

The Structure of Interdisciplinary Science: Uncovering and Explaining Roles in Citation Graphs

Eoghan Cunningham^{1,2} and Derek Greene^{1,2}

¹ School of Computer Science, University College Dublin, Ireland

² Insight Centre for Data Analytics, University College Dublin, Ireland

Abstract. Role discovery is the task of dividing the set of nodes on a graph into classes of structurally similar roles. Modern strategies for role discovery typically rely on graph embedding techniques, which are capable of recognising complex local structures. However, when working with large, real-world networks, it is difficult to interpret or validate a set of roles identified according to these methods. In this work, motivated by advancements in the field of explainable artificial intelligence (XAI), we propose a new framework for interpreting role assignments on large graphs using small subgraph structures known as graphlets. We demonstrate our methods on a large, multidisciplinary citation network, where we successfully identify a number of important citation patterns which reflect interdisciplinary research.

Keywords: role discovery, node embedding, citation network, explainable artificial intelligence

1 Introduction

In light of the perceived importance of interdisciplinary research, many studies have been conducted that quantify article interdisciplinarity in an effort to identify relevant research trends and to explore their impact. The most widely-accepted methods for measuring interdisciplinarity assess *knowledge integration* or *knowledge diffusion* using citation information, thus measuring interdisciplinarity as some function of the balance, diversity, and dissimilarity of the disciplines identified in an article’s cited papers [17, 19] or citing papers [16, 28]. To implement these metrics, each paper in a research corpus must be assigned to an explicit research topic or subject category, for which sources are numerous, inconsistent, and sometimes unavailable. Subject categories are most commonly assigned to papers according to the journals in which they are published. However, these assignments rarely agree with underlying citation community structure [17]. There is evidence that interdisciplinary research can be identified in a corpus according solely to the citation structure. Specifically, it has been shown that frameworks that encode the *structural role* of articles in a citation graph can predict interdisciplinary interactions more accurately than

those that encode only the proximity between papers [6]. In light of this, we explore the potential for modern graph learning methods to identify the citation structures associated with interdisciplinary research.

Numerous approaches have been proposed in the literature for the task of structural role discovery on graphs, where nodes on a graph are divided into classes of *structurally equivalent* nodes [21]. Early approaches in this area relied on graph sub-structures known as *graphlets* or *motifs* [12]. In large graphs, where we wish to identify higher-order structural features, counting large graphlets is very expensive. The majority of recent approaches employ a representation learning technique known as node embedding, which substitutes graphlets with alternative structural features like degree distributions [20] and diffusion wavelet patterns [8]. These methods are designed to learn higher-order structural relationships than those that can be discovered by small graphlets. However, in many cases, these alternative approaches come at the cost of interpretability. When applied to graphs that are too large to be visualised reasonably, it is difficult to understand the substantive meaning of a given set of structural roles.

While embedding methods for role analysis have previously been shown to be capable of grouping nodes into known roles or structures (such as those in synthetic graphs or transport networks [8, 20]), it remains unclear as to how these roles should be interpreted or validated when applied to real-world graphs with unknown roles. Moreover, different role discovery methods learn different sets of structural roles, depending on the application many or none of these clusterings may be valid. As such, it is critical that we can compare the roles discovered by different methods.

The core contribution of this work is a new framework for explaining a set of discovered roles using graphlet *orbits*, as described in detail in Section 3. Later in Section 4 we apply this framework to a large, multidisciplinary citation network to extract and interpret sets of structural roles. In the case of each paper, we compute Rao-Stirling diversity scores to indicate an articles interdisciplinarity. Crucially, in addition to interpreting a set of candidate roles, we explore the distributions of IDR scores assigned to papers in different clustering to assess if any of the candidate clusterings have grouped papers according to their interdisciplinarity.

2 Related Work

2.1 Measuring Interdisciplinarity

Interdisciplinary research is most commonly defined as research activity which integrates expertise, data or methodologies from two or more distinct disciplines. Accordingly, many studies assign an interdisciplinary research (IDR) score to a research article, which is calculated as a function of the balance, diversity and dissimilarity of the disciplines identified in the articles reference list [17, 19]. Alternatively, some studies compute a similar score on the disciplines identified in an article’s *citing* literature, instead measuring IDR according to an articles impact/influence across multiple disciplines [16, 28]. A popular function for

measuring IDR is the Rao-Stirling Diversity index [27]

$$D = \sum_{i,j(i \neq j)} p_i p_j d_{ij} \quad (1)$$

where IDR is measured as a pairwise summation over all pairs of disciplines identified in a set of articles cited by (or citing) some focal paper. Here p_i and p_j represent the proportion of references to disciplines i and j respectively, while d_{ij} refers to some precomputed distance between the disciplines. This metric, and its many variants, are reliant on explicit topic or discipline categories for research papers, such as those provided by Web of Science, Scopus, or Microsoft Academic. Such explicit categorisations for research papers, especially those assigned at a journal level, are problematic [1, 13]. Moreover, the inconsistencies evident across many of these subject taxonomies [25] may confirm that no singular, correct categorisation exists. However, recent graph learning methods may be capable of identifying the citation structures associated with interdisciplinary research, without any knowledge of explicit discipline categorisation.

2.2 Local Structure and Role Embeddings

Research developed in social science studied local graph structures using small graph patterns such as triads, cycles and stars [15]. More recent research in computational science further developed these methods and proposed the term ‘motif’ – a subgraph pattern (or *graphlet*) which is significantly over-represented in a graph [12]. Motif and graphlet counts represent powerful features for expressing graph structure, and have been employed in graph learning tasks such as node classification and anomaly detection [7]. When using graphlet counts to describe nodes, it is more information to consider the position in the graphlet at which the node occurs. Equivalent positions within graphlets are called *automorphism orbits* or simply orbits. Figure 1 illustrates a subset of graphlets with 2, 3, 4 and 5 nodes, and includes each of the distinct orbits on these graphlets, as they were enumerated by [18].

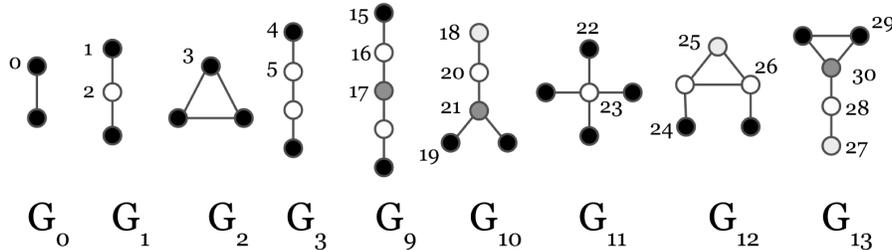


Fig. 1: Examples of graphlets with 2, 3, 4, and 5 nodes. Graphlets and orbits are enumerated according to [18]

Role discovery is the task of grouping nodes which share similar structural patterns in a graph into distinct classes [21]. Many modern approaches to role discovery rely on graph embedding, where nodes are transformed into low-dimensional vector representations [3, 10, 20, 8]. Specifically, graph embedding methods for the purpose of role discovery (often called “role embedding”) learn dense vector representations for nodes such that nodes that are structurally similar will be embedded nearby in the embedding space (i.e., will have similar representations). A clustering of the role embedding space thus represents a set of discovered roles. However, if this network cannot be visualised, it is difficult to interpret the roles. Moreover, with numerous approaches to role embedding (e.g. [3, 8, 10, 20]), and many possible clusterings of each embedding space, we require some approach to explain a set of discovered roles so that they can be compared and validated.

2.3 Explanation via Surrogate Models

In recent years, much research has been conducted in the field of *explainable artificial intelligence* (XAI) towards a goal of understanding and explaining so-called ‘*black box*’ models [2]. One popular approach is to use a global surrogate model: “an interpretable model that is trained to approximate the predictions of a black box model” [14]. As such, the decisions made by a previously uninterpretable system can be explained through interpretations coming from the surrogate model.

Some classification models (such as logistic regression models or decision tree-based classifiers) are interpretable by definition, as any input feature’s effect on the model’s classification can be measured (for example, using regression coefficients). However, many model-agnostic methods of interpretation have also been developed. Generally, these methods propose different means of perturbing the input data, and exploring its effect on the model’s output. For example, *Partial Dependence Plots* (PDPs) [9] offer a graphical representation of a model’s prediction as a function of its arguments or inputs. Broadly speaking, a PDP shows how a feature influences the prediction of a model, on average, for every value that it can take. The *Accumulated Local Effects* plot (ALE) [4] provides a similar illustration of a feature’s effect, except it accounts for possible correlations between features, and as a result is more robust than the PDP in some applications. In 2021, global surrogate models were employed to provide feature-based explanations for communities detected on graphs [24]. Specifically, this work assessed the importance of different graph features using *permutation importance* [5], where the values for a feature are permuted throughout the dataset and the effect on model performance is measured to indicate its importance.

3 Methods

In this section we outline a general framework for uncovering and evaluating structural roles on a graph. The complete process is illustrated in Figure 2.

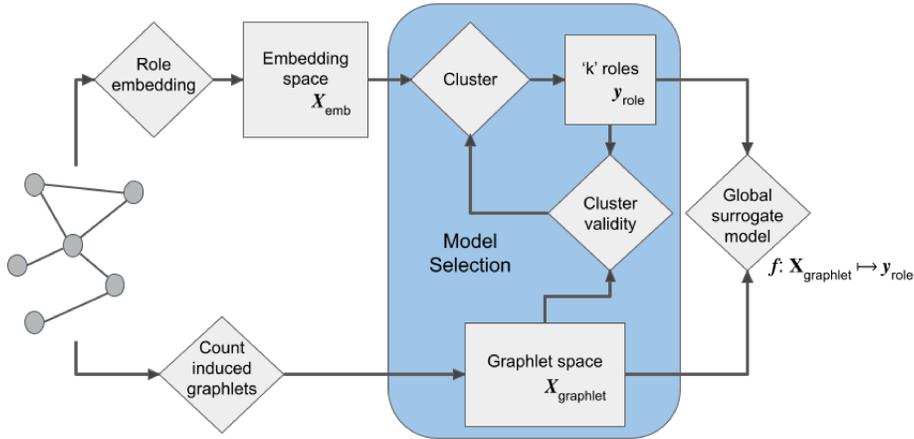


Fig. 2: Overview of the workflow for the proposed framework.

We begin with a graph $G = (V, E)$, from which we wish to identify a set of discrete structural roles. We employ a role embedding algorithm to map each node $v \in V$ to an embedding space $X_{emb} = \mathbb{R}^{128}$, thus a clustering of X_{emb} is considered a set of discovered roles y_{role} . Additionally, we represent the same nodes in the graphlet space $X_{graphlet} = \mathbb{Z}^{73}$ derived from G , where each node v is represented by a *bag-of-orbits* vector $x^v = \{x_0^v, x_1^v, \dots, x_{72}^v\}$, with x_i^v denoting the number of times node v is involved in induced graphlet orbit i . We use a vocabulary of 73 orbits, which we count on graphlets of size 2 to 5 using the ORCA method [11]. We refer to graphlets and orbits according to the enumeration by [18]. The graphlet space is first used to validate any set of roles we identify in the graph. By clustering the embedding space, we group nodes into k discrete roles, which we can evaluate using cluster validity metrics calculated on the graphlet space. Employing various role embedding algorithms and clustering methods, we identify a set of candidate clusterings (or roles) according to the separation they achieve in the graphlet space.

Graphlet-orbit counts are a powerful language for describing the structural features of nodes on a graph [7, 18]. As such, a clustering (according to node embeddings) which is valid in the graphlet-orbit space should offer a reasonable set of structural roles. Additionally, we employ graphlet orbits to interpret a set of discovered roles. For a candidate clustering y_{role} , we fit a surrogate model $f: X_{graphlet} \mapsto y_{role}$. By modelling the role assignment in the graphlet space, we can explore the feature importance and effect of the different graphlet orbits in role assignments, according to many *model-agnostic* explanation techniques from the field of XAI [14]. In Section 4, we demonstrate graphlet-based explanation using permutation importance [5] and accumulated local effects plots [4]. As we will show, highlighting important or discriminatory orbits can offer a visual means of understanding the structure of a role in the graph.

4 Application

In this section we apply the methods described in Section 3 to extract and interpret sets of structural roles in a citation network. In the case of each paper, we compute Rao-Stirling diversity scores to indicate an articles interdisciplinarity and thus explore the distributions of IDR scores assigned to papers in different roles. We identify a set of roles which has grouped papers according to their interdisciplinarity. Finally, we use graphlets to interpret the structure of these more interdisciplinary roles, and highlight certain citation structures which are specific to interdisciplinary research.

4.1 Data

In order to discover the citation structures of interdisciplinary research, we require a large, dense citation network that contains research from a diverse set of disciplines. In addition, we require that each paper can be assigned to a subject category or discipline, according to an established taxonomy. We construct a novel citation network using Microsoft Academic Graph [26] citation data from a seed set of journal papers. This sets consist of samples of articles from Scopus indexed journals, stratified according to their All Science Journal Categories (ASJC). The graph contains samples of 1,500 articles published between 2017 and 2018 in Scopus indexed journals with the ASJCs ‘*Computer Science*’, ‘*Mathematics*’, ‘*Medicine*’, ‘*Chemistry*’, ‘*Social Sciences*’, ‘*Neuroscience*’, ‘*Engineering*’, and ‘*Biochemistry, Genetics and Molecular Biology*’. We maximise the completeness of the graph by including all available referenced articles that are published in Scopus indexed journals. In this manner, we produce a dense, multidisciplinary citation network, such that each article can be categorised according to the ASJC of the journal in which it was published. Later, these discipline categories can be used to compute article interdisciplinarity according to Rao-Stirling diversity of disciplines identified in both an articles citing and cited papers. In total, the citation graph contains 41,895 papers (nodes) and 129,159 citations (undirected edges).

4.2 Role Discovery

For each article in the citation graph, we learn role embeddings according to 4 approaches: (i) *Role2Vec* [3], (ii) *Struc2Vec* [20], (iii) *RolX* [10], (iv) *GraphWave* [8], and we cluster each embedding space using k -means clustering for values of k between 2 and 19. Articles clustered according to their role embeddings represent a set of structural roles.

4.3 Role Interpretation

Validation. Figure 3 shows the cluster validity of the roles discovered according to the 4 role embedding methods. Each set of roles represents a k -means clustering of the embedding space, which is then transformed to the graphlet space,

where we measure the validity of the clustering using silhouette score [23]. We plot the silhouette score for each embedding for values of k in the range [2, 19]. Silhouette scores can take a value in the range [-1,1]. A high score represents dense, separate clusters, a score of 0 indicates an overlapping clustering, while a negative score indicates an incorrect clustering. According to these scores, we choose 3 candidate roles to demonstrate interpretation and explanation: (i) Struc2Vec ($k = 6$) which is an outlier in the Struc2vec roles and achieves an overlapping clustering, (ii) RolX ($k = 3$) which has the highest silhouette score for all approaches with more than 2 clusters, and (iii) GraphWave ($k = 3$) which achieves a positive score.

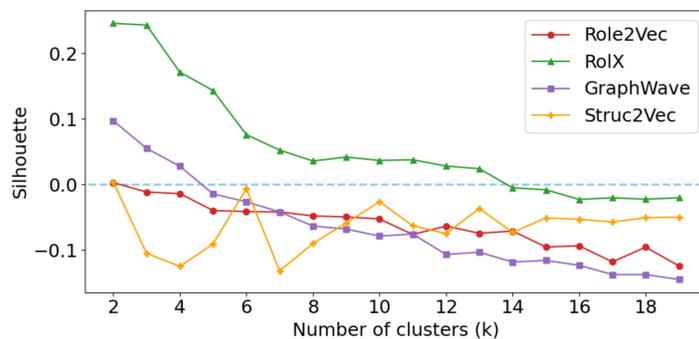


Fig. 3: Silhouette score for k -means clusterings of the different embedding spaces. Silhouettes are calculated according to the clusterings if we describe each node by their bag-of-orbits vector, rather than their embedding.

Surrogate models. We fit a random forest classifier to model role assignments according to the bag-of-orbits vectors for each node. That is, we learn a function $f : X_{graphlet} \mapsto y_{role}$ for each of the candidate roles y_{role} . We choose a random forest classifier as we anticipate that structural roles may be non-linear in the graphlet space and may rely on combinations and interactions of features to model higher-order structures. As the orbit counts follow a power-law distribution, we log-transform all features in the graphlet space. Table 1 reports the 5 most informative features for each model according to permutation importance. In the case of the Struc2Vec ($k = 6$) roles, the overlap between clusters in the graphlet space is evident. The only informative features (with non-zero permutation importance) are small, local orbits – the approach is blind to deeper, more complex structures. The RolX ($k = 3$) approach, which best separates the nodes in the graphlet space, has grouped nodes according to larger, simple structures. The most informative orbits: 15, 4, and 1, each refer to chains of varying length. Finally, GraphWave ($k = 3$) appears to have grouped the nodes according to more complex, higher-order structures. Many of the features important to role classification in the GraphWave case (27, 24, and 18) contain combinations of chains, stars and cliques. Depending on the domain or application in which we

employ role discovery, any one of these sets of roles may be the most valid or useful. However, without modelling the role assignments in the graphlet space, we are unable to understand which structures are being grouped in the discovered roles. We will use the GraphWave roles to demonstrate further explanation and explore the structure of interdisciplinary papers.

| | Struc2Vec ($k = 6$) | RolX ($k = 3$) | GraphWave ($k = 3$) |
|---|-------------------------|-------------------------|-------------------------|
| 1 | 0 (0.111 \pm 0.0005) | 15 (0.113 \pm 0.0016) | 27 (0.022 \pm 0.0006) |
| 2 | 2 (0.010 \pm 0.0001) | 4 (0.007 \pm 0.0002) | 15 (0.014 \pm 0.0002) |
| 3 | 3 (0.006 \pm 0.0001) | 1 (0.006 \pm 0.0003) | 17 (0.013 \pm 0.0006) |
| 4 | 5 (0.002 \pm 0.0001) | 27 (0.005 \pm 0.0002) | 24 (0.013 \pm 0.0004) |
| 5 | 16 (0.000 \pm 0.0000) | 19 (0.004 \pm 0.0003) | 18 (0.013 \pm 0.0004) |

Table 1: Top 5 most important features for each surrogate model. Orbits are ranked by permutation importance, which is included in parentheses.

Roles and their structure. The GraphWave method clusters 35,136 papers into role 0, 16,453 papers into role 1, and 306 papers into role 2. Figure 4 contains Accumulated Local Effects (ALE) [4] plots for 3 features/orbits and their effect on classification to each of the 3 GraphWave roles. We illustrate the ALE of orbits 27 (the end of a chain adjacent to a clique) and 17 (the middle of a long chain) as two of the most important structures (according to permutation importance). We also include the ALE of orbit 0 (node degree) as a valuable reference; it is useful to confirm that the roles are indeed separating nodes according to more complex features and not simply the number of edges. The ALE plot for orbit 27 shows that for low-to-mid values of that orbit count, a node will be classified as role 0. However, if a node’s count for orbit 27 exceeds a threshold, it will be classified as role 1. We suppose 2 scenarios when a focal node’s count for orbit 27 (a chain adjacent to a clique) will become large: (i) the node is adjacent to a large community – each triangle in which the node at position 30 participates will increase the count, (ii) the node exists at the center of a barbell graphlet, i.e. on a longer chain between two or more communities. We illustrate these scenarios in Figure 5. There should exist some threshold value of 27 beyond which a node must exist on the chain between two communities. For example, if a focal node has a count for orbit 27 that is greater than the count of triangles (orbit 3) for the node at position 30, then the focal node must be adjacent to a second community (scenario (iii)). This threshold will be represented by the greatest value of orbit 3 in the graph. We include this value for reference in Figure 4. Beyond this threshold, a node is more likely to be classified in role 1. Accordingly, we conclude that a node/paper that is on the end of a chain adjacent to a community will be assigned to role 0, while a node that exists on a bridge between two communities will be assigned to role 1.

In order to identify the structures specific to the smallest role (role 3), we can fit another surrogate model on only the nodes in cluster 1 and 2. ALE plots for this model are included in Figure 6. In this case we find that orbit 27 does not meaningfully distinguish between the two roles. Instead, orbit 18 (the end

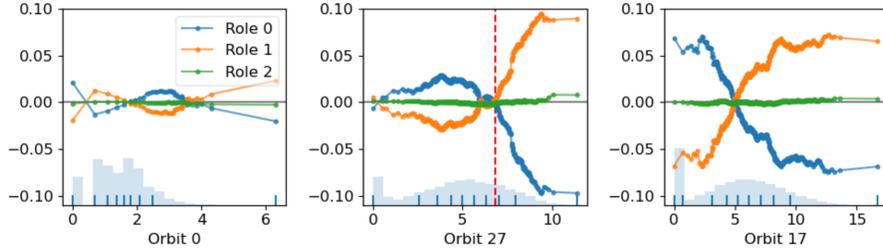


Fig. 4: Accumulated Local Effect [4] plots for a surrogate model which classifies nodes to GraphWave roles according to graphlet orbit counts. The figure shows the effect of 3 features on role classification: orbits 27 and 17; the most important features as measured by permutation importance, and orbit 0 (node degree); which we include for reference. In the case of orbit 27, we highlight the maximum value of orbit 3 that was observed in the graph.

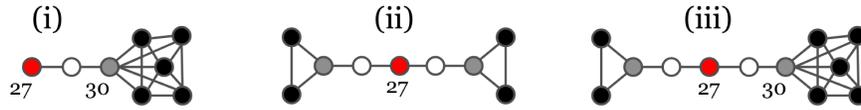


Fig. 5: Higher order structures containing orbit 27. Here (i) and (ii) represent likely structures for nodes with high counts for orbit 27. When the count for orbit 27 exceeds a threshold, we infer structure (iii).

of a chain adjacent to a star) is the most informative feature, and, for very high values of this orbit count, nodes will be assigned to role 2. Such nodes likely represent the centre of a bridge between large communities that are less densely connected (i.e., containing many open triads). We conclude this to be an important structure for role 3.

4.4 Interdisciplinary Roles

Figure 7 plots the IDR scores for the papers in each of the GraphWave roles. IDR is calculated according to the Rao-Stirling diversity of the ASJC categories identified in an article's citing papers. As Rao-Stirling IDR scores may be biased according to the number of articles in the summation, we bin nodes by degree, and plot IDR distributions for each role, within each bin. Specifically, we log-transform the node degrees and group nodes into 10 bins of equal width, within which we plot IDR distributions for each role if the bin contains more than 50 papers from each role. According to these plots, we note that the structural roles identified by GraphWave have grouped papers into clusters with different IDR distributions. Even when we account for node degree (a potential bias of the Rao-Stirling IDR score), papers assigned to structural roles 1 and 2 have

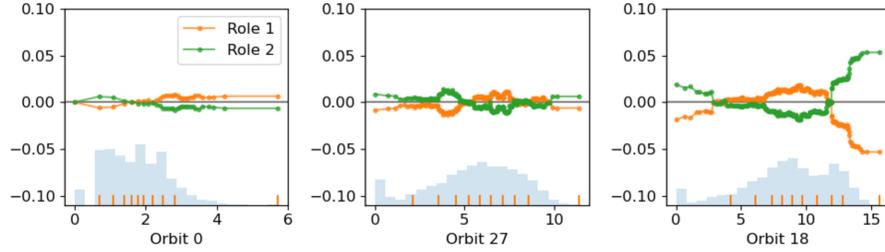


Fig. 6: Accumulated Local Effect [4] plots for a surrogate model which classifies nodes to GraphWave roles 1 or 2 according to graphlet orbit counts. The figure plots the effect of 3 features: orbits 0; node degree and 27; previously the most important feature for the global model, and orbit 18; the most important feature in the surrogate model.

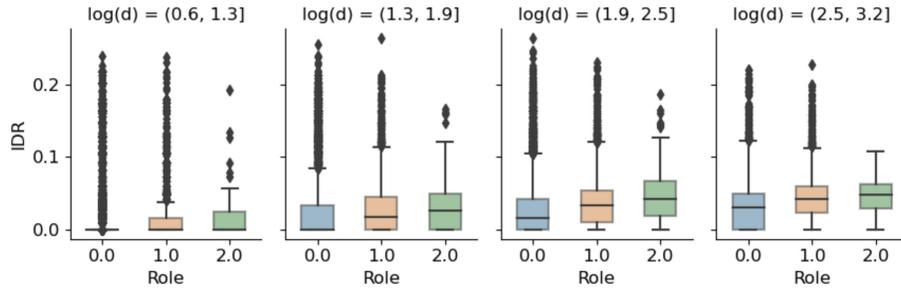


Fig. 7: Interdisciplinarity scores (IDR) computed for papers in each of the 3 GraphWave clusters. IDR is computed as the Rao-Stirling [27] diversity of the research disciplines identified in an articles citing papers. We bin papers/nodes according to the log of their degree, and compare distributions within each bin.

consistently greater IDR distributions than those assigned to role 0. We recall some of the important structures that were identified for roles 1 and 2: (1) a bridge between densely connected communities, and (2) a bridge between large, sparsely connected communities. We conclude these to be citation important structures associated with interdisciplinary research.

5 Discussion and Conclusions

Many modern methods for role discovery in graphs rely on node embeddings [22]. We have applied a number of these methods to learn structural roles on a research citation network. While these methods have previously been shown to be capable of grouping nodes into known roles (e.g., in synthetic graphs or transport networks [20, 8]), it remains unclear how roles could be understood or validated when applied to graphs with unknown roles. In this work we outlined

a framework for interpreting a set of discovered roles using graphlets and orbits. We leveraged methods from the field of explainable AI to explore the subgraph structures that are specific to discovered roles. We demonstrated our approach on a citation network, where we identified important structures specific to interdisciplinary research. It is evident in our analysis that different role discovery methods learn different sets of structural roles. In different applications many or all of these clusterings may be valid, but it is critical that we can compare the roles discovered by different methods. While our framework is general, and applicable to explanation and validation in all role discovery tasks, we highlight the utility of structural role embeddings in mapping interdisciplinary research.

For the task of identifying and mapping IDR interactions, structural paper embeddings could be augmented by considering additional, non-structural information, such as article or abstract text. This could provide a richer paper representation, without imposing a predefined or static disciplinary classification on the graph. There is also scope for improving upon our proposed framework. Many *model-agnostic* approaches have been developed for explaining surrogate models [14], which could be applied to interpret the role assignments in the graphlet space. For example, second-order effects of pairs of features can be calculated in a similar manner to the ALE analysis we have included [4]. Combinations of graphlets could be highly effective in modelling higher-order, more complex graph structures. One possible limitation of our current framework is the number of correlated features in the graphlet space. In future applications it may be possible to reduce the set of graphlet orbits to a minimal set of uncorrelated features via traditional feature selection techniques.

Acknowledgments. This research was supported by Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289_P2.

References

1. G. Abramo, C. A. D’Angelo, and L. Zhang. A comparison of two approaches for measuring interdisciplinary research output: The disciplinary diversity of authors vs the disciplinary diversity of the reference list. *Journal of Informetrics*, 12(4):1182–1193, 2018.
2. A. Adadi and M. Berrada. Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE Access*, 6:52138–52160, 2018.
3. N. K. Ahmed, R. A. Rossi, J. B. Lee, T. L. Willke, R. Zhou, X. Kong, and H. El-dardiry. role2vec: Role-based network embeddings. In *Proc. DLG KDD*, pages 1–7, 2019.
4. D. W. Apley and J. Zhu. Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(4):1059–1086, 2020.
5. L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
6. E. Cunningham and D. Greene. Assessing network representations for identifying interdisciplinarity. *arXiv preprint arXiv:2203.12455*, 2022.
7. P. Cunningham, M. Harrigan, G. Wu, and D. O’Callaghan. Characterizing ego-networks using motifs. *Network Science*, 1(2):170–190, 2013.

8. C. Donnat, M. Zitnik, D. Hallac, and J. Leskovec. Learning structural node embeddings via diffusion wavelets. In *Proc. 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1320–1329, 2018.
9. J. H. Friedman. Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, pages 1189–1232, 2001.
10. K. Henderson, B. Gallagher, T. Eliassi-Rad, H. Tong, S. Basu, L. Akoglu, D. Koutra, C. Faloutsos, and L. Li. Rolx: structural role extraction & mining in large graphs. In *Proc. 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1231–1239, 2012.
11. T. Hočevár and J. Demšar. Combinatorial algorithm for counting small induced graphs and orbits. *PLoS one*, 12(2):e0171428, 2017.
12. R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.
13. S. Milojević. Practical method to reclassify Web of Science articles into unique subject categories and broad disciplines. *Quantitative Science Studies*, 1(1):183–206, 02 2020.
14. C. Molnar. *Interpretable machine learning*. Lulu.com, 2020.
15. J. L. Moreno. Who shall survive?: A new approach to the problem of human interrelations. 1934.
16. A. Porter and D. Chubin. An indicator of cross-disciplinary research. *Scientometrics*, 8(3-4):161–176, 1985.
17. A. Porter and I. Rafols. Is science becoming more interdisciplinary? measuring and mapping six research fields over time. *Scientometrics*, 81(3):719–745, 2009.
18. N. Pržulj. Biological network comparison using graphlet degree distribution. *Bioinformatics*, 23(2):e177–e183, 2007.
19. I. Rafols and M. Meyer. Diversity and network coherence as indicators of interdisciplinarity: case studies in bionanoscience. *Scientometrics*, 82(2):263–287, 2010.
20. L. F. Ribeiro, P. H. Saverese, and D. R. Figueiredo. struc2vec: Learning node representations from structural identity. In *Proc. ACM KDD 2017*, 2017.
21. R. A. Rossi and N. K. Ahmed. Role discovery in networks. *IEEE Transactions on Knowledge and Data Engineering*, 27(4):1112–1131, 2014.
22. R. A. Rossi, D. Jin, S. Kim, N. K. Ahmed, D. Koutra, and J. B. Lee. On proximity and structural role-based embeddings in networks: Misconceptions, techniques, and applications. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 14(5):1–37, 2020.
23. P. J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J. Computational and Applied Mathematics*, 20:53–65, 1987.
24. S. Sadler, D. Greene, and D. Archambault. Selecting informative features for post-hoc community explanation. In *International Conference on Complex Networks and Their Applications*, pages 297–308. Springer, 2021.
25. Z. Shen, F. Chen, L. Yang, and J. Wu. Node2vec representation for clustering journals and as a possible measure of diversity. *Journal of Data and Information Science*, 4(2):79, 2019.
26. A. Sinha, Z. Shen, Y. Song, H. Ma, D. Eide, B.-J. Hsu, and K. Wang. An overview of Microsoft Academic Service (MAS) and applications. In *Proc. 24th international conference on World Wide Web*, pages 243–246, 2015.
27. A. Stirling. A general framework for analysing diversity in science, technology and society. *Journal of the Royal Society Interface*, 4(15):707–719, 2007.
28. R. Van Noorden. Interdisciplinary research by the numbers. *Nature*, 525(7569):306–307, 2015.