated in Fig. 2, GraphSAGE learns node embeddings for attributed graphs (where nodes have features or attributes) through aggregating neighboring node attributes. The aggregation parameters are learned by encouraging node pairs co-occurring in short random walks to have similar representations. Many GNN models learn functions that generate the embeddings for a node, which sample and aggregate feature and topological information from the node’s neighborhood. However, the benefit of training a




Figure 1: The methodology of predicting the link existence in a car competition network using graph neural network model

GraphSAGE model, in contrast to other GNN methods, is its inductive behavior, which is necessary for engineering applications. Most other GNN methods are transductive, which means they can only generate embeddings for a single fixed graph. If a completely new product comes up in a dynamically evolving graph, these transductive approaches cannot generalize to such unseen nodes. In contrast, GraphSAGE is an inductive method that leverages the attribute information of a new node to efficiently generate representations on previously unseen data. The detailed algorithm of GraphSAGE is shown in Algorithm 1. Interested readers are encouraged to read [32] for details of the algorithm.

To train a GraphSAGE model, the inputs are the product attributes (*i.e.* node features) and the network structure (*i.e.* adjacency matrix) of the product co-consideration network. Then for each node, the GNN models are able to encode nodes into lower-dimensional space in the node embedding stage. For example, as illustrated in Fig.1, nodes i and j can be represented by vectors i and j , which carry the information of node i 's and j 's features and local neigh-

borhoods, respectively.

Edge embeddings Using a GNN-trained embedding for nodes, one can also learn the representation for all possible links (edges) in the network. This is done by aggregating every possible pair of node embeddings. We use the dot product of vectors i and j to find the edge embeddings. Note that other symmetric operations such as addition can also be used to aggregate two node embeddings to give an edge embedding. In our experiments, we did not find significant differences among aggregation methods on the final prediction performance. After training, edges between similar nodes are expected to be closer to each other in the edge embedding space.

Once we learn the edge embeddings, they can be used as an input to any machine learning model, which can be trained to predict whether an edge exists or not, which is discussed next.

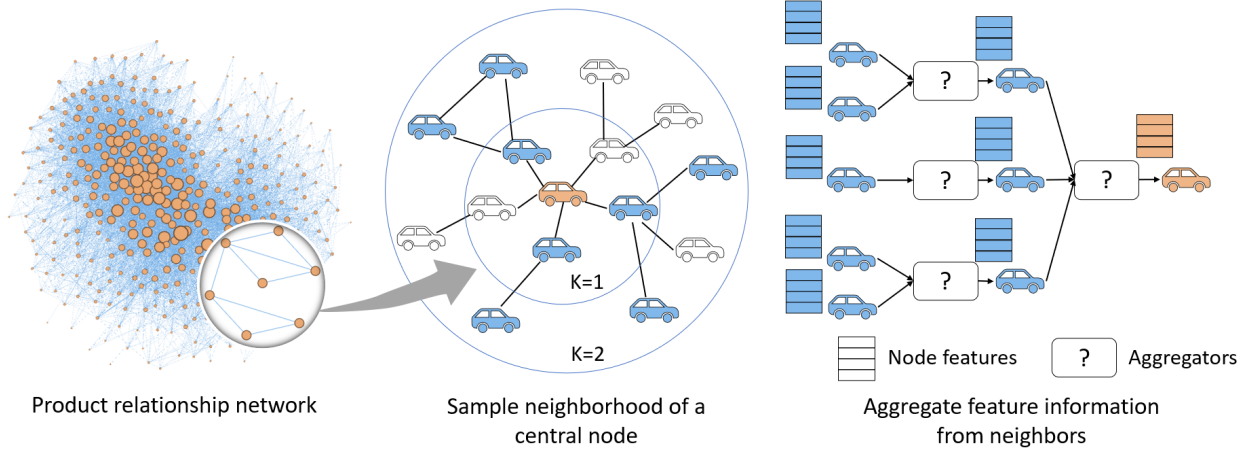


Figure 2: Illustration of sampling and aggregation in GraphSAGE method. A sample of neighboring nodes contributes to the embedding of the central node.

Algorithm 1: GraphSAGE embedding generation (i.e., forward propagation) algorithm from [32]

Input : Graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$; input features $\{\mathbf{x}_v, \forall v \in \mathcal{V}\}$; depth K ; weight matrices $\mathbf{W}^k, \forall k \in \{1, \dots, K\}$; non-linearity σ ; differentiable aggregator functions $\text{AGGREGATE}_k, \forall k \in \{1, \dots, K\}$; neighborhood function $\mathcal{N} : v \rightarrow 2^{\mathcal{V}}$

Output: Vector representations \mathbf{z}_v for all $v \in \mathcal{V}$

```

1  $\mathbf{h}_v^0 \leftarrow \mathbf{x}_v, \forall v \in \mathcal{V}$ ;
2 for  $k = 1 \dots K$  do
3   for  $v \in \mathcal{V}$  do
4      $\mathbf{h}_{\mathcal{N}(v)}^k \leftarrow \text{AGGREGATE}_k(\{\mathbf{h}_u^{k-1}, \forall u \in \mathcal{N}(v)\})$ ;
5      $\mathbf{h}_v^k \leftarrow \sigma(\mathbf{W}^k \cdot \text{CONCAT}(\mathbf{h}_v^{k-1}, \mathbf{h}_{\mathcal{N}(v)}^k))$ 
6   end
7    $\mathbf{h}_v^k \leftarrow \mathbf{h}_v^k / \|\mathbf{h}_v^k\|_2, \forall v \in \mathcal{V}$ 
8 end
9  $\mathbf{z}_v \leftarrow \mathbf{h}_v^K, \forall v \in \mathcal{V}$ 

```

3.3 Classification Model for Link Prediction

The link prediction problem can be viewed as a binary classification problem, where the goal is to predict whether a link candidate exists in the network (Class 1 or a positive edge) or does not exist (Class 0 or a negative edge). During the GNN model training, we can also train a downstream classification model to predict link existence, given the edge embedding as an input.

For each pair of nodes, the classification model takes the edge embeddings as input and whether the link exists or not as labels. Any classification model, such as logistic regression, k-nearest neighbors, and naive Bayes classifiers, can be integrated with the GNN model to predict the link existence. We used a multilayer perceptron (MLP) model for this work. Note that in the training process, the GNN model and the classification model are trained simultaneously for the supervised learning task. To avoid imbalanced training of the classification task for networks with very few edges, the two classes are balanced by sub-sampling the negative edges (an edge which does not exist in the training data).

3.4 Validation Networks

After the training was completed, we tested the performance of the model in predicting links for an unseen network. The model can be tested on two different types of networks. In one case, the initial network was divided into two parts by randomly sampling edges. The GNN model was tested to predict links for the held-out links. In the second case, we trained the model on one network and tested it on another completely unseen network. However, this presents new challenges, which are discussed next.

3.5 Adjacency Prediction Model

While GNN-based link prediction methods are typically used to find missing links from a graph, they cannot be directly applied to a completely unknown network. However, in engineering design applications, it is possible to train a model on products in Year 1 and make predictions about Year 2, which may have new products and evolved versions of previous products. Applications may require that predictions for links between products are made where training and testing networks belong to different domains, time periods, or locations. This presents a circularity problem, as a typical GNN, including GraphSAGE, needs at least a partial adjacency matrix as an input, to predict the complete adjacency matrix.

We overcame this issue by developing a method to predict an approximate adjacency matrix using a separate machine learning model, which is referred to as the adjacency prediction model in Fig.1. The predicted adjacency is used to identify a few neighbors of each node, which are used in the GNN as a partial adjacency matrix.

There are several ways of predicting the adjacency matrix, given the node attributes. A naïve way would be to find all the nodes in the new graph, which also appeared in the training dataset, and copy their adjacency information. However, such a model performs poorly, as all the new nodes have no neighbors, therefore, the GNN cannot make accurate predictions about them.

Instead, we used a similarity-based K-nearest neighbor method in the adjacency prediction model.

The similarities among product nodes are measured by the cosine distance of all car features. Using these similarities for each node, K most similar nodes from the graph are selected as neighbors. This gives us the approximate adjacency matrix, where each node is connected to its K-nearest neighbors. The benefit of this approach is that all nodes in the co-consideration network are connected to some other nodes. While the choice of K is subject to the modeler, we seek an appropriate number to keep the density of the network comparable with a typical co-consideration product network in the training network.

Note that other machine learning methods can also be used to output an approximate adjacency matrix. For instance, one can train a classification model with the average car attributes as input and a binary output corresponding to link existence. Our preliminary analysis showed that classification models (*e.g.* logistic regression) did not perform as well as the nearest-neighbor approach. This may be attributed to classification models not finding sufficient neighbors for all nodes. Our method overcame this limitation by assigning the same number of neighbors to all nodes, which also yields good empirical results.

3.6 Metrics for Link Prediction

With the trained GNN model and classification model, we predicted the co-consideration network in the subsequent years based on the new node features and the approximate adjacency prediction model. The link prediction can be regarded as a binary classification model, which predicts the probability of the target link existence to be Yes or No. To evaluate the performance of the classification model, we analyzed the confusion matrix (which describes the performance of a classifier) and the receiver operating characteristic (ROC) curve, which plots the true positive rate and false positive rate). To compare different models, we used the area under the curve (AUC) metric, which measures the area underneath the ROC curve, and provides an aggregated measure of the performance across all possible classification thresholds. The AUC ranges in value from 0 to 1, a higher AUC value indicates a better classification model.

3.7 Permutation-based Feature Importance

In the engineering design domain, besides forecasting the future market competition, it is important to understand the dominant features in product competition. Therefore, we investigated the importance of different design attributes in the GNN method.

Feature importance is the increase in model error when the feature’s information is destroyed. Permutation feature importance measures the increase in the prediction error of the model after we permuted the feature’s values, which breaks the relationship between the feature and the true outcome. We measured the importance of a feature by calculating the increase in the model’s prediction error after permuting the feature. A feature is “important” if shuffling its values increases the model error, because in this case the model relied on the feature for the prediction. A feature is “unimportant” if shuffling its values leaves the model error unchanged, because in this case the model ignored the feature for the prediction.

Feature importance based on the training data tells us which features are important for the model in the sense that it depends on them for making predictions. Permutation feature importance does not require retraining the model. Some other methods suggest deleting a feature, retraining the model, and then comparing the model error. Since the retraining of a machine learning model can take a long time, “only” permuting a feature can save a lot of time. Permutation feature importance is linked to the error of the model.

Outline of the permutation importance algorithm:

1. Inputs: fitted predictive model m , tabular dataset (training or validation) D .
2. Compute the reference score s of the model m on data D using all features
3. For each feature j (column of D):
 - (a) For each repetition k in $1, \dots, K$
 - i. Randomly shuffle column j of dataset D to generate a corrupted version of the data named $\tilde{D}_{k,j}$

- ii. Compute the score $s_{k,j}$ of model on corrupted data $\tilde{D}_{k,j}$

- (b) Compute importance i_j for feature f_j defined as: $1 - \frac{s_{k,j}}{s}$

4 Results and Discussion

In this section, we demonstrate the use of the GNN approach to study the Chinese car market. We use car survey data provided by the Ford Motor Company as a test example. By training a network model, we can predict the future market competition even though car attributes are changing and new products are introduced. Meanwhile, the feature importance in the car competition network is examined for the training network, which can be reported back to designers to make strategic design changes.

4.1 Data Description

Our dataset contains customer survey data from 2012 to 2016 in the China market. In the survey, there were more than 40,000 respondents each year who specified which cars they purchased and which cars they considered before making their final car purchase decision. Each customer indicated at least one and up to three cars which they considered. The dataset resulting from the survey also contains many attributes for each car (*e.g.* price, power, brand origin, and fuel consumption) and many attributes for each customer (*e.g.* gender, age).

4.2 Link Prediction for Car Co-Consideration Network

In this example, we used our method to build a model that predicts co-consideration links in the car dataset aforementioned. The problem is treated as a supervised link prediction problem on a homogeneous network with nodes representing cars (with attributes such as engine size and categorical body type) and links corresponding to car-car co-consideration relationship.

Network construction To study car co-consideration, we started by creating a car co-consideration network based on customers’ survey responses in the 2013 survey data. The network consists of 388 unique car models as network nodes. The link between a pair of nodes (denoting cars) is all allocated based on the presence of the car co-consideration by at least one customer.

The input car attributes As demonstrated in the Methodology section, the car attributes and co-consideration network adjacency matrix serve as the input of the GNN and classification models, and the link existences are labels to judge the training performance. Our experiment studied 29 car attributes, which were manually chosen. The list of attributes contain all the effective engineering attributes (*e.g.* fuel consumption, engine size) and car types (*e.g.* body type, market segmentation) available in the survey dataset. The attributes are listed in Table 6. Note that the attributes are both continuous and categorical. The categorical variables are transformed via a one-hot encoder which converts categorical variables into vectors (after one-hot encoding 29 features lead to 210 features), and the continuous variables are normalized to vary between 0 and 1.

Experimental settings In the training process, we built a model with the following architecture. First, we built a two-layer GraphSAGE model that takes labeled node pairs corresponding to possible co-consideration links, and outputs a pair of node embeddings for the two cars of the pair. These embeddings were then fed into a link classification layer, which first applied a binary operator to those node embeddings (dot product) to construct the embedding of the potential link. The thus-obtained link embeddings are passed through the dense link classification layer to obtain link predictions - the probability for these candidate links to actually exist in the network. The entire model was trained end-to-end by minimizing the loss function of choice (*e.g.*, binary cross-entropy between predicted link probabilities and true link labels, with true/false links having

labels 1/0) using stochastic gradient descent (SGD) updates of the model parameters, with minibatches of training links fed into the model.

We specified the minibatch size (number of node pairs per minibatch) and the number of epochs for training the model to 20 and 100. As for the number of sampled neighbors, we set the sizes of 1-hop and 2-hop neighbor samples for GraphSAGE to be 20 and 10. For the GraphSAGE part of the model, we selected hidden layer sizes of 20 for both GraphSAGE layers and a bias term, and a dropout rate of 0.3. We stacked the GraphSAGE and prediction layers and defined the binary cross-entropy as the loss function. Parameters were chosen based on initial analysis on a validation set. Our code will be made public on Github for other researchers to replicate our results.

Predicting missing links in the 2013 network

We split our input graph into a training graph and a test graph. We used the training graph for training the model (a binary classifier that, given two nodes, predicts whether a link between these two nodes should exist or not) and the test graph for evaluating the model’s performance on hold-out data. Each of these graphs will have the same number of nodes as the input graph, but the number of links will differ (be reduced) as some of the links will be removed during each split and used as the positive samples for training/testing the link prediction classifier.

The prediction performance along with the training performance is first measured by a confusion matrix in Table 2. The right-hand part of Table 2 shows the confusion matrix of 2013 prediction on held-out links. It includes 4 different combinations of predicted and actual classes. The 609 in the top-left cell is the true negative (predict negative and it’s true), the 502 in the top right is the false positive (predict positive and it’s false). The associated percentages indicate that for all pairs of nodes without link existence (actual class = 0), 54.82% are predicted correctly whereas 45.18% are not. Meanwhile, the 75 in the bottom left is the false negative (predict negative and it’s false), and 1036 in the bottom right is the true negative (predict negative and it’s true),

Table 2: Confusion matrix in predicting 2013 with 29 features. Average F1-score for 2013 is 0.74. AUC for 2013 train is 0.84 and test is 0.84. True Negative Rate (TNR) and True Positive Rate (TPR) are shown in brackets.

Actual Class	2013 training prediction		2013 test prediction on held-out links	
	0	1	0	1
0	5390 (TNR 53.90%)	4610 (FPR 46.10%)	609 (TNR 54.82%)	502 (FPR 45.18%)
1	592 (FNR 5.92%)	9408 (TPR 94.08%)	75 (FNR 6.75%)	1036 (TPR 93.25%)

which suggested that for all pairs of nodes with link existence (actual class = 1), 93.25% are predicted correctly while 6.75% are not. We further calculate other evaluation metrics to quantify classification performance. The F1 score, which measures the test accuracy in an unbalanced class, was 0.74 for the predicted missing links (the range of the F1 score was [0, 1]), while the AUC was 0.84 for both training set and held-out test set. Basically, the AUC tells how capable the model is when distinguishing between classes. The higher the AUC, the better the model is. The problem of over-fitting is avoided because the AUCs for both training set and test set are comparable.

Predicting entire network for 2014 Once the trained model is converged, the learned parameters for the GNN model and the classification model can be used to predict the co-consideration network in the following years. As a test dataset, the car co-consideration network in 2014 is predicted. First, the 2014 car models set, which have an intersection with the 2013 car set and also have newly emerged cars, acts as the input of the prediction process without any link information. Then, through the adjacency prediction model, an approximate adjacency matrix based on the similarities of nodes is generated. Next, the node features and approximate prediction model are fed into the GNN model and followed by the classification model, the link existence of each pair of nodes is forecasted with a certain probability threshold.

Likewise, we computed the confusion matrix for

the predicted 2014 co-consideration network in Table 3, and calculated the F1 score as 0.65. Furthermore, we scoped out the AUC-ROC curve (in Fig. 3) at various threshold settings. The overall AUC is 0.80.




Figure 3: AUC-ROC curve to predict 2014 co-consideration network with 6 attributes and 29 attributes

Predicting entire network for 2015 Hitherto, we have predicted the 2014 co-consideration network based on the training data in 2013. However, as 2014 succeeded 2013, the market structure did not change

Table 3: Confusion matrix in predicting 2014 and 2015 with 29 features. F1-score for 2014 is 0.65 and 0.65 for 2015. AUC for 2014 is 0.80 and 0.80 for 2015

		2014 test prediction on unseen network		2015 test prediction on unseen network	
		0	1	0	1
Actual Class	0	42633 (TNR 61.73%)	26435 (FPR 38.27%)	45735 (TNR 61.28%)	28893 (FPR 38.72%)
	1	1811 (FNR 15.17%)	10124 (TPR 84.83%)	2195 (FNR 16.43%)	11167 (TPR 83.57%)

dramatically. Among 389 cars in 2013 and 403 cars in 2014, there are 296 cars in common. Therefore, to further assess the prediction capability for the model, we predict the 2015 co-consideration network using the trained model (2013 training data) with the car attributes and similarity-based adjacency matrix.

The predicted results are recorded and evaluated in Table 3 and Fig. 3, where the F1 score is 0.65 and AUC is 0.80. Compared to the prediction results in 2014, the prediction in 2015 maintains an equivalent performance, which is an indication of model robustness.

Predicting with six attributes In order to make a fair comparison of the previous traditional statistical network models (e.g., ERGMs), we used the same set of input attributes (only 6 attributes in [13]) and compared the AUC of each model. Besides, as previous studies used a subset of cars and did not make predictions for newly emerged car models, we took the intersection of 2013 and 2014 cars (296 cars in total) for our analysis.

When only six car features were utilized in the training and prediction model, we obtained the prediction results for 2014 data for GNN and ERGM, respectively, in Table 4. In the confusion matrix, we observed that in ERGM, the true positive rate (the ratio of true positive to all actual positive) is 79.81% and the true negative rate (the ratio of true negative to all actual negative) is 40.51%. Both of the values are lower than those predicted by GNN. Furthermore, the F1 score of the ERGM is merely 0.31, which is almost half the F1 score of the GNN

model. The AUC for ERGM prediction is 0.68, which is also less than the corresponding value of 0.78 for the GNN model. All of the evidence suggested that the prediction model of GNN performs better than the traditional statistical network models.

Then we summarized all the AUCs to compare in Table 5. Notice that the ERGM with 29 attributes does not associate with an AUC value because the model does not get converged with so many attributes. Meanwhile, we did not run the six attributes prediction for the 2015 data on the GNN because the common car set for 2013 and 2014 is no longer suitable for the 2015 car market. It is apparent from the comparison that the GNN models perform better than the ERGM model with a higher AUC and F1 score, and GNN models can accommodate larger networks more design attributes and introduction of unseen nodes in the study of product relationship.

4.3 Interpretability of Attributes

To inspect the feature importance, we applied the permutation method to find the decrease in a model score when a single feature value is randomly shuffled. Permutation importance calculation repeats the process with multiple shuffles to ensure the accuracy. We ran 50 permutations for each feature in the training data and calculated the drop in performance. The results are shown in Table 6. We found that the make of the car, the body type, and the segment are the most important attributes for the GNN to predict ties.

Table 6 shows that 14 of the 29 attributes have no positive effect on the model prediction. Note that

Table 4: Confusion matrix in predicting 2014 with six features and 296 cars for using the GNN method and the ERGM method. F1 score is 0.60 for the GNN model and 0.31 for the ERGM model, and the AUC is 0.78 for the GNN model and 0.68 for the ERGM model.

		2014 prediction class GNN		2014 prediction class ERGM	
		0	1	0	1
Actual Class	0	20336 (TNR 54.95%)	16675 (FPR 45.05%)	14993 (TNR 40.51%)	22018 (FPR 59.49%)
	1	867 (FNR 13.04%)	5782 (TPR 86.96%)	1384 (FNR 20.82%)	5265 (TPR 79.18%)

Table 5: Comparing train AUC and test AUC in different years, different models and different sets of attributes. AUC in link prediction. The goal is to predict the entire network (all existing and non-existing edges) in a 0/1 classification task

Number of attributes	Train AUC (2013)	Test AUC (2014)	Test AUC (2015)	Test AUC (ERGM)
29 attributes	0.84	0.80	0.80	NA
Six attributes	0.81	0.78	NA	0.68

negative values are returned when a random permutation of a feature’s values results in a better performance metric compared to the performance before a permutation is applied. This means the model does not rely on feature’s which have negative values when predicting links for the training data. We observe that most continuous values, such as engine size, price, fuel consumption, and power do not have high importance.

It is noteworthy that the permutation methods on feature importance can be applied to either training data or test data. In the end, one needs to decide whether one wants to know how much the model relies on each feature for making predictions (training data) or how much the feature contributes to the performance of the model on unseen data (test data).

5 Implications for Design

A car is an expensive commodity, and customers usually consider multiple options before deciding which car to buy. This decision may be influenced by many factors, such as the customer’s budget, driving needs,

required and necessary features, the popularity of nearby car models, brand, past experience, the influence of cars owned or recommended by family and friends, *etc.*

From a manufacturer’s perspective, it is important to understand the market competition and develop strategies to improve their market share. The proposed model can support manufacturers in the following aspect:

First and foremost, the prediction capability of the GNN model facilitates the forecast of future market competition when a new car is introduced or the attributes of an existing car change. The model can be used by designers to anticipate the outcomes of a design change or a design release. For example, when a new car is released, the model can predict what other cars will be considered concurrently (co-consideration link existence). Therefore, designers or manufacturers can use this information to develop their design strategy. In addition, it is noticeable that the true positive rate for the prediction is over 80% for all the results shown, which shows there is a considerably high probability that an actual link exists that will test positive. This indicates that competition in the

Table 6: Car attributes type and feature importance

Attribute	Variable Type	Importance	Sample Values
Make	Categorical, Nominal	$2.44 \cdot 10^{-2}$	Audi, Ford
Body Type and Number of Doors	Categorical, Nominal	$1.26 \cdot 10^{-2}$	2 Door Coupe
Segment (Detailed)	Categorical, Nominal	$1.15 \cdot 10^{-2}$	CD Premium Car
Segment Number	Categorical, Nominal	$8.9 \cdot 10^{-3}$	1, 2, 3
Segment (Combined)	Categorical, Nominal	$8.1 \cdot 10^{-3}$	B, C
Market Category	Categorical, Nominal	$5.9 \cdot 10^{-3}$	Small Size
Body Type	Categorical, Nominal	$5 \cdot 10^{-3}$	Coupe
Community	Categorical, Nominal	$4.5 \cdot 10^{-3}$	1, 2, 3
Brand Origin	Categorical, Nominal	$4.1 \cdot 10^{-3}$	European, Japanese
Import	Categorical, Binary	$3.8 \cdot 10^{-3}$	[0, 1]
Lane Assistance	Categorical, Binary	$1.5 \cdot 10^{-3}$	[0, 1]
Third row of seats	Categorical, Binary	$1.4 \cdot 10^{-3}$	[0, 1]
Park Assistance	Categorical, Binary	$6 \cdot 10^{-4}$	[0, 1]
AWD	Categorical, Binary	$3 \cdot 10^{-4}$	[0, 1]
Leather Seats	Categorical, Binary	$1 \cdot 10^{-4}$	[0, 1]
EngineSize log	Numerical, Continuous	0	10.4409
Alloy Wheels	Categorical, Binary	$-2 \cdot 10^{-4}$	[0, 1]
Fuel Consumption	Numerical, Continuous	$-2 \cdot 10^{-4}$	8.216
Fuel per Power	Numerical, Continuous	$-2 \cdot 10^{-4}$	0.066
Luxury	Categorical, Binary	$-3 \cdot 10^{-4}$	[0, 1]
Autotrans	Categorical, Binary	$-3 \cdot 10^{-4}$	[0, 1]
Year of Data	Numerical, Discrete	$-4 \cdot 10^{-4}$	2013,2014
Price log	Numerical, Continuous	$-4 \cdot 10^{-4}$	16.0406
Stability Control	Categorical, Binary	$-5 \cdot 10^{-4}$	[0, 1]
Fuel Type	Categorical, Nominal	$-5 \cdot 10^{-4}$	ICE
Power log	Numerical, Continuous	$-7 \cdot 10^{-4}$	6.7535
Side Airbags	Categorical, Binary	$-7 \cdot 10^{-4}$	[0, 1]
Navigation	Categorical, Binary	$-8 \cdot 10^{-4}$	[0, 1]
Turbo	Categorical, Binary	$-1.4 \cdot 10^{-3}$	[0, 1]

future can be well captured by the prediction model.

Secondly, the feature importance results shed light on understanding the key features in the co-consideration network formation. The results of the feature importance in Table 6 show that some features, such as make, body type, import, lane assistance, third row, park assistance, and AWD, have a higher impact on the product co-consideration network, whereas other features, such as turbo and navigation, are not key factors in making predictions. Knowing these factors and introducing interventions to change them for future product iterations can enable a car manufacturer to affect the competition relationships, leading to a larger market share. However, we should warn that it is imprudent to make definitive conclusions from regression models without real-world validation. Nevertheless, our analysis sheds light on key factors that customers may be considering while making their purchase decisions.

6 Future Work and Limitations

This work demonstrates the efficacy of GNNs in modeling products and the relationships between them. The findings of this study have several important implications for future practice, which will be discussed next.

Predict link strength in weighted and directed networks using GNN The current link existence prediction model consists of a GNN model and a classification model. Similarly, we can add a regression model as the downstream task instead of a classification model. The metrics of measuring the link strength prediction model could use root mean square error (RMSE) and mean absolute error (MAE). The prediction of link strength will enable designers to more precisely evaluate the effects of potential designs on market demand compared to merely predicting the existence of links.

Predict network structure for multi-dimensional networks with heterogeneous links using GNN To further capture the relationship between customers and products, a

multi-dimensional customer-product network can model the heterogeneous edges to model customers' considerations and choices simultaneously. In this work, we focused on undirected edges. However, future work will analyze directed edges to study the final choice of a customer from within a set of options.

Limitations First of all, this study is limited by the nature of survey data. The link existence between a pair of products is measured by the customers' consideration behavior. However, with the restriction of survey data, this study only samples a small portion of the real car market. The network we studied has a density of 14.73%, which leads to an unbalanced dataset with most links classified as 0. To overcome the issue, we randomly selected a subset of samples from the original dataset to match the samples coming from both classes in the training process.

Another limitation lies in the interpretability of feature importance. When two features are correlated and one of the features is permuted, the model will still have access to the feature through its correlated feature. This will result in a lower importance value for both features, whereas they might actually be important. The problem is common in many interpretable machine learning problems, and our work is no exception to it.

Thirdly, a notable drawback of GraphSAGE is that sampled nodes might appear multiple times, thus potentially introducing a lot of redundant computation. With the increase of the batch size and the number of samples, the number of redundant computations increases as well. Moreover, despite having nodes in memory for each batch, the loss is computed on a single batch of them, and, therefore, the computation for the other nodes is also in some sense wasted. Further, the neighborhood sampling used in GraphSAGE is effective in improving computing and memory efficiency when inferring a batch of target nodes with diverse degrees in parallel. Despite this advantage, the default uniform sampling can suffer from high variance in training and inference, leading to sub-optimum accuracy. While new architectures inspired by GraphSAGE attempt to reduce computa-

tion time and performance variation, we did not focus on finding the best architecture for improving the computational efficiency, as it was not central to our focus area.

7 Conclusions

We present a systematic method to study and predict relationship between products by using the inductive graph neural network models.

This paper makes the following key contributions:

1. We show that neural network models, which can embed each node of a graph into a real vector, can capture node feature and graph structure information, to enable machine learning applications on complex networks. This is the first attempt of implementing GNNs to predicting product relationships.
2. We show that GNN models have better link prediction performance than ERGMs, both for held-out links from the same year and predicting the entire network structure for future years.
3. We overcome a limitation of GNN by proposing a new method to predict links between unseen cars for future years.
4. We show the scalability of the GNN method by modeling the effect of a large number of continuous and categorical attributes on link prediction.
5. We use permutation-based methods to find the importance of attributes to help design decisions.

In future work, we aim to make predictions on the product relationship strength and extend the current work on more complex network structures to investigate the relationship among customers and products.

References

- [1] Braha, D., Suh, N., Eppinger, S., Caramanis, M., and Frey, D., 2006. “Complex engineered systems”. In *Unifying Themes in Complex Systems*. Springer, pp. 227–274.
- [2] Holling, C. S., 2001. “Understanding the complexity of economic, ecological, and social systems”. *Ecosystems*, **4**(5), pp. 390–405.
- [3] Hoyle, C., Chen, W., Wang, N., and Koppelman, F. S., 2010. “Integrated bayesian hierarchical choice modeling to capture heterogeneous consumer preferences in engineering design”. *Journal of Mechanical Design*, **132**(12).
- [4] Newman, M. E., 2003. “The structure and function of complex networks”. *SIAM review*, **45**(2), pp. 167–256.
- [5] Simon, H. A., 1977. “The organization of complex systems”. In *Models of discovery*. Springer, pp. 245–261.
- [6] Wasserman, S., Faust, K., et al., 1994. “Social network analysis: Methods and applications”.
- [7] Wang, M., Chen, W., Fu, Y., and Yang, Y., 2015. “Analyzing and Predicting Heterogeneous Customer Preferences in China’s Auto Market Using Choice Modeling and Network Analysis”. *SAE International Journal of Materials and Manufacturing*, **8**(2015-01-0468), pp. 668–677.
- [8] Fu, J. S., Sha, Z., Huang, Y., Wang, M., Fu, Y., , and Chen, W., 2017. “Modeling Customer Choice Preferences in Engineering Design using Bipartite Network Analysis”. In Proceedings of the ASME 2017 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference.
- [9] Sha, Z., Huang, Y., Fu, S., Wang, M., Fu, Y., Contractor, N., and Chen, W., 2018. “A Network-Based Approach to Modeling and Predicting Product Co-Consideration Relations”. *Complexity*, **2018**.
- [10] Sha, Z., Saeger, V., Wang, M., Fu, Y., and Chen, W., 2017. “Analyzing Customer Preference to Product Optional Features in Supporting Product Configuration”. *SAE International Journal of Materials and Manufacturing*, **10**(2017-01-0243).

- [11] Wang, M., Chen, W., Huang, Y., Contractor, N. S., and Fu, Y., 2016. “Modeling customer preferences using multidimensional network analysis in engineering design”. *Design Science*, **2**.
- [12] Wang, M., Sha, Z., Huang, Y., Contractor, N., Fu, Y., and Chen, W., 2016. “Forecasting Technological Impacts on Customers’ Co-Consideration Behaviors: A Data-Driven Network Analysis Approach”. In ASME 2016 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, pp. V02AT03A040–V02AT03A040.
- [13] Cui, Y., Ahmed, F., Sha, Z., Wang, L., Fu, Y., and Chen, W., 2020. “A weighted network modeling approach for analyzing product competition”. In International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, Vol. 84003, American Society of Mechanical Engineers, p. V11AT11A036.
- [14] Sha, Z., Huang, Y., Fu, J. S., Wang, M., Fu, Y., Contractor, N., and Chen, W., 2018. “A network-based approach to modeling and predicting product coconsideration relations”. *Complexity*, **2018**.
- [15] Stokes, J. M., Yang, K., Swanson, K., Jin, W., Cubillos-Ruiz, A., Donghia, N. M., MacNair, C. R., French, S., Carfrae, L. A., Bloom-Ackermann, Z., et al., 2020. “A deep learning approach to antibiotic discovery”. *Cell*, **180**(4), pp. 688–702.
- [16] Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., and Philip, S. Y., 2020. “A comprehensive survey on graph neural networks”. *IEEE transactions on neural networks and learning systems*.
- [17] Jain, A., Liu, I., Sarda, A., and Molino, P., 2019. Food Discovery with Uber Eats: Using Graph Learning to Power Recommendations. <https://eng.uber.com/uber-eats-graph-learning/>. [Online; accessed 01-March-2021].
- [18] Wang, J., Huang, P., Zhao, H., Zhang, Z., Zhao, B., and Lee, D. L., 2018. “Billion-scale commodity embedding for e-commerce recommendation in alibaba”. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 839–848.
- [19] Wang, M., Sha, Z., Huang, Y., Contractor, N., Fu, Y., and Chen, W., 2018. “Predicting product co-consideration and market competitions for technology-driven product design: a network-based approach”. *Design Science*, **4**.
- [20] Shocker, A. D., Ben-Akiva, M., Boccara, B., and Nedungadi, P., 1991. “Consideration set influences on consumer decision-making and choice: Issues, models, and suggestions”. *Marketing letters*, **2**(3), pp. 181–197.
- [21] Hauser, J. R., and Wernerfelt, B., 1990. “An evaluation cost model of consideration sets”. *Journal of consumer research*, **16**(4), pp. 393–408.
- [22] Roberts, J. H., and Lattin, J. M., 1991. “Development and testing of a model of consideration set composition”. *Journal of Marketing Research*, **28**(4), pp. 429–440.
- [23] Gaskin, S., Evgeniou, T., Bailiff, D., and Hauser, J., 2007. “Two-stage models: Identifying non-compensatory heuristics for the consideration set then adaptive polyhedral methods within the consideration set”. In Proceedings of the Sawtooth Software Conference, Vol. 13, Citeseer, pp. 67–83.
- [24] Dieckmann, A., Dippold, K., and Dietrich, H., 2009. “Compensatory versus noncompensatory models for predicting consumer preferences”.
- [25] Damangir, S., Du, R. Y., and Hu, Y., 2018. “Uncovering patterns of product co-consideration: A case study of online vehicle price quote request data”. *Journal of Interactive Marketing*, **42**, pp. 1–17.

- [26] Wang, M., Huang, Y., Contractor, N., Fu, Y., Chen, W., et al., 2016. “A network approach for understanding and analyzing product co-consideration relations in engineering design”. In DS 84: Proceedings of the DESIGN 2016 14th International Design Conference, pp. 1965–1976.
- [27] Zhou, J., Cui, G., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., and Sun, M., 2018. “Graph neural networks: A review of methods and applications”. *arXiv preprint arXiv:1812.08434*.
- [28] Kipf, T. N., and Welling, M., 2016. “Semi-supervised classification with graph convolutional networks”. *arXiv preprint arXiv:1609.02907*.
- [29] Atwood, J., and Towsley, D., 2015. “Diffusion-convolutional neural networks”. *arXiv preprint arXiv:1511.02136*.
- [30] Cao, S., Lu, W., and Xu, Q., 2016. “Deep neural networks for learning graph representations”. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 30.
- [31] Qiu, J., Dong, Y., Ma, H., Li, J., Wang, K., and Tang, J., 2018. “Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec”. In Proceedings of the eleventh ACM international conference on web search and data mining, pp. 459–467.
- [32] Hamilton, W. L., Ying, R., and Leskovec, J., 2017. “Inductive representation learning on large graphs”. *arXiv preprint arXiv:1706.02216*.
- [33] Doshi-Velez, F., and Kim, B., 2017. “Towards a rigorous science of interpretable machine learning”. *arXiv preprint arXiv:1702.08608*.
- [34] Molnar, C., 2020. *Interpretable machine learning*. Lulu. com.
- [35] Du, M., Liu, N., and Hu, X., 2019. “Techniques for interpretable machine learning”. *Communications of the ACM*, **63**(1), pp. 68–77.
- [36] Ribeiro, M. T., Singh, S., and Guestrin, C., 2016. ““ why should i trust you?” explaining the predictions of any classifier”. In Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining, pp. 1135–1144.
- [37] Breiman, L., 2001. “Random forests”. *Machine learning*, **45**(1), pp. 5–32.
- [38] Fisher, A., Rudin, C., and Dominici, F., 2019. “All models are wrong, but many are useful: Learning a variable’s importance by studying an entire class of prediction models simultaneously.”. *Journal of Machine Learning Research*, **20**(177), pp. 1–81.
- [39] Altmann, A., Tološi, L., Sander, O., and Lengauer, T., 2010. “Permutation importance: a corrected feature importance measure”. *Bioinformatics*, **26**(10), pp. 1340–1347.
- [40] Putin, E., Mamoshina, P., Aliper, A., Korzinkin, M., Moskalev, A., Kolosov, A., Ostrovskiy, A., Cantor, C., Vijg, J., and Zhavoronkov, A., 2016. “Deep biomarkers of human aging: application of deep neural networks to biomarker development”. *Aging (Albany NY)*, **8**(5), p. 1021.
- [41] Matin, S., Farahzadi, L., Makaremi, S., Chelgani, S. C., and Sattari, G., 2018. “Variable selection and prediction of uniaxial compressive strength and modulus of elasticity by random forest”. *Applied Soft Computing*, **70**, pp. 980–987.
- [42] Farinosi, F., Giupponi, C., Reynaud, A., Ceccherini, G., Carmona-Moreno, C., De Roo, A., Gonzalez-Sanchez, D., and Bidoglio, G., 2018. “An innovative approach to the assessment of hydro-political risk: A spatially explicit, data driven indicator of hydro-political issues”. *Global Environmental Change*, **52**, pp. 286–313.