

Referral prediction in Healthcare using Graph Neural Networks

REGINA DUARTE, Instituto Superior Tecnico, Universidade de Lisboa, Portugal

QIWEI HAN, NOVA School of Business and Economics, Nova University of Lisbon, Portugal

CLAUDIA SOARES, NOVA School of Science and Technology, Nova University of Lisbon, Portugal

Medical referrals between primary care doctors and specialists affect many aspects of patient care, such as quality of care, patient satisfaction, and health care costs. In this work, we hypothesize that the social network of medical doctors can influence the referral process. We analyse primary-specialty referrals through transaction data over medical appointments gathered between 2012 and 2017 in all hospital centers of an European private healthcare provider. In this paper we uncover patterns and hidden mechanisms in primary-specialty referrals to improve the efficiency of health care services, and propose a link prediction model for recommendation using Graph Neural Networks. First, we carried out exploratory data analysis and searched for patterns or interesting features of the data and afterwards, we created and analysed the doctor's social network and the referrals network. We learned the representation of the referral network using a Graph Neural Network (GNN). Finally, we discussed the nodes representation of the referral network which evidences our main hypothesis. This work addresses the discovery of important patterns in medical referrals to, in future work, improve the efficiency of collaboration within organizations.

CCS Concepts: • **Computing methodologies** → **Neural networks**; • **Applied computing** → **Health care information systems**; • **Information systems** → **Recommender systems**.

Additional Key Words and Phrases: medical referral, graph neural networks, social networks for recommendation

ACM Reference Format:

Regina Duarte, Qiwei Han, and Claudia Soares. 2021. Referral prediction in Healthcare using Graph Neural Networks. 1, 1 (September 2021), 14 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Medical referral between primary care physicians (PC) and specialists (SC) represents the formal mechanism in the health system to address the need of patients for specialty care. It is defined as the transfer or sharing of responsibility over the patient healthcare, between medical doctors, generally from the primary care doctor to a specialist doctor [1]. It may affect many aspects of patient care, such as quality of care, patient satisfaction, health care costs, etc. [2]. Many qualitative studies on the subject focused on cooperation and conflict between physicians [3] and on how personal relationships could interfere on the referral process [4, 5], while others investigate new collaboration models between general practitioners and specialists [6]. A recent review paper focuses on the competences to promote effective collaboration in patient care between primary care and specialist medical staff [7], evidencing how the handover between primary and secondary care can be an important factor for the health outcome of the patient [8]. The existing

Authors' addresses: Regina Duarte, reginaduarte@tecnico.ulisboa.pt, Instituto Superior Tecnico, Universidade de Lisboa, Lisbon, Portugal; Qiwei Han, qiwei.han@novasbe.pt, NOVA School of Business and Economics, Nova University of Lisbon, Carcavelos, Portugal; Claudia Soares, claudia.soares@fct.unl.pt, NOVA School of Science and Technology, Nova University of Lisbon, Caparica, Portugal.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

literature typically leveraged the patient consultation history extracted from insurance claims data to construct the patient sharing network between physicians based on the shared patients [2, 9]. The patient sharing network essentially operationalizes an informal information-sharing network in which physicians provide care to shared patients. However, this network does not necessarily conform to the formal organizational structure that physicians are affiliated with, and thus may provide valuable insights in explaining the referral mechanism. A previous paper applied community detection algorithms to the referral network in Canada [10], although this is a very interesting study, it is based on specially created survey data, which is known to be prone to biases and is necessarily small in number of answers.

In this paper, we hypothesize that medical doctors' informal social networks can influence the referral process. In other words, doctors' referral decisions may be limited to their social contacts and do not always benefit patient-centered care. This implies that network structure metrics derived from the doctors' social network can serve as informative features to boost the predictive performance of a model for referral recommendations [11, 12]. As such, after we present and analyse the data we use, we create two networks: 1) the referral network connecting PC to SC if a patient consults a PC and then an SC within a month, and 2) the social network of all doctors according to their similar profiles. Then, we learn the representation of the referral network using a Graph Neural Network (GNN) [13] and discuss the findings. The main objective of the study is to uncover hidden mechanisms in the primary-specialty referrals using features extracted from informal social network of doctors, which may help health organizations to improve the referral process through recommendations [14].

2 DATA COLLECTION AND DATA WRANGLING

To conduct this analysis and hypothesis validation, we began to analyse a large-scale patient consultation data from a Portuguese private healthcare provider to understand more about the referral patterns. The data gathered includes a data set about the consultations that were done between 2012 and 2017 with over 12 million consultations and around 1.4 million patients. Per consultation we have access to (1) the demographics information of the patient like the gender, age, their nationality and address; (2) which doctor the patient have the appointment with; (3) the date of the occurrence and (4) the Hospital where the consultation took place. In parallel, we also have access to information about each doctor: their gender, age, education, the hospital where the doctor works, their first specialty (and the second one if it is the case). We encounter 3,632 physicians whose 389 were primary care doctors (PC) and 1,313 were Specialty care doctors (SC) The rest of the doctors had an unknown specialty (1,930 doctors) through 7 hospitals.

2.1 Exploratory Data Analysis

For a better understanding of the data, we started with some confirmations of tendencies that would be expected. One example of that is the patient age distribution. As expected and shown in figure 1, there are more elder and younger patients, because those are the ages that requires more health care. The number of consultations throughout the year, shown in figure 2 is another expected tendency that presents a higher number of appointments in the cold months. Concerning the appointments by specialization, we find out that the number of appointments for each specialty follows different tendencies. Figure 3 shows a great increase in the number of family doctors consultations in 2012 while the increase of other specialties is rather slow. It is also interesting to point out that the specialty radiology has a high number of consultations per year.

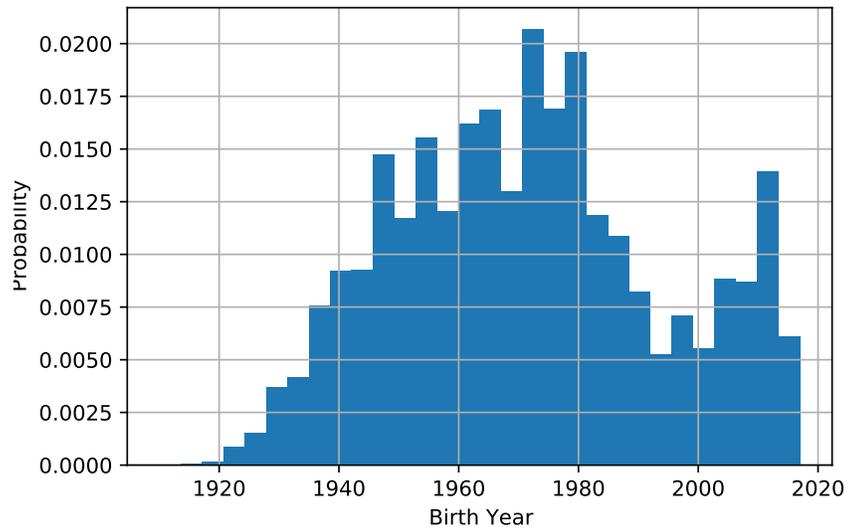


Fig. 1. Distribution of birth year of the patients. It is clear that elder patients and the children are more common patients.

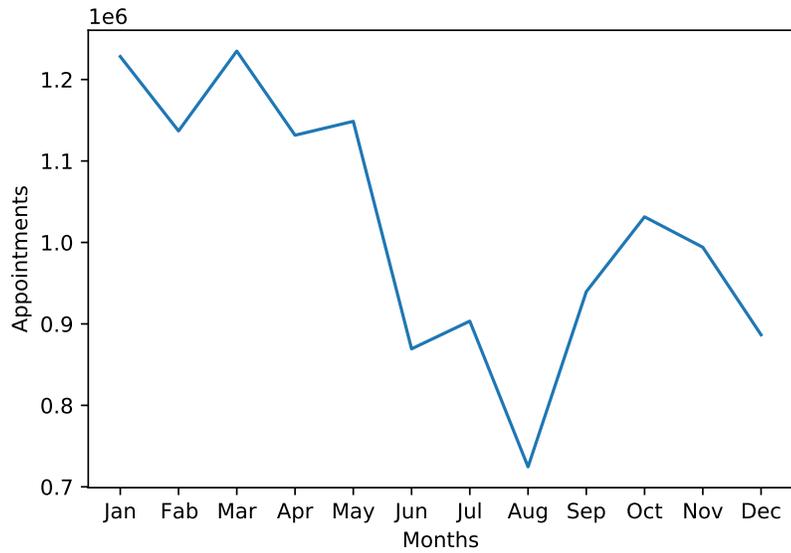


Fig. 2. Number of consultations throughout the year. There are a lot more appointments in the cold months, and in December the number is also relatively low.

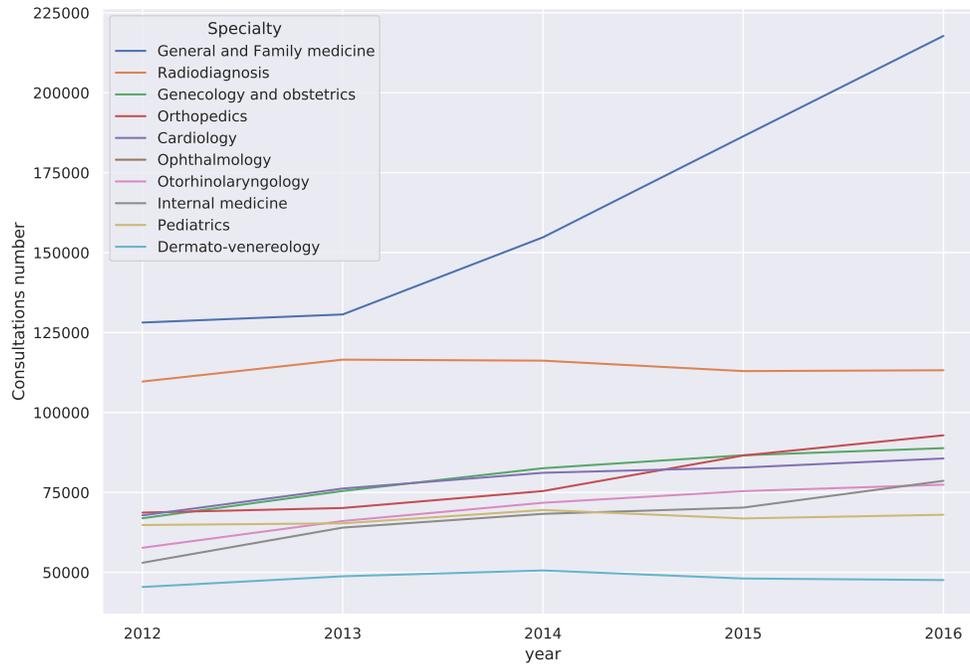


Fig. 3. Number of appointments per specialization throughout the years. Here is only shown the nine most requested specialties. It shows a great increase in the number of family doctors consultations in 2012 while the increase of other specialties is rather slow. It is also interesting to point out that the specialty radiology has a high number of consultations per year

An interesting analysis is the gender distribution of the doctors based on their age. Figure 4 shows the birth year distribution of the doctors for the top 3 most frequented faculties highlighted by gender. The analysis suggest that in the decades of 70/80 the percentage of women pursuing medicine increased. This tendency is correlated with the dictatorial state of Portugal that ended in 1974. After this event, women eventually start to be the predominant gender in the higher education institutions, not only in the health sciences but in a vast fields in the academia. This phenomenon is well documented in [15], chapter 4.

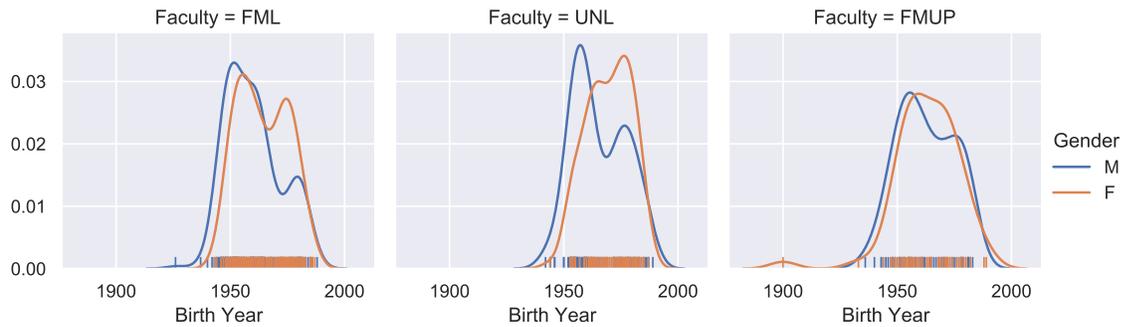


Fig. 4. Doctors Age distribution for the 3 most frequented universities. FML stands for Faculty of medicine of Lisbon University, FMUP stands for Faculty of medicine of Porto University and UNL stand for Nova University of Lisbon. In FML and UNL there is a clear growth of female medicine students in the decades of 70/80

A further analysis shows that the number of doctors that one patient has had appointments with follows a power law distribution (figure 5). Being so, a considerable amount of patients only had consultations with one doctor, so the data related to those patients turns out irrelevant to exploring the referrals mechanism.

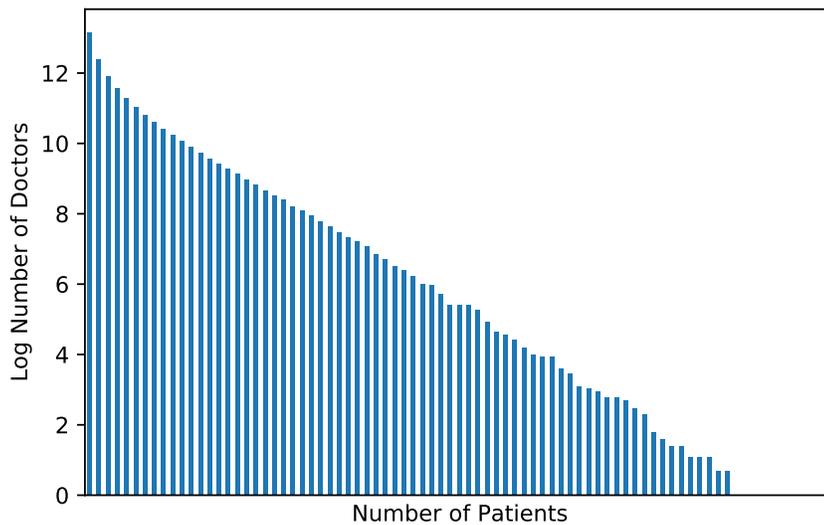


Fig. 5. Number of physicians (in log scale) that patients have appointments with. We see a power-law phenomenon, where only a few patients consulted many doctors and most patients consult one to three clinicians.

Joining the information of the two data sets leads to the first step towards the analysis of the referrals. In figure 6 is shown the cumulative percentage of interactions through the time interval. An interaction is a pair of appointments

of the same patient. The first appointment occurs first in time and is necessarily a primary care consultation and the second one is a specialty appointment of any kind. There were only considered interactions with the minimum time interval between the two appointments. (e.g. an appointment with the familiar doctor in January followed by an appointment in cardiology in February leads to an interaction with a time interval of one month). In figure 7 this idea is further explored.

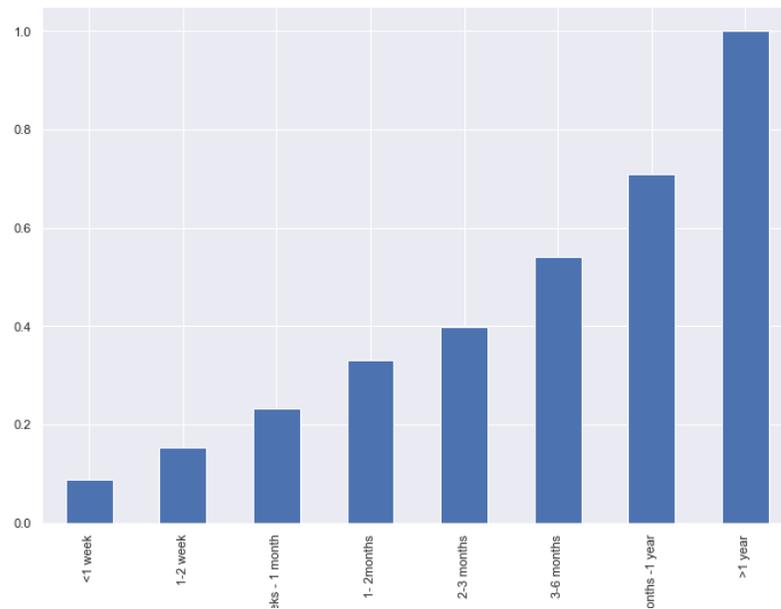


Fig. 6. Cumulative distribution of the time interval of the doctor-doctor interactions, through time. Interactions are defined as consecutive same patient appointments from a PC appointment to a SC one. For the construction of the referral network we consider interactions that occur in less than one month which represents approximately 22% of the interactions.

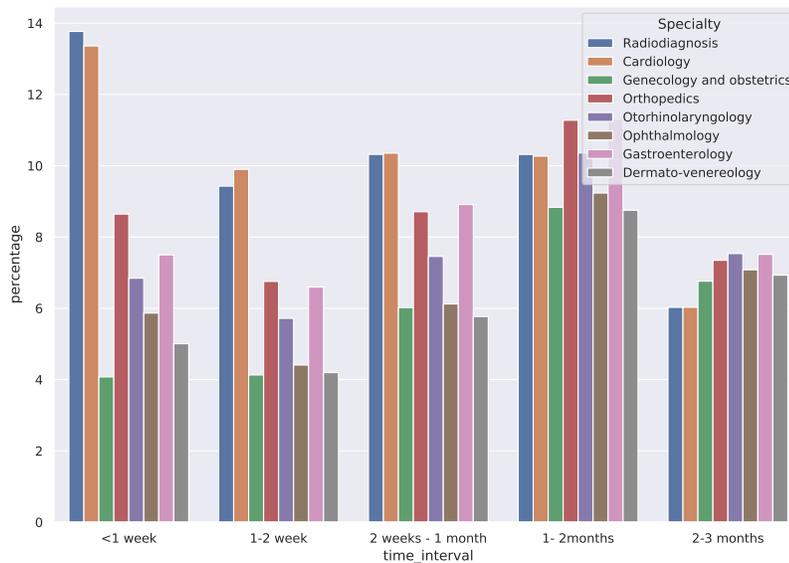


Fig. 7. Percentage of interactions of each specialty that were done per time interval. Among other things, it's clear that Cardiology and radio-diagnosis are the specialties that have more appointments in less than a week after a primary care appointment - around 13% of all interactions within these specialties occur in less than a week. In the other hand, in orthopedics, is more probable to have an appointment with a specialist one a two months after the primary care consultation. It is an evidence of the urgency of some specialties in comparison with others.

3 METHODOLOGY

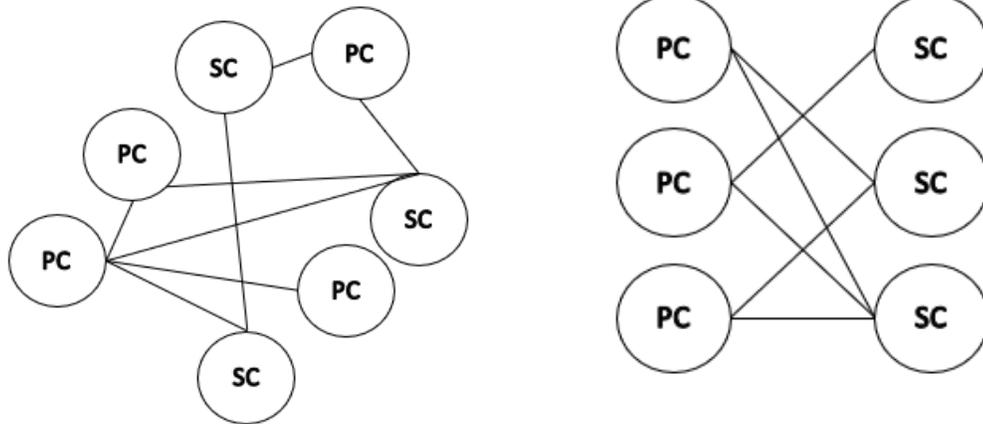
3.1 Medical Referrals as a Social Network

With the data wrangling and initial analysis done, we constructed (1) a *referral weighted bipartite graph*, to analyse deeper the referrals, and (2) a *social graph* of the doctors. In Fig.8 we present a schema of the two graphs.

Following common standard used in [16, 17], the referral graph was built using only the interactions that occurred within a month. The links are the referrals made from the Primary Care to the Specialist Care doctors, so there are only links connecting one PC doctor to another SC doctor. As such, in the referral graph 294 PC are connected with 839 SC through 34,249 edges. The edge weight on the referral network represents the number of patients that PC refers to the SC. Importantly, many physicians do not link to the referral network, which raises potential inefficiency concerns for their lack of involvement in the referral process.

The social network was built from the information about the doctors – an edge connects two doctors if they share similar profiles such as whether they receive their medical degree from the same institution at the same time, perform the residency internship in the same hospital or the location where they studied and where they are working in, to induce connections from work and education.

We calculated summary network metrics for the resulting referral and social networks, respectively. For example, the degree of nodes in the referral network follows the power-law distribution. This means that: (1) there are few physicians



(a) Social interactions: two doctors are connected if they share similar profiles, e.g., did medical school in the same institution at the same time or work together in the same hospital.

(b) Referrals: a primary care (PC) doctor is connected to a specialty doctor (SC) if there is a patient that had an appointment with the SC up to one month after an appointment with the PC.

Fig. 8. Representation of the two graphs constructed. The Social network in (a) models the social interactions between doctors. Two doctors are connected if they know each other from their work or education. Similarly in (b) the Referral network models the medical referrals done by PC to SC.

that receive a huge amount of referrals (if the node is an SC) or that refer to a lot of specialty doctors (if the node is a PC) (2) the vast majority of the SC only receive referrals from few PC or PC only refer to a few SC. Such observations are subject to several possible explanations. The SC with high in-degree can be those with high popularity, while PC with low out-degree may have relatively limited social contacts. Meanwhile, we obtained the average clustering coefficient for the referral network (0.149) to measure the fraction of the number of observed squares to the total number of possible squares in the network. This represents an essential precondition for the referral network to exhibit small-world structure and suggests that physicians in the referral network have a higher tendency to cluster together.

3.2 Referral Prediction using Graph Embeddings

Node embeddings learned from graph-structured data provide low-dimensional vector representations for each node using its graph neighborhood [18]. It has showed to be very useful for numerous machine learning applications, such as node classification, clustering, and link prediction. We adopt the GraphSAGE model to generate node embeddings for the referral network, because the model can leverage node attributes to jointly learn the structure of each node's neighborhood together with the distribution of node features in the neighborhood [13]. The GraphSAGE model is an unsupervised representation learning model that can learn node embeddings from large graphs using jointly the graph structure and the node features. This is done by solving a classification task: predict whether a node pair is likely to co-occur in a random walk on the graph. To solve the task are generated positive and negative node pairs. The positives pairs are generated from performing random walks and the negatives are generated by randomly selected nodes according to a distribution. With this simple classifier, the model can learn inductive mappings from attributes of nodes and their neighbours to node embeddings while preserving the structural similarities of nodes and their features.

These GraphSAGE node embeddings are the base for the link prediction tasks we perform on the referral network to test the hypothesis of social features of doctors influence the referral process. To this end we use the Stellargraph Library which is a Python library for machine learning on graphs and networks that already has implemented the GraphSAGE model for both tasks of link prediction and node embeddings [19].

From the social network of doctors, we computed three centrality measures for each node: betweenness, eigen vector, and degree centrality, and added them as features to GraphSAGE to accomplish the link prediction task, based on the referral network. Firstly, we learned an unsupervised graph representation of the referral network. The GNN was trained with 1,134 nodes and 27,743 edges, and tested with the same number of nodes and 30,825 edges. Features used were physician gender and age, plus the centrality features for the social network-aware experiment. We trained for 20 epochs, layer sizes of 20 by 20 and a dropout rate of 0.3. A neural network model was trained for link prediction optimizing with Adam (learning rate of $1e-3$) over binary cross-entropy loss. In comparison, we also evaluate the link prediction task on the referral network in the *absence* of social network information.

To make these findings robust, we trained other two models. The first one, another neural network for the link prediction task on the referral network using, instead, Attri2Vec embeddings for the nodes [20]. This strategy was done to ensure that the results were not dependent of the node embedding model used. In contrast to GraphSAGE, the Attri2Vec node embeddings are learnt by performing a linear/non-linear mapping on node content features. The model generates node pairs of node target and node context and uses them to learn the representation on the target node with the existence of the context node. This is done using a deep learning model where the objective is to minimize the binary cross-entropy loss function with regards to the predicted node pair labels and true pair labels using stochastic gradient descent. The Attri2Vec model is implemented in the Stellargraph Python library as well. Our model was also trained with 1,134 nodes and 27,743 edges, and tested with the same number of nodes and 30,825 edges and were used the same features as the GraphSage GNN model. We trained for 10 epochs, layer sizes of 128 and without a normalization technique. The optimizer used was again Adam but with the learning rate of $1e-2$ over binary cross-entropy loss.

The second one model we trained served as a baseline. We use Node2Vec [21] to get the node embeddings and performed the link prediction task with Logistic Regression afterwards. The Node2Vec model was one of the firsts embedding models for nodes in a graph. The algorithm computes a vector representation of a node based on random walks in the graph. The neighborhood nodes of the graph are also sampled through deep random walks. This algorithm performs a biased random walk procedure in order to efficiently explore diverse neighborhoods. It is based on a similar principle as Word2Vec. The algorithm follows a strategy to sample the random walks that requires 4 major parameters: **Number of walks** – Number of random walks to be generated from each node in the graph; **Walk length** – How many nodes are in each random walk; **P** – Return hyperparameter and **Q** – Inout hyperparameter.

The embeddings were produced with a number of walks equals to 10, a walk length equals to 80 and P and Q equal to 1. The output dimension of the embeddings was 128. With the embeddings, we compute four different operations between two node embeddings for the link prediction task and saw what performs better when used in the Logistic regression classifier. We use L1 distance, L2 distance, the average and the hadmard operator. We train the model with 4623 examples; The model selection for the distance was done with 1541 examples and the test set was composed by 6848 examples.

In the table 1 is described the summary of the models constructed and their features.

Model	Train set numbers	Test set numbers	Hyperparameters
GraphSage	1134 nodes 27743 edges	1134 nodes 30825 edges	20 by 20 layer network Epochs 20 Dropout 0.3 Adam learning rate 1e-3
Attri2Vec	1134 nodes 27743 edges	1134 nodes 30825 edges	128 layer network Epochs 10 No regularization Adam learning rate 1e-2
Node2Vec	1134 nodes 4628 edges	1134 nodes 6848 edges	Number of walks 10 Wlak length 80 $P = Q = 1$ Adam learning rate 1e-2

Table 1. Summary of all the models trained. Node2Vec model served as a baseline.

3.3 Node embeddings

The node embeddings produced from the models described in the previous section (GraphSAGE and Attri2Vec) are in itself relevant to uncover patterns in the referral graph and to show, once more, the importance of doctor’s social features in the referral mechanism. Following this reasoning, we extracted the 20-dimensional vector embeddings of nodes produced by the GNN GraphSage model and used visualization techniques to get a sense about what was happening. We used several techniques to reduce the dimensionality of the embeddings from 20 to 2, namely U-MAP [22], T-SNE [23] and ISOMAP [24]. In this analysis we search for patterns in the node embeddings of the referral graph to find similar structures and explanations for the medical referrals’ behavior. The same analysis was done for the node embeddings produced by Attri2Vec.

4 RESULTS

4.1 Referral Prediction using Graph Embeddings

The results of the medical referral link prediction for all the models are presented in the tables 2. For all schemes, we show the accuracy and loss on the test set. The findings suggest that the information on social relationships of doctors does improve the predictive power of the model. Using the GraphSage model, accuracy increases about 0.18 with the added information about the social network of doctors and as expected, the loss suffered decreases as well. Similarly, the results concerning the Attri2Vec model, show also an increase in the accuracy and a decrease in the loss function in the presence of social features, confirming the contribution of these features for the referral prediction. A comparison between the models shows a slightly better performance of the Attri2Vec on the network with social information – being that the case, the node embedding analysis presented in the next section will focus the embeddings of this model.

4.2 Node Embeddings

The visualizations techniques of the Attri2Vec node embeddings used are intrinsically different and, therefore, the results presented are of a different nature, but there are some structures and takeaways in common. In the figures 9 and

Model	Social Network	Test loss	Test accuracy
GraphSage	Without information	1.0629	0.5232
	With information	0.6939	0.7072
Attri2Vec	Without information	0.6790	0.6876
	With information	0.6254	0.7268
Node2Vec	Without information	—	0.5828
	With information	—	0.6396

Table 2. Link prediction task metrics, when social network features are added to referral information with different model embeddings - GraphSage, Attri2Vec and Node2Vec as the baseline.

10, we present the node embeddings for the model with social features and the model without them respectively for the UMAP technique. We depict results in 4 subplots, each one highlighting one important aspect of the embeddings. Firstly, we highlight the role of the clinician — if the doctors are primary care doctors (**PC**) or specialty care doctors (**SC**); secondly, in another subplot, we colored the gender distribution over the embeddings. The third and fourth aspects are the age of the doctors — decomposed by decade of birth and the hospitals where the clinicians work.

These visuals evidence, clearly, that with the presence of social features the node embeddings present more structure — we can distinguish easily a more detailed cluster structure on the embeddings from the model with social features. Such a granular presentation could be a reason for the success of the link prediction task with added social features information. In contrast, the visuals without the social information, does not show such a defined structure.

Looking at the doctor’s role perspective, PC and SC doctors are well distributed across the main clusters of the embeddings in the two models, which makes sense due to the fact that the referrals are done from a PC to a SC, so if a cluster represents a group of doctors who are connected and in contact through referrals, it must have the two types of doctors.

In what concerns the hospitals where the doctors work in, we notice that the node embeddings without the social features do not exhibit a clear pattern we can extract. In the other hand, in the figures representing the node embeddings from the model with social features there are visible patterns. The doctors who work in the same hospital are mostly close together — signaling that this characteristic is useful for the referral mechanism, albeit not the only one, which is something expected beforehand. Moreover, the hospital where the doctors work in is not a direct input feature to the model. Hence, we can argue that the social features — the ones added to the input — can bear that information.

The age and gender perspectives are in itself revealing of some curious patterns. In both embeddings is clear a separation between female and male doctors but this structure is more expressive in the model without the social features. From this we can extract two main conclusions: (1) The model with social features uses information other than age and gender, and doing so, the embeddings are enriched, more complex and are a better proxy to the reality and (2) The gender, nevertheless, is also a relevant feature to the referrals and there is a underlying pattern of the doctors to referring to doctors of the same gender. Regarding the age of the doctors, this principle doesn’t apply as strongly. In

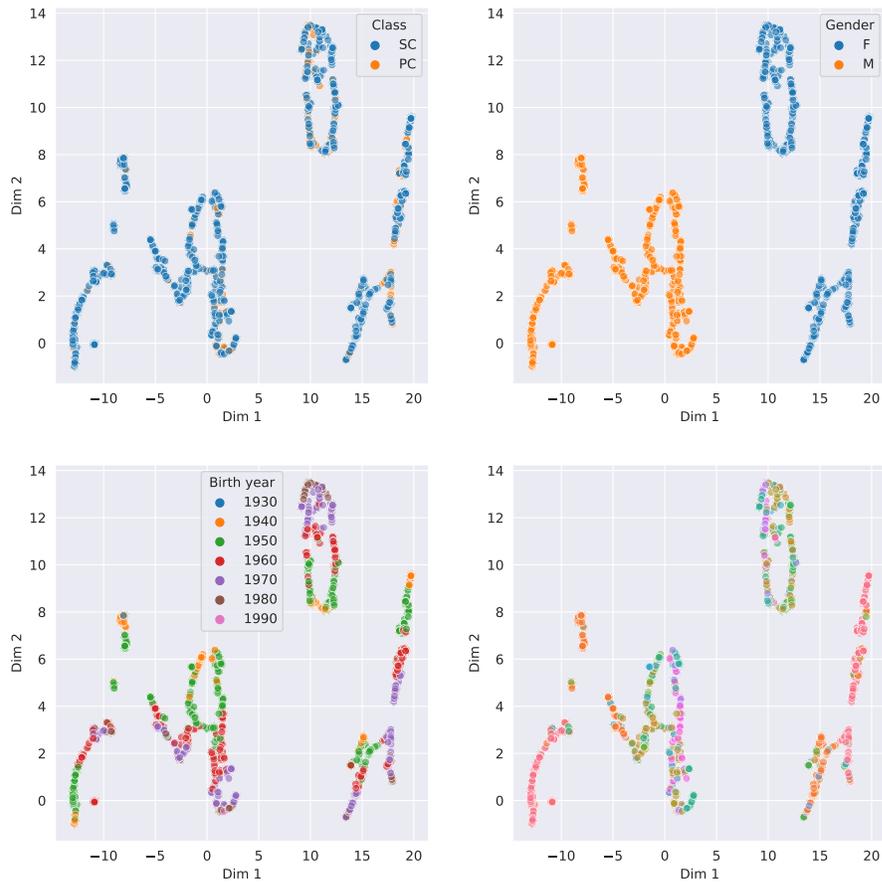


Fig. 9. Node embeddings extracted from the Attri2Vec model with the social features reduced to two dimension using U-MAP. In the top left plot the node embeddings separate primary care doctors (PC) from specialty doctors (SC); in the top right the embeddings are colored by gender; in the bottom left the node embedding is colored by birth year of the doctors whereas in the right, the colors represent each hospital. These 4 subplots clearly present some structure in the embeddings. Each cluster could represent a group of doctors who reference to each other within the group.

the embeddings from the model with social features the age don't seem to be good differentiator while in the simpler model that is the case. These results enforce the importance of the doctor's social features for the referral mechanism and show that without the social information the embeddings are not strong enough to explain the referrals. The node embeddings for the ISOMAP and TSNE techniques are very similar in what concerns the analysis done regarding the four dimensions described above.

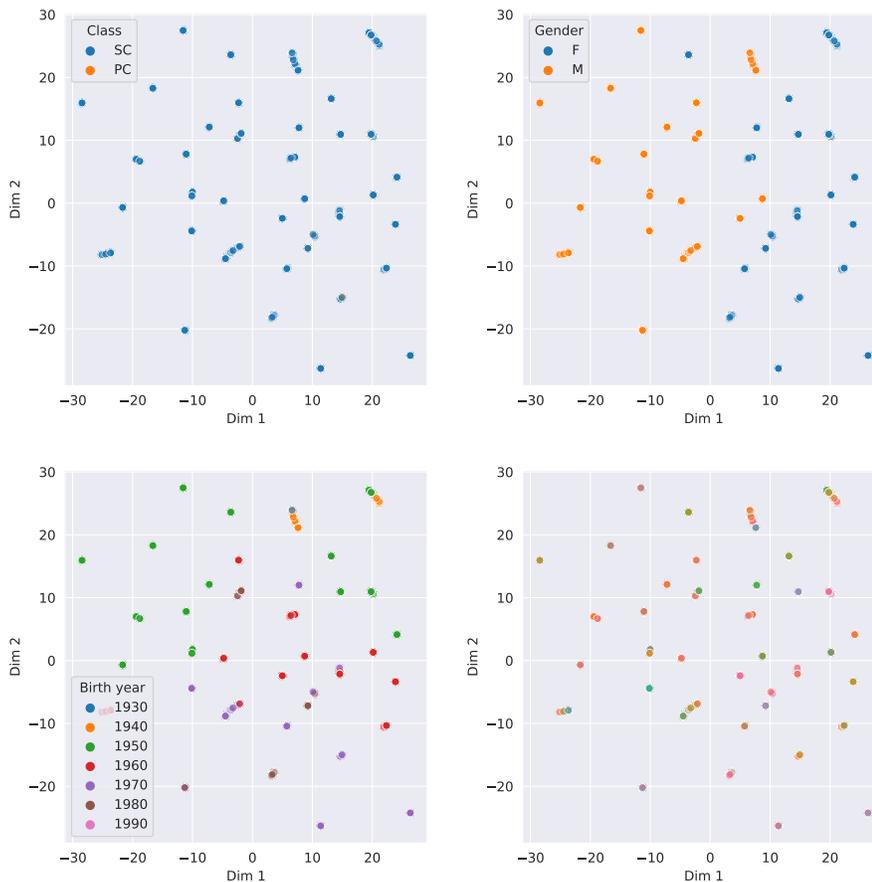


Fig. 10. Node embeddings extracted from the Attri2Vec model without the social features reduced to two dimension using U-MAP. In the top left plot the node embeddings separate primary care doctors (PC) from specialty doctors (SC); in the top right the embeddings are colored by gender; In the bottom left, the node embeddings are colored by birth years of the doctors whereas in the right, the colors represent each hospital. This embeddings compared to the embeddings from the model with social features lack more structure — which enforces the importance of social features for the referral prediction. With this embeddings is not clear if the birth age is, even in a small amount, correlated with the prediction of the referral links (despite the fact that in the ISOMAP and TSNE this relation occurs).

5 CONCLUSION

We uncovered patterns in the referral mechanism of physicians and tested the hypothesis that features computed from the social network of the clinicians are fundamental to the referral of patients from primary care to specialist care doctors. In this paper we backed this intuitive statement with evidence from a real dataset of a private European healthcare provider. We produced a link prediction model on the referral graph with social network features with improved accuracy.

As this work confirms, the referral mechanism is biased towards the social network of the doctors, which can have potential disadvantages and can not serve the best interest of the patients. A valuable future development would be the construction of an *unbiased* recommendation system that would serve the best patient’s interests.

REFERENCES

- [1] Shortell, S.M., Anderson, O.W.: The physician referral process: a theoretical perspective. *Health services research* **6**(1), 39 (1971)
- [2] Barnett, M.L., Christakis, N.A., O'Malley, A.J., Onnela, J.P., Keating, N.L., Landon, B.E.: Physician Patient-Sharing Networks and the Cost and Intensity of Care in US Hospitals. *Medical Care* **50**(2), 152–160 (2012). doi:10.1038/mp.2011.182.doi
- [3] Marshall, M.N.: How well do general practitioners and hospital consultants work together? a qualitative study of cooperation and conflict within the medical profession. *British Journal of General Practice* **48**(432), 1379–1382 (1998)
- [4] Muzzin, L.J.: Understanding the process of medical referral: part 5: communication. *Canadian Family Physician* **38**, 301 (1992)
- [5] Chan, T., Sabir, K., Sanhan, S., Sherbino, J.: Understanding the impact of residents' interpersonal relationships during emergency department referrals and consultations. *Journal of graduate medical education* **5**(4), 576 (2013)
- [6] Berendsen, A.J., Benneker, W.H., Meyboom-de Jong, B., Klazinga, N.S., Schuling, J.: Motives and preferences of general practitioners for new collaboration models with medical specialists: a qualitative study. *BMC Health Services Research* **7**(1), 1–9 (2007)
- [7] Janssen, M., Sagasser, M.H., Fluit, C.R., Assendelft, W.J., de Graaf, J., Scherpier, N.D.: Competencies to promote collaboration between primary and secondary care doctors: an integrative review. *BMC family practice* **21**(1), 1–13 (2020)
- [8] Ekwegh, U., Dean, J.: Improving care planning and communication for frail older persons across the primary–secondary care interface. *Future Healthcare Journal* **7**(3), 23 (2020)
- [9] Appel, A.P., de Santana, V.F., Moyano, L.G., Ito, M., Pinhanez, C.S.: A social network analysis framework for modeling health insurance claims data. *arXiv preprint arXiv:1802.07116* (2018)
- [10] Almansoori, W., Addam, O., Zarour, O., Elzohbi, M., Sarhan, A., Kaya, M., Rokne, J., Alhaji, R.: The power of social network construction and analysis for knowledge discovery in the medical referral process. *Journal of organizational computing and electronic commerce* **24**(2-3), 186–214 (2014)
- [11] An, C., O'Malley, A.J., Rockmore, D.N.: Referral paths in the U.S. physician network. *Applied Network Science* **3**(1) (2018). doi:10.1007/s41109-018-0081-4
- [12] Almansoori, W., Murshid, A., Xylogiannopoulos, K.F., Alhaji, R., Rokne, J.: Electronic medical referral system: Decision support and recommendation approach. In: 2012 IEEE 13th International Conference on Information Reuse & Integration (IRI), pp. 572–577 (2012). IEEE
- [13] Hamilton, W.L., Ying, R., Leskovec, J.: Inductive representation learning on large graphs. In: NIPS (2017)
- [14] Han, Q., Ji, M., Martinez de Rituerto de Troya, I., Gaur, M., Zejnilovic, L.: A hybrid recommender system for patient-doctor matchmaking in primary care. In: 2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA), pp. 481–490 (2018)
- [15] Neave, G., Amaral, A.: Higher Education in Portugal 1974 -2009. Springer, ??? (2009)
- [16] Landon, B.E., Keating, N.L., Onnela, J.P., Zaslavsky, A.M., Christakis, N.A., James O'Malley, A.: Patient-sharing networks of physicians and health care utilization and spending among medicare beneficiaries. *JAMA Internal Medicine* **178**(1), 66–73 (2018). doi:10.1001/jamainternmed.2017.5034
- [17] An, C., O'Malley, J.A., Rockmore, D.N., Stock, C.D.: Analysis of the U.S. Patient Referral Network. *Statistics in medicine* **37**(5), 847–866 (2018). doi:10.1002/cncr.27633.Percutaneous
- [18] 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. KDD '16, pp. 855–864 (2016)
- [19] Data61, C.: StellarGraph Machine Learning Library. GitHub (2018)
- [20] Zhang, Y.J.Z.X.e.a. D.: Attributed network embedding via subspace discovery. *Data Mining and Knowledge Discovery* (33), 1953–1980 (2019). doi:10.1007/s10618-019-00650-2
- [21] A. Grover, J.L.: Node2vec: Scalable feature learning for networks. *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)* (2016)
- [22] McInnes, L., Healy, J., Saul, N., Großberger, L.: Umap: Uniform manifold approximation and projection. *Journal of Open Source Software* **3**(29) (2018)
- [23] Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(11) (2008)
- [24] Tenenbaum, J.B., De Silva, V., Langford, J.C.: A global geometric framework for nonlinear dimensionality reduction. *science* **290**(5500), 2319–2323 (2000)

ACKNOWLEDGMENTS

This work was funded by Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2019, UIDB/00124/2020, Social Sciences DataLab, PINFRA/22209/2016, UIDB/04516/2020 NOVA LINCS), POR Lisboa and POR Norte (Social Sciences DataLab, PINFRA/22209/2016), and (DSAIPA/AI/0087/2018).