

BILATERAL TRADE MODELLING WITH GRAPH NEURAL NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

Bilateral trade agreements confer preferred trading status between participating countries, enabling increased trade and potential economic growth. Predicting such trade flows often serve as important economic indicators used by economists and policymakers with impactful ramifications in economic policies adopted by respective countries. However, traditional approaches to predicting potential trade partners are through the use of gravity methods which are cumbersome to define due to the exponentially growing number of constants that need to be considered. In this work, we present a framework for directly predicting bilateral trade partners from observed trade records using *graph representation learning*. Furthermore, we show as a downstream task that modeling bilateral trade as a graph allows for the classification of countries into various income levels. Empirically, we observe accuracies of up to 98% for predicting trading partners and 68% on income level classification.

1 INTRODUCTION

International trade involves the exchange of goods, capital and services between countries and where two countries are involved, it is referred to as *bilateral trade*. Often, the deficit and surpluses created via bilateral trade represent important economic development indicators, which drive the adoption of specific domestic economic policies –i.e. relaxation of restrictions and trade barriers, in either country. Consequently, various models have been employed by economists to understand trade patterns and factors that account for the observed trade activities between countries. For instance, the Ricardian model, introduced the idea of comparative advantage of countries, whereby a country exports more of the goods they can produce at a lower cost (Ricardo, 1817). In a similar vein, the “factor of abundance” argues that the trading behaviour of a country is influenced by what they confidently produce in abundance (Steiner, 2015).

The most popular method with empirical benefits is known as the Gravity Model of trade, which is motivated by Newton’s law of gravitation. In essence, the gravity model relates the bilateral trade flows between two countries using the respective gross domestic product (GDP) of each country while taking into account the geographical distance. Intuitively, trade flow is high when participating countries have high GDP’s and are geographically close to each other (Deardorff, 1998). While the gravity model is an effective empirical measure for bilateral trade flow, it lacks both a theoretical justification (Anderson, 2011), as well as suffers from practical limitations. In particular, model performance is dictated by defining handcrafted features such as cultural differences and political terms that require significant domain knowledge.

Present work. In this paper, we take a data driven approach to modelling bilateral trade. We first make the observation that trade flows can naturally be interpreted as a graph wherein countries are nodes and edges represent countries undertaking bilateral trade. We leverage recent advances in graph representation learning and predict trade links between countries, crucially without first estimating trade flow heuristics. We further analyze the graphical structure of trade relationships between countries and use it to power a supervised learning approach to predict income levels of countries using graph neural networks (GNN’s). Empirically we observe 98% and 68% accuracies in the predicting bilateral trade links and income levels respectively.

2 BACKGROUND

2.1 GRAVITY MODEL OF TRADE

Inspired by Newton’s law of universal gravitation, the gravity model provides a theoretical approach to representing the numerical trade strength between any two countries. The gravity model is used to compute the magnitude of trade strength, —i.e. trade flow, between any two countries. The magnitude of this trade value increases with increasing respective net income or GDP and decreases with increasing distance (Deardorff, 1998; Chaney, 2013). It is expressed as

$$F_{ij} = M \frac{GDP_i \cdot GDP_j}{D_{ij}}, \quad (1)$$

where F_{ij} is the trade flow between countries i and j , GDP_i is the GDP of country i , M is a proportionality constant and D_{ij} is the geographical distance between countries i and j . A more convenient way to deal with this equation is to express it in log and introduce coefficients and placeholder variables to account for other unanticipated factors such as cultural influences and political terms. Equation 1 can then be expressed as follows:

$$\ln F_{ij} = c_0 + c_1 \ln GDP_i + c_2 \ln GDP_j + c_3 \ln D_{ij} + c_4 d + c_5 P_{ij} + \epsilon_{ij},$$

where c_k are hand engineered constants, P_{ij} is a political influence term and ϵ_{ij} is an error correction term. (Helpman et al., 2008)

2.2 COUNTRY CLASSIFICATION

The World Bank defines four income groups for all countries in the world namely: high-, upper middle, lower middle and low-income. The division into these income groups is based on the total annual income called the gross national income (GNI) per capita. The GNI of a country gives an idea of its economic strength and weaknesses and in general, the standard of living of the average citizen. Countries are classified into various income groups if their GNI falls within a certain threshold as shown in table 1 (Team, 2017). This classification by income levels can be used to measure progress over time or analyse data for countries falling into the same income groups.

Table 1: Country GNI and income group - World Bank, 2017.

Table 2: Summary of data features and representation.

Income Group	GNI threshold	Feature	Notation	Size	Representation
lower	1,006 and below	nodes	\mathcal{V}	111	countries
lower middle	1,006 - 3,955	node features	\mathbf{X}	38	population etc. ¹
upper middle	3,956 - 12,235	edges	\mathbf{A}	476	trade indicator
high	12,235 and above	edge weights	\mathcal{E}	476	net trade val (USD)
		node labels	\mathcal{Y}	4	income group

2.3 GRAPH NEURAL NETWORKS

Given a simple graph $\mathcal{G} = (\mathcal{V}, \mathbf{X}, \mathbf{A}, \mathcal{E})$ that contains a single type of relation and no self-loops with individual entities, —i.e. countries, referred to as nodes \mathcal{V} with real valued features $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times D}$. Here, $|\mathcal{V}|$ and D are the number of nodes and the dimensionality of node features respectively. The dense square matrix $\mathbf{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ is known as the adjacency matrix and $\mathbf{a}_{ij} = 1$ if there exists an edge between nodes i and j and zero otherwise.

Graph Neural Networks (GNNs) are a family of approaches that aim to generalize neural networks, developed for Euclidean data, to graphs (Scarselli et al., 2009). Common tasks on graph structured data involve link prediction, where the goal is to predict potential edges between nodes, and node classification where the classifier has to predict a label $y \in \mathcal{Y}$ (Kipf & Welling, 2016a; Bhagat et al., 2011; Zhang & Chen, 2018). While there are many other interesting problems with graph structured data, we focus on the problem of link prediction and node classification and direct the interested reader to the following surveys for more information (Hamilton et al., 2017).

3 DATA AND TASKS

We now outline our proposed approach to modelling bilateral trade using graph representation learning. We first describe the basic components of a typical graph and relate each to a component in our data. Our primary goals are to (1) predict potential trade partners and (2) classify countries into their respective income groups. These are essentially multi-class node classification and link prediction tasks. We trained Graph Neural Network models and baseline models to perform the aforementioned tasks. In the next subsection, we talk about how we collected the data, the representation approach we took and the GNN models adopted for the downstream tasks.

3.1 DATA COLLECTION AND REPRESENTATION

The data used for this study are taken from two different sources: the United Nations Comtrade Database and Kaggle. The United Nations Comtrade Database² is an international trade database containing the reporter-partner trade statistics collated for about 170 countries over a certain period of time. These trade statistics include (1) imports, exports, re-exports and re-imports, (2) commodities exchanges between the trade partner and reporter, (3) trade value in US dollars.

The data retrieved from Kaggle, on the other hand, had the profile of specific countries which included the geographical, financial, geological and other information. To ensure consistency, we targeted data for a particular year and in all found the trade and profile information for a total of 111 countries together with their income groups used as target labels.

Each country considered as a node had 38 node features being the profile collected for each country. We then used the net trade balance between the countries to construct an adjacency matrix such that there is an edge between the countries if the trade balance was not zero and vice versa. An entry of 1 is assigned if there is an edge and 0 otherwise. The trade values are in US dollars and constitute the edge weight matrix. A summary of features and the representation is provided in Table 2.

4 EXPERIMENTS

We evaluate our approach to modelling bilateral trade using GNN’s in two settings: link prediction of withheld edges between trading countries and income classification of countries (node classification) as a separate downstream task. We train all GNN based models using PyTorch Geometric (Fey & Lenssen, 2019) using 80% and 20% for training and test sets. For optimization of model parameters we use the Adam optimizer (Kingma & Ba, 2014), while hyperparameters are tuned using Bayesian optimization (Snoek et al., 2012; Brochu et al., 2010).

Setup and Baselines. In addition to GNNs, we test a multi-layer perceptron (MLP) and a logistic regression model as baseline models for node classification. Both MLP and the logistic regression model consume node features as input, but critically do not have access to any underlying graph structure in the form of an adjacency matrix. We hypothesize that effective learning requires the utilization of local neighborhood information available via the adjacency matrix —i.e. adoption of specific domestic trade policies in one country can influence similar policies to its neighbors.

4.1 LINK PREDICTION

In the link prediction task, we used the graph autoencoder (GAE) and variational graph autoencoder (VGAE) (Kipf & Welling, 2016b) to learn a latent representation of the input graph. The input graph having few observed edges and hence a sparse adjacency matrix. This input graph is fed into the link prediction model which then reconstructs a new adjacency matrix representing a new neighbourhood structure for each node in the graph. We evaluate the reconstruction accuracy using the area under the curve (AUC) and average precision (AP) metrics. In Table 4, we report high AUC and AP scores for both GAE and VGAE. This is indicative of how well these models are reliably able to reconstruct the adjacency matrix from the learned latent representation.

²<https://comtrade.un.org/data/>

4.2 NODE CLASSIFICATION

We used the graph convolution network (GCN) (Kipf & Welling, 2016a), ChebNet (Defferrard et al., 2016), graph attention network (GAT) (Veličković et al., 2018) and attention-based graph neural network (AGNN) (Thekumparampil et al., 2018) models described in Section 2 to perform a multi-class node classification with the input graph. We split the data into train and test sets. We do this by randomly masking out 20% of the total nodes and edge indices from the adjacency matrix and using the 80% for training. The learned model is then used to predict labels for the masked out set of nodes. We then report the classification accuracy on the test accuracy as a measure of model performance. We summarize results denoting mean results after 100 runs in Table 3. We report that the GCN has the best accuracy score after 1200 epochs where we do an early stopping to avoid deterioration of model performance as was observed. ChebNet also compares competitively with the linear baseline model.

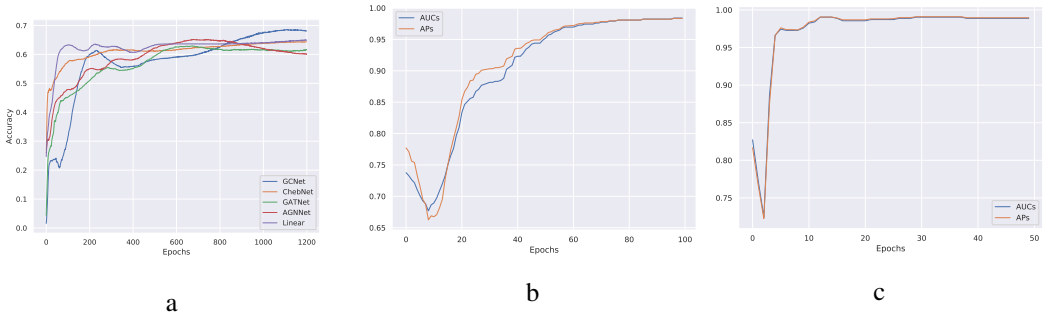


Figure 1: (a) Model performance on node classification over 100 runs (b) AUC and AP scores on reconstruction with GAE (c) AUC and AP scores on reconstruction with VGAE

Table 3: Results for multi-class node classification task

	GCN	ChebNet	GAT	AGNN	Linear	Logistic Reg
Test Accuracy	0.6812	0.6436	0.6158	0.6003	0.6491	0.5758

Table 4: Results for link prediction task

	GAE	VGAE
AUC	0.9840	0.9888
Average Precision	0.9835	0.9896

5 DISCUSSION AND CONCLUSION

We approach downstream tasks — income level classification and potential trade partners prediction among countries — as a graph representation learning by leveraging their historical bilateral trade relationships. We use graph neural networks to classify countries and predict a new neighbourhood structure for each country. Our approach naturally points to a new direction machine learning particularly graph representation learning and application in the field of Economics. From the results we obtained, we confirm that our approach does well for the intended tasks and could be improved for future trade analysis. Since we only considered a standard static graph from data of trade activities between countries in a particular year, interesting future work is to use a time-series data of trade activities over a finite length of time in dynamic graph analysis. This will encourage analysis of (1) how countries evolve from one class to the other with time and (2) how trade activities change with time between countries. In the latter case, a temporal prediction of an edge between any two countries (nodes) will be an indication of how likely they will partner in trade in future.

REFERENCES

- James E Anderson. The gravity model: Annual review of economics. 2011.
- Smriti Bhagat, Graham Cormode, and S. Muthukrishnan. Node classification in social networks. *CoRR*, abs/1101.3291, 2011. URL <http://arxiv.org/abs/1101.3291>.
- Eric Brochu, Vlad M. Cora, and Nando de Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *CoRR*, abs/1012.2599, 2010. URL <http://arxiv.org/abs/1012.2599>.
- Thomas Chaney. The gravity equation in international trade: An explanation. Working Paper 19285, National Bureau of Economic Research, August 2013. URL <http://www.nber.org/papers/w19285>.
- Alan Deardorff. Determinants of bilateral trade: Does gravity work in a neoclassical world? *The Regionalization of the World Economy*, pp. 7–32, 1998. doi: <https://doi.org/10.3386/w5377>.
- Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. *CoRR*, abs/1606.09375, 2016. URL <http://arxiv.org/abs/1606.09375>.
- Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- William L Hamilton, Rex Ying, and Jure Leskovec. Representation learning on graphs: Methods and applications. *arXiv preprint arXiv:1709.05584*, 2017.
- Elhanan Helpman, Marc Melitz, and Yona Rubinstein. Estimating Trade Flows: Trading Partners and Trading Volumes*. *The Quarterly Journal of Economics*, 123(2):441–487, 05 2008. ISSN 0033-5533. doi: 10.1162/qjec.2008.123.2.441. URL <https://doi.org/10.1162/qjec.2008.123.2.441>.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016a.
- Thomas N Kipf and Max Welling. Variational graph auto-encoders. *NIPS Workshop on Bayesian Deep Learning*, 2016b.
- David Ricardo. *On the Principle of Political Economy and Taxation*, chapter 1. Batoche Books, 52 Eby Street South, Kitchener, Ontario, N2G 3L1 Canada, 3 edition, 1817.
- F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, Jan 2009. ISSN 1941-0093. doi: 10.1109/TNN.2008.2005605.
- Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical bayesian optimization of machine learning algorithms. In *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 2, NIPS’12*, pp. 2951–2959, Red Hook, NY, USA, 2012. Curran Associates Inc.
- Patrick Steiner. Determinants of bilateral trade flows, 2015.
- World Bank Data Team. New country classifications by income level, 2017. URL <https://blogs.worldbank.org/opendata/new-country-classifications-income-level-2017-2018>. Last accessed 31 December 2019.
- Kiran K. Thekumparampil, Sewoong Oh, Chong Wang, and Li-Jia Li. Attention-based graph neural network for semi-supervised learning, 2018. URL <https://openreview.net/forum?id=rJg4YGWRb>.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph Attention Networks. *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rJXMpikCZ>. accepted as poster.

Muhan Zhang and Yixin Chen. Link prediction based on graph neural networks. *CoRR*, abs/1802.09691, 2018. URL <http://arxiv.org/abs/1802.09691>.