

# Harnessing the Power of Ego Network Layers for Link Prediction in Online Social Networks

MUSTAFA TOPRAK, CHIARA BOLDRINI, ANDREA PASSARELLA, and MARCO CONTI, CNR-IIT, Italy

Being able to recommend links between users in online social networks is important both for the platforms themselves and for third parties leveraging social media information to grow their business. Predictions are typically based on unsupervised or supervised learning, often leveraging simple yet effective graph topological information, such as the number of common neighbors. However, we argue that richer information about personal social structure of individuals might lead to better predictions. In this paper, we propose to leverage well-established social cognitive theories to improve link prediction performance. According to these theories, individuals arrange their social relationships along, on average, five concentric circles of decreasing intimacy. We postulate that relationships in different circles have different importance in predicting new links. In order to validate this claim, we focus on popular prediction algorithms (both unsupervised and supervised) and we extend them to include social-circles awareness. We validate the prediction performance of these social-aware algorithms against their baseline versions, leveraging two Twitter datasets comprising a community of video gamers and generic users. We show that social-awareness generally provides significant improvements in the prediction performance. In addition, social-awareness is able to deliver good predictions also in cases where resources (e.g., storage, computing) are limited and other approaches perform poorly. Finally, we show that social-awareness can be used in place of using a classifier (which may be costly or impractical) for targeting a specific category of users.

CCS Concepts: • **Human-centered computing** → **Social recommendation; Social networks**; • **Computing methodologies** → *Supervised learning; Unsupervised learning*.

Additional Key Words and Phrases: link prediction, social circles, Dunbar's model, Twitter

## ACM Reference Format:

Mustafa Toprak, Chiara Boldrini, Andrea Passarella, and Marco Conti. 2020. Harnessing the Power of Ego Network Layers for Link Prediction in Online Social Networks. 1, 1 (March 2020), 29 pages. <https://doi.org/10.1145/1122445.1122456>

## 1 INTRODUCTION

In 2019, the number of social media users worldwide has reached 3.484 billion, up 9% on a year-on-year basis [19]. Social media are now large ecosystems, where we gather to interact with people we already know or to join communities of people that share our interests. Examples of the latter are Science Twitter<sup>1</sup> or Bookstagram<sup>2</sup>, which are simply open, unofficial sub-communities, within Twitter and Instagram, of accounts (mostly) talking about a specific topic. Clearly, being able to

<sup>1</sup><https://blogs.plos.org/absolutely-maybe/2017/04/28/whos-who-on-science-twitter-and-who-counts/>

<sup>2</sup>[https://www.huffpost.com/entry/bookstagram-how-readers-changed-the-way-we-use-instagram\\_b\\_59f0aaa2e4b01ecaf1a3e867](https://www.huffpost.com/entry/bookstagram-how-readers-changed-the-way-we-use-instagram_b_59f0aaa2e4b01ecaf1a3e867)

---

Authors' address: Mustafa Toprak, [mustafa.toprak@iit.cnr.it](mailto:mustafa.toprak@iit.cnr.it); Chiara Boldrini, [chiara.boldrini@iit.cnr.it](mailto:chiara.boldrini@iit.cnr.it); Andrea Passarella, [andrea.passarella@iit.cnr.it](mailto:andrea.passarella@iit.cnr.it); Marco Conti, [marco.conti@iit.cnr.it](mailto:marco.conti@iit.cnr.it), CNR-IIT, Via G. Moruzzi 1, Pisa, 56124, Italy.

---

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](mailto:permissions@acm.org).

© 2020 Association for Computing Machinery.

XXXX-XXXX/2020/3-ART \$15.00

<https://doi.org/10.1145/1122445.1122456>

pair together people with similar interests is crucial for social-driven and content-driven platforms like Facebook, Twitter, Instagram, Youtube, etc. This spurred intense research on link prediction, whereby algorithms suggest potential new links with users that are *similar* to each other. Solving this problem is equivalent to finding a needle in a haystack: online social networks (OSN) feature millions of users, the number of potential links between them grows with the square of the number of users, but the actually existing links are quite sparse.

Link prediction algorithms can be broadly divided into unsupervised and supervised algorithms. Unsupervised methods typically associate a score to each potential new link, and then only the links with a score above a predefined threshold – or the top-K ones – are used for recommendation. Within this class, the most popular strategy for prediction is based on estimating the similarity between pairs of nodes, and then leveraging this similarity for prediction (the higher the similarity, the higher the probability that the two nodes should be connected). This similarity is typically computed using topological information about the network, be it local information (such as the number of common neighbors) or global one (such as the length of paths connecting two nodes). Despite their apparent simplicity, local similarity-based prediction algorithms perform generally well, even when compared to more complex approaches [36], and are also computationally efficient. A supervised learning strategy typically treats link prediction as a binary classification problem: potential new edges are labelled as positives if the classifier believes they will appear in the future, as negatives in the opposite case. These classifiers are trained using a set of existing and missing edges, described with a set of features (e.g., the heuristics used by similarity-based approaches for computing similarity). Standard learning approaches can be used in this case (such as logistic regression, Random Forest, etc.). Supervised learning has been shown to outperform unsupervised benchmarks in [25, 47, 51]. Recently, self-supervised link prediction methods have been proposed, which automatically extract latent features that reflect the relevant structure of the graph. These approaches mainly rely on graph embeddings [13, 23] or graph neural networks [10, 27, 51], and have been shown to perform very well on the link prediction task.

One of the well-known mechanisms that drive link formation is the *homophily* principle, which states that individuals tend to bond with those that are more similar to them [38]. Common-neighbors-based algorithms leverage this principle for predicting new links. However, findings from anthropology show that not all bonds are equal: some relationships are just acquaintances and do not imply a significant cognitive engagement. With respect to the meaningful relationships, each individual organises them into concentric *social circles* where the intimacy progressively decreases from the innermost to the outermost circle [26, 54]. This model of how humans organise their relationships stems from the *social brain hypothesis*, which links our social life to cognitive constraints related to the size of our neocortex [14]. This hierarchical social structure is known to impact significantly on who we trust [44], on the way information spreads in online social networks [4], and on the diversity of information that can be acquired by users [3]. This layered social structure is typically represented using an ego network graph [16, 26, 32, 37], in which the individual, referred to as *ego*, is at the center of the graph, and the edges connect her to the peers (called *alters*) with which she interacts. Figure 1 illustrates the classical ego network structure. The ego-alter tie strength is typically computed as a function of the frequency of interactions between the ego and the alter. The ego network structure is characterised by a striking regularity, both for offline and online relationships: the typical sizes of the layers are 1.5, 5, 15, 50, 150, with an approximately constant ratio of 3 between the size of consecutive layers. Please note that the second layer includes all the alters in the first layer, the third layer includes all the alters in the first and second layers, and so on.

Since social links are not all the same, simply considering the set of common neighbors (or analogous topological metrics), disregarding the importance of each social link, might not give a

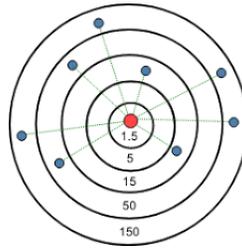


Fig. 1. Layered structure of the ego network. The red dot at the center corresponds to the ego.

precise estimation of the homophily between two individuals. Thus, in this paper, we set out to improve link prediction algorithms (both supervised and unsupervised) with information about the different importance of social relationships. To this end, we exploit well-established models in the anthropology literature that describe the relative importance of our social bonds by grouping them into different social circles. Please note that the social circles we refer to are very loosely related to the concept of communities in social networks. In fact, social circles are (local) groupings entirely from the perspective of the ego, they include only nodes with which the ego is directly connected, and the links between the alters (which are crucial for regular communities) do not play a role at all here. Note also that our goal is not to propose a new prediction algorithm, but to add social-awareness to popular prediction solutions. To the best of our knowledge, the only other work in the literature that investigates ego networks in relation to link prediction is [2]. However, in the context of that work, an ego network is a different structure from the one studied here (only the links between the alters are considered), social cognitive models are not leveraged, and the effect of using information on the different social circles for link prediction is not investigated.

While the general problem of link prediction entails recommending any of the useful links that a user may want to create with other users, in certain applications we may want to recommend links between users that share a specific common interest together. Predicting links in such contextualised communities might be much more relevant (and valuable) than predicting links between generic users. This could be the case, for example, of a third party service that wants to connect users interested in a specific topic. In this paper, we showcase one application of this concept by focusing on a set of gamers on Twitter, and we investigate how to make new link recommendations leveraging key properties of the way users allocate their cognitive capacity to social relationships. As discussed in Section 3, we believe focusing on such a target set of users, characterised by a specific type of online engagement, is particularly significant for link recommendation in online social networks. Then, we apply the same method to generic Twitter users, in order to assess the efficacy of the proposed solution also beyond the contextualised communities settings.

The key findings presented in the paper are the following:

- In the vast majority of cases, regardless of the prediction approach (unsupervised or supervised), the specific heuristic or learning algorithm, and the metric (precision, AUC, F1 score) considered, leveraging social circles information outperforms the corresponding baseline in which circles are ignored.
- The contribution of the different social circles to the achieved precision varies with the recommendation algorithm considered. In unsupervised similarity-based approaches, innermost circles help policies with little or no penalization of high degree nodes, while outermost circles boost the performance of algorithms relying on high-degree nodes penalization. The precision of supervised strategies is generally considerably increased when only the innermost circle, corresponding to the small set (composed of 3-5 people typically) of the most intimate relationships, is considered.

- Social-awareness is very effective when only little information can be used to perform recommendation, such as in the case of resource limitations (e.g., storage or computational resources). In this situation, making predictions based on the edges in the innermost circles (equivalent to considering just about a dozen neighboring nodes per ego) is generally at least as effective as using the whole network.
- Leveraging social circles information provides the same performance as using additional classifiers on nodes, which might be impractical or costly to use.
- When used for link prediction, our preliminary results show that explicit features that leverage social circle information are able to beat latent features extracted using graph embedding methods. This generally happens when the innermost circles are used, thus combining resource savings with top prediction performance.

The rest of the paper is organised as follows. We review the related work in Section 2, considering both the literature on social cognitive constraints and that on link prediction. In Section 3 we present the two datasets that we leverage in our analysis. The ego-aware unsupervised and supervised link prediction approaches that we study are presented Section 4. Their prediction performance are then evaluated in Section 5. In Section 6 we derive the time and space complexity of our ego-aware policies with respect to their baselines. Finally, in Section 7 we conclude the paper.

## 2 RELATED WORK

### 2.1 Dunbar’s model and ego networks

According to the *social brain hypothesis* from anthropology [14], the social life of primates is constrained by the size of their neocortex. Specifically, for humans, the typical group size is estimated around an average of 150 members, a limit that goes under the name of Dunbar’s number. This limit is related to the cognitive capacity that humans are able to allocate to nurturing their social relationships. The 150 *friends*, in fact, do not include acquaintances, but only people with which a coherent quality relationship is entertained. In a modern society, this entails, at the very least, exchanging birthday or Christmas cards every year. Unsurprisingly, humans are not fair in how they distribute their cognitive attentions among the 150 important persons they have around them. To the contrary, it is possible to group our social relationships in circles of increasing intimacy, as illustrated in Figure 1. Typically, each of us has at least four circles of intimacy [26, 54]: the innermost one (*support clique*) includes our close family, the *sympathy group* comprises all those people whose death tomorrow would leave us deeply affected, the *affinity group* includes colleagues, extended family, people you hang out with, and finally the *active network* is made up of all the people you meaningfully interact with at least once a year. The circles are conventionally concentric, with the inner layers contained in the outer ones. Many people also feature an additional inner layer (contained in the support clique) comprising on average 1.5 alters for which they have a very high emotional investment [15].

The social brain theory was developed for the offline world, in which relationships were nurtured via face-to-face interactions, letters, or landline phone calls. Many had postulated that this theory would not carry over to the online world, where OSN allow us to engage conveniently with a huge number of people scattered across the globe. However, Dunbar’s model stood the test of the cyberworld as well: Dunbar’s number hold for email communications [24], mobile phone calls [40], and Twitter [20]. More importantly, the same layered structure of social relationships has emerged on both Facebook and Twitter [15]. These findings are extremely important. In fact, since ego network structures are known to impact significantly on the way information spreads in OSN, and on the diversity of information that can be acquired by users [3], embedding these models of

human cognition into services for OSN may drastically improve the quality of service provided to the users.

## 2.2 Link prediction

The literature on link prediction algorithms is large. Here we only summarise the main approaches and we refer the interested readers to [35, 36] for detailed surveys on the topic. Following the taxonomy proposed by [36], link prediction approaches can be classified into similarity-based methods, algorithmic methods, preprocessing methods, and probabilistic & statistical methods (the latter are out of the scope of the paper and will not be discussed further).

Similarity-based methods make up the largest class of link prediction algorithms proposed in the literature. The rationale of this approach is that nodes are more likely to form links with nodes that are similar to them. This idea is grounded in the widespread and well-documented social phenomenon of homophily [38]. Algorithms in this class differ in how they define the similarity between nodes. We can broadly distinguish them based on whether they use local information, global information, or a hybrid combination of the two (this approach is one of the least popular, hence we will not treat it further). For a given node pair, local similarity-based solutions rely on node neighborhood-related structural information, such as the number of common neighbors [30], an inverse function of the degree of common neighbors [1, 53], or a Preferential Attachment index [30]. Local similarity-based approaches are very efficient, even on large networks, due to their easy parallelization. Their main theoretical limitation is that they are able to predict only new links between neighbors-of-neighbors. However, their prediction accuracy is typically good, and comparable to that obtained with more complex approaches. Their counterpart, global similarity-based approaches, rely on metrics computed considering the whole network topology. These metrics focus typically on the possible paths inside the network, such as the shortest paths [29], paths with different lengths by means of the Katz Index [30], random walks [33], or community membership [43].

Algorithmic methods map the link prediction problem into well-known algorithmic approaches. Within this category, classifier-based methods treat the link prediction problem as a binary classification problem [12, 25, 50, 51]. Each node pair can be characterised with a variety of attributes (including whether the link will appear or not in the future) and this labelled set can be fed to virtually any classifier (such as decision trees, SVM, k-nearest neighbors, random forest, neural networks and, recently, graph neural networks). The advantage of this approach is that it can be extended and adapted to any new attribute that one wants to test, and that it generally significantly outperforms similarity-based methods [12, 25, 47]. Also, while the heuristics of similarity-based approaches typically work well in human social networks (which feature well-studied structures), recent supervised proposals (such as graph neural networks) are much more flexible when the network under study is of a different nature, such as for protein-protein interaction networks [51]. The drawback of supervised link prediction algorithms is that they do not scale very well (as they are computationally intensive) and that they may suffer from the class imbalance problem, typical of link prediction. While classifier-based solutions are by far the most popular in the algorithmic methods category, other methods have been proposed that are based on evolutionary algorithms [7] and matrix factorization solutions [39].

Preprocessing methods are considered meta-approaches, as they are intended to be used in conjunction with other algorithms, for which they provide preprocessing/pre-filtering intended to remove some noise in the network. As an example, the clustering method discussed in [29] suggests the removal of the weakest links (those between nodes with few or zero common neighbors). The solution that we propose in this paper falls into this category. The meta-approach that we

investigate is based on considerations related to how people distribute their social capacity across their relationships, rather than on pure graph-related properties.

Orthogonally to this classification, we can also distinguish between approaches that consider the weights of the links or not. Among the former, we mention [53, 55]. As discussed in [34], though, the performance of weighted indices is often worse than their unweighted counterpart. For this reason, in this work we do not consider the weight of links besides what is needed to compute the ego network structure.

### 3 THE DATASETS

As discussed in Section 1, in this paper we investigate social circle-aware link prediction considering both a community of users with shared interests (gaming, in our case) and generic users. Before presenting the algorithms for predictions in Section 4, here we first introduce the datasets that we use as case study. Without loss of generality, some definitions required in the prediction algorithms will be discussed with reference to these datasets. However, please note that each such definition can be used as-is for any other dataset describing different communities, and the algorithm works unchanged in all such cases.

In the gaming-related dataset, we consider as shared topic of interest “indie games”. An independent video game, or indie game, is a video game that is often created without the financial support of a publisher, although some games funded by a publisher are still considered “indie”. Indie games often focus on innovation and rely on digital distribution. Growing a community of interested users is thus a critical success factor for them. The goal of our link prediction task is to recommend users interested in indie games to other users interested in indie games. In order to collect an initial community of users interested in indie games, we first had to identify some relevant indie games. To this purpose, we referred to Steam, a digital game distribution platform (<https://store.steampowered.com/>), and SteamSpy (<https://steamspy.com/>), which provides statistics about the games on Steam. Among the indie games listed as the most popular during October 2017, we were able to identify the Twitter accounts of 133 of them. We then downloaded the timelines of these games using the Twitter Search API. From the downloaded timelines, the 400 most frequently used hashtags have been extracted and used to monitor Twitter using the Twitter Streaming API. This allowed us to identify a set of 8,932 users engaging in game-related conversations. We labelled these users as *gamers*. We have then collected (June 2018) their timeline (most recent 3200 tweets) using the Twitter Search API.

In the generic user dataset, we focus on a group of users whose interests are not associated with a specific domain. To this aim, we used a dataset presented in [5], containing the timelines (most recent 3200 tweets) of 1,930,802 Twitter users obtained by means of a snowball sampling starting from US President Barack Obama in November 2012. Please refer to [5] for further details on the collection process.

We filter the 8,932 gamers and 1,930,802 generic users and extract their ego network (including their social circles) using the same methodology as in [8]. Please note that, to this aim, we simply need a weighted social graph, whose edge weights correspond to the contact frequency between the corresponding nodes. In order to extract reliable information, we filter users according to the same policy used in [8]. We only consider users whose Twitter activity is regular (i.e., they post, on average, at least one tweet every 3 days for at least 50% of the total number of months of their activity) and stationary (i.e., they are not in their initial stage of engagement with the Twitter platform, which typically features a transient spike of activity). The strength of ego-alter relationships is inferred from the frequency of direct tweets (mentions, reply, or retweets) between the ego and the alters. The optimal number of social circles per users is obtained using the Mean Shift clustering algorithms to group such frequencies. Approximately one third of the gamers

Table 1. Gaming-related dataset summary

# of gamer nodes $ V_e $	2,995
# of domain-specific nodes $ V_d $	70,859
# of all nodes $ V  =  V_e \cup V_d \cup V_n $	470,485
# of gamer edges $ E_e $	2,614
# of domain-specific edges $ E_d $	154,581
# of edges $ E $	1,004,011

Table 2. Generic users dataset summary

# of ego nodes $ V_e $	3000
# of all nodes $ V  =  V_e \cup V_n $	278,510
# of ego edges $ E_e $	7,158
# of edges $ E $	567,739

passed all the filters described above (relying only on their last 3200 tweets), and this left us with 3,061 gamers with reliable ego network information. On the other hand, approximately 7% of the generic users passed all the filters described, amounting to 148,105 generic users with reliable ego network information. When we focus on an ego  $i$ , who is a gamer in gaming-related dataset or a generic user in generic user dataset, its alters can now be associated with the social circle of  $i$  to which they belong. In the gaming-related dataset, the alters of ego  $i$  can be generic users, other gamers, or even games (e.g., one of the initial 133 games). This implies that some nodes are *domain-specific* (in our case, gaming-related), while others are generic. On the other hand, in generic user dataset, all alters of ego  $i$  are also generic users. Generic nodes make prediction more challenging with respect to the case of homogeneous nodes (as, e.g., in [30]).

In order to enrich the set of domain-specific nodes in the gaming-related dataset, we also downloaded all followers of the initial 133 games, thus bringing in additional 25,014 nodes, for a total of 31,091 domain-specific nodes (which include games and gamers without ego network). Thus, in the end, we have three classes of nodes in gaming-related dataset: gamer users with ego network ( $V_e$ ), regular (unlabelled) users without ego network ( $V_n$ ), and domain-specific users that may have or may have not an ego network ( $V_d$ , including games and gamers for which ego network information is not available or not reliable because they didn't satisfy the filters discussed previously). In generic user dataset, we have two classes of nodes: generic users with ego network ( $V_e$ ) and generic users without ego networks ( $V_n$ ) that are basically alters. We further filter these graphs as suggested in [30] by considering only the nodes that are connected to the giant component. In order to have a generic user dataset with a similar size of gaming-related dataset and to reduce its computational intensity, we sampled 3,000 egos from the giant component of the generic users network by snowball sampling [21].

The final datasets are summarised in Table 1 and Table 2. The imbalance ratio, i.e., the ratio between the negatives (potential links that do not exist in practice) and the positives (links actually existing) in the class of both gamers and generic users is around  $10^4:1$ , which is quite high. This is a well-known problem in link recommendation for online social networks (remember from Section 1 that links are sparse). This implies that special care should be taken in the evaluation, as explained in detail in Section 5.2.

#### 4 LINK PREDICTION BASED ON SOCIAL CIRCLES

We denote with  $\mathcal{G} = (V, E)$  the graph modelling the relationships in our dataset. As discussed in Section 3, in the gaming-related dataset, vertices can belong to one of three classes: *gamers* (denoted with  $V_e$ ), *domain-specific users* (denoted with  $V_d$ ), and *unlabelled users* (i.e., for which we know nothing regarding their involvement in gaming related activities, denoted with  $V_n$ ). In the generic user dataset, vertices can belong to one of two classes: users with ego network (denoted with  $V_e$ ) and users without ego network (denoted with  $V_n$ ). Since our goal is to evaluate the contribution of social circles to link prediction, we will focus on the subset of nodes for which social circles could be computed. Thus, predictions will be made for the set of possible edges between nodes in  $V_e$ .

#### 4.1 Similarity-based unsupervised learning with circle-awareness

In order to showcase the effect of social circles for unsupervised link prediction, we focus on the simple, yet effective class of similarity-based approaches. Like for all similarity-based link prediction algorithms, the goal of the proposed algorithm is to associate each non-existing link between two users  $i, j$  in  $V_e$  with a score proportional to the likelihood that the link will actually be formed in the future. Hence, the goal of the link predictor is to associate with each edge in  $V_e \times V_e - E_e$  a probability/confidence that the link will indeed appear in the future. Qualitatively, this is equivalent to *suggesting users to other users* for possible interactions, friendships, etc.

In order to study the effect of the social circles on link prediction, we propose to *slice* the social graph based on the membership to a specific circle (recall that our concept of social circles is different from that of communities). Let us focus on a user  $i \in V_e$ . We denote with  $C_1^{(i)}, C_2^{(i)}, \dots$ , the sets comprising the neighbors of  $i$  that belong to  $i$ 's first, second,  $\dots$ , social circle, respectively. Thus, the graph  $\mathcal{G}$  sliced according, e.g., to circle  $C_3$ , is the graph including only links between egos and their alters up to layer 3 (remember that  $C_3$  includes  $C_1$  and  $C_2$ ). More formally, it can be defined as  $\mathcal{G}_{C_3} = (V, E_{C_3})$ , where  $E_{C_3} = \{e_{ij} \in E : \exists i \in V_e, \exists j \in V \text{ s.t. } j \in C_3^{(i)}\}$ .

Social-circles slicing can be extended to include category-based slicing as well for domain-specific datasets. For example, from  $\mathcal{G}_{C_3}$  we can retain only domain-specific nodes and edges (i.e., nodes in  $V_e \cup V_d$  and edges in  $E_e \cup E_d$ ). We generalise the notation related to the slicing preprocessing by denoting with  $\omega$  the specific slicing considered, and with  $\mathcal{G}_\omega$  the resulting graph. In Section 5, we will discuss the performance of link prediction both with social-based slicing and with social-based plus category-based slicing. We anticipate here that when the social circles information is not used for the prediction, leveraging the category information provides a significant advantage. Vice versa, social-circles awareness makes information on the category less relevant (and this is extremely important, since the extraction of categories typically requires an additional, domain-specific, classifier or manual labelling).

In order to perform link predictions, for each user-user pair  $i, j$  for which a link does not exist in  $E_e$ , we do the following:

**DEFINITION 1 (UNSUPERVISED CIRCLES-AWARE LINK PREDICTION).** *For a fixed  $K$ , the social circles-aware link prediction algorithm suggests to a user  $i$  the top  $K$  users  $j$  in the following ranked list:*

$$L_\omega = \{j \in V_e \wedge (i, j) \notin E_e : \text{sim}_\omega(i, j)\}, \quad (1)$$

where  $\omega$  denotes the selected slicing and  $\text{sim}_\omega(i, j)$  is the similarity computed on  $\mathcal{G}_\omega$ .

The definition of a new similarity function is out of the scope of the paper. What we want to do here is to evaluate popular similarity functions available in the literature when they are enriched with knowledge about the ego network circles. In order to isolate the effect of social-circles awareness, we need approaches that are simple (so it is easy to gauge the role of social circles) yet effective. For these reasons, our choice fell on the strategies described below, often used in the related literature. In summarising them, we denote with  $\Gamma_\omega(i)$  the neighborhood of node  $i$  in  $\mathcal{G}_\omega$ .

- *Common neighbors (CN)* [30]: the similarity is given by the number of common neighbors between users  $i$  and  $j$  in  $\mathcal{G}_\omega$ :

$$\text{sim}_\omega(i, j) = |\Gamma_\omega(i) \cap \Gamma_\omega(j)|. \quad (2)$$

- *Jaccard's Coefficient (JC)* [30]: the similarity is computed as the Jaccard similarity of the set of common neighbors of users  $i$  and  $j$  in  $\mathcal{G}_\omega$ :

$$\text{sim}_\omega(i, j) = \frac{|\Gamma_\omega(i) \cap \Gamma_\omega(j)|}{|\Gamma_\omega(i) \cup \Gamma_\omega(j)|}. \quad (3)$$

- *Adamic-Adar (AA)* [1]: the Adamic-Adar similarity reduces the importance of common neighbours having high degree:

$$sim_{\omega}(i, j) = \sum_{z \in \Gamma_{\omega}(i) \cap \Gamma_{\omega}(j)} \frac{1}{\log(|\Gamma(z)|)}. \quad (4)$$

- *Resource Allocation (RA)* [53]: the RA score is similar to the Adamic-Adar one, but it penalises even more the common neighbors with high degree:

$$sim_{\omega}(i, j) = \sum_{z \in \Gamma_{\omega}(i) \cap \Gamma_{\omega}(j)} \frac{1}{|\Gamma(z)|}. \quad (5)$$

Among these four link prediction algorithms, RA is the one that consistently performs better in the related literature [34, 52, 55]. In Sections 5.4 and 5.5 we will investigate if this is still the case when the policies can leverage the knowledge of the social circles.

## 4.2 Supervised learning with social-circle awareness

Each topological metric discussed in the previous section captures a single possible mechanism yielding to the formation of new links in the network. [31] discuss how this unsupervised approach can take advantage of adding supervised learning on top of it. They show that simple ranking of heuristics is outperformed by supervised classifiers, because the latter have the capability of identifying multiple differentiating boundaries in the similarity score domain, even when just a single heuristic is used as feature.

Thus, in this section, we cast our social-aware link prediction problem into a supervised learning problem. This entails computing a vector of features for each user pair  $i, j$ . Each of the metrics in Section 4.1 becomes a feature that describes the user pair. Each pair is also labelled to mark whether the link exists or not. In order to test the different performance of different supervised approaches, in this work we consider the following learning algorithms (in parentheses we report the link prediction papers in which they have been previously used): logistic regression [42], Random Forest [31], decision trees, naïve Bayes, and SVM [25], the latter with both linear and polynomial kernels. We use their R implementations (`glm`, `randomForest`, `rpart`, `kLaR`, `kernlab`, respectively) together with the `caret` package for training and test. Parameter optimization is applied on the training set with 10-fold cross validation.

## 5 EVALUATION

In this section we carry out the performance evaluation of the link prediction approach proposed in Section 4 using the dataset described in Section 3.

### 5.1 Training and test data

There are two main approaches in the literature to obtaining a test dataset for the evaluation. The first one is based on randomly removing some existing links from the graph (then the capabilities of the proposed algorithms to predict these removed links are measured). This approach has been often used in the related literature [34, 50, 51, 53, 55]. The second approach is based on obtaining two temporal snapshots of a graph (longitudinal dataset). In this case, the performance analysis of the proposed link prediction algorithms measures how good each algorithm is in predicting the links that were not present in the first temporal snapshot but appeared in the second one. This approach is used in [25, 30, 46] and it is recommended as a general best practice in [49]. In this work, we will rely on the second approach, as it able to capture the real process behind link formation.

For obtaining longitudinal data, we downloaded (December 2019) the Twitter timeline of the gamers in  $E_e$  one year and a half after the initial download, and timeline of the generic users in  $E_e$  eight years after its initial download. We then extracted the links between gamers in the gaming-related dataset and between generic users in the generic users dataset that have appeared in the meanwhile. These new 843 links between gamers and new 1216 links between generic users constitute the set  $E_{new}$  of links to be predicted, while  $E_{old}$  contains the links existing in the first temporal snapshot.

While the above considerations are sufficient for unsupervised link prediction, with supervised learning we also need to assign the negatives (i.e., missing links) to the train and test sets. To this aim, we split them 90%-10% between train and test. Clearly, the set of edges selected for test set may offer only a partial view of the performance (i.e, they may be easier or harder than average to predict). For this reason, we perform a  $k$ -fold cross-validation [35], with  $k = 10$ , each time selecting a new 10% of the negative links for the test set. How to aggregate the performance metrics obtained from the  $k$  folds is a problem discussed in [17]. In general, it can be done in two ways. With microaveraging, true positives, false positives, true negatives, and false negatives from the  $k$  folds are pooled into one contingency table, then the metrics are computed. With macroaveraging, metrics are computed per fold, then averaged. Important performance metrics (such as the F1 score) are unbiased under microaveraging [17], hence we opted for this approach.

Using the 90% of negatives for training entails training supervised schemes on millions of negatives and only thousands of positives (Table 1, Table 2). While the impact of class imbalance is typically limited for similarity-based unsupervised approaches like the ones discussed in Section 4.1, it can create serious problems for the scalability and reliability of supervised learning. In order to mitigate the problem, the common approach in the related literature [25, 46, 51] is to undersample the negative class. Despite its widespread use, this technique is not without drawbacks [49], and it can lead to overly optimistic results. In order to provide reliable measurements, we will show both the results obtained with undersampling and the results obtained with the complete negative class. In the latter case, some learning strategies may not converge in a reasonable time (due to the scalability problems mentioned above), and we will exclude them from the evaluation.

## 5.2 Evaluation metrics

Similarly to the related literature [30, 35, 55], unsupervised link prediction algorithms are evaluated using a top-K analysis, i.e., we compare their performance in predicting  $K$  new links. This allows for a fair comparison among the different approaches, as it avoids fixing a similarity threshold for approaches in which the same threshold may have a different meaning. In order to span a reasonable  $K$  range around the number of positives in the test set (which, as we discussed in Section 5.1, are 843 and 1216, for the gaming and generic users datasets respectively), we consider  $K \in \{100, 843, 1000\}$  for the gaming-related dataset and  $K \in \{100, 1216, 1500\}$  for the generic users dataset. For a given slice  $\omega$  and for a fixed  $K$ , each unsupervised link predictor described in Section 4.1 outputs a ranked list  $L_\omega^K \subset L_\omega$  of  $K$  pairs in  $V_e \times V_e - E_{old}$ , which are the newly predicted links, ranked by confidence in prediction (in our case, the similarity value). Then we can compute classic metrics such as:

- $TP = |L_\omega^K \cap E_{new}|$
- $FP = |L_\omega^K - E_{new}|$
- $FN = |E_{new} - L_\omega^K|$
- $TN = |V_e \times V_e - E_{old}| - |L_\omega^K \cup E_{new}|$ .

Supervised approaches automatically label all the elements in the test set (without the need to set the  $K$  or a similarity threshold), hence the number of positives predicted is part of the workings of the supervised approach. In this case, thus, we do not pick the top  $K$  potential edges, but we

directly rely on the labelling of the trained predictor for the test edges. It is then straightforward to compute the above metrics (for example,  $TP$ s are those edges in  $E_{new}$  that are marked as  $new$  by the supervised prediction algorithm).

Traditionally, binary classifiers are evaluated using the ROC curve, which plots the True Positive Rate ( $\frac{TP}{TP+FN}$ ) against the False Positive Rate ( $\frac{FP}{FP+TN}$ ). However, the scenario we are considering clearly suffers from the class imbalance problem, since the actual edges between nodes are significantly fewer than all the possible edges that there could exist between them. In this situation, the ROC curve, while being insensitive to the class imbalance<sup>3</sup>, tends to provide overoptimistic results [18]. The related literature on evaluating systems with class imbalance suggests using metrics like precision ( $= \frac{TP}{TP+FP}$ ), recall (aka TPR) and  $F_1$  score. As already argued by [46], in the context of link prediction, precision is more important than the other metrics (such as recall), because if the precision is high, one can live with some false negatives. Therefore, the focus of this evaluation will be mostly on precision. For the completeness of results, we also provide, when relevant, the F1 score and the AUC of the precision-recall curve. Note also that we are interested in comparing the performance obtained using the baseline approach (no social circles information, all edges are considered) and the other strategies in which only edges belonging to the ego network circles are considered.

Precision and recall are evaluated based on metrics (TP, FP, TN, FN) which corresponds to specific realizations of a target phenomenon. Thus, the question arises of how confident we can be about the results obtained on a particular collection that is the result of random sampling (as in the case of cross-validation or subsampling of negatives for supervised prediction). [22] proposes a Bayesian approach to the problem. The basic idea is to model TP, FP, FN, TN as if they were sampled from a multinomial distribution, with unknown success probability per category  $\pi_i$ . Based on the properties of the multinomial distribution, we know that individual counts are binomial (i.e., equivalent to the number of successes in a sequence of  $n$  independent experiments, each with success probability  $p$ ). Specifically, it can be proved that TP are binomial with parameters  $TP+FP$  (corresponding to the number of trials) and success probability  $p$  (corresponding to the precision). Using Bayes' theorem and well-known priors associated with binomials, it can be derived that the posterior of the precision  $p$  is Beta distributed:

$$p|\mathcal{D} \text{ (aka the posterior distr.)} \sim \text{Beta}(TP + \lambda, FP + \lambda). \quad (6)$$

A similar formula is obtained for the recall  $r$ . In these formulas,  $\lambda$  is a parameter whose value depends on the prior considered:  $\lambda = \frac{1}{2}$  corresponds to Jeffrey's non-informative prior (the one we use here),  $\lambda = 1$  to the uniform prior. Credible intervals<sup>4</sup> for  $p$  and  $r$  can then be obtained as the intervals containing 95% of the Beta distribution. The formula for the credible interval of the F1 score is slightly more complicated, so we do not report it here. Microaveraging and credible intervals are also used in [47].

### 5.3 Ego network analysis

As anticipated in Section 3, we extract the ego network structure of the gamers and generic users in  $V_e$  using the methodology described in [8]. The resulting distribution of the optimal circle number of gamer egos and generic user egos can be seen in Figure 2. The mean value is 4.94 for gamers and 4.73 for generic users, and the median and mode values are equal to 5 for both datasets. The plot

<sup>3</sup>Differently from the accuracy, which should not be used when imbalance is an issue. In fact, given the accuracy formula  $\frac{TP+TN}{P+N}$ , it is easy to see that correct classification in the most numerous class will hide anything about the sparse class.

<sup>4</sup>Using confidence intervals instead of credible intervals would not be the best choice in this case, because precision, recall, and F1 score belong to  $[0, 1]$ , hence, especially for values close to either 0 or 1, they depart significantly from the normality assumption behind confidence intervals.

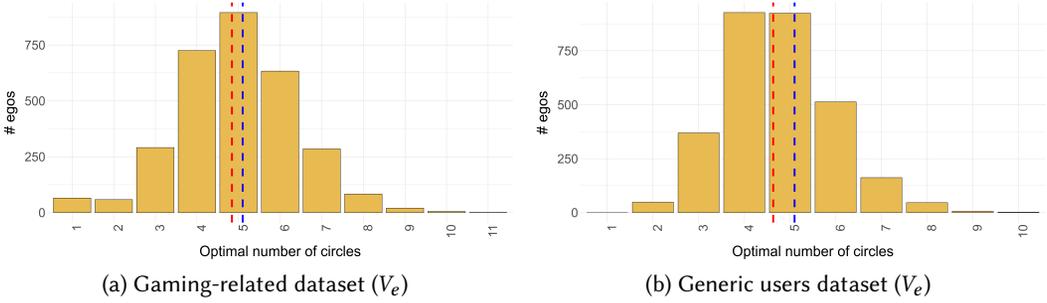


Fig. 2. Distribution of the optimal number of social circles for the gamer egos and generic user egos. Red and blue lines represent the mean and median values, respectively.

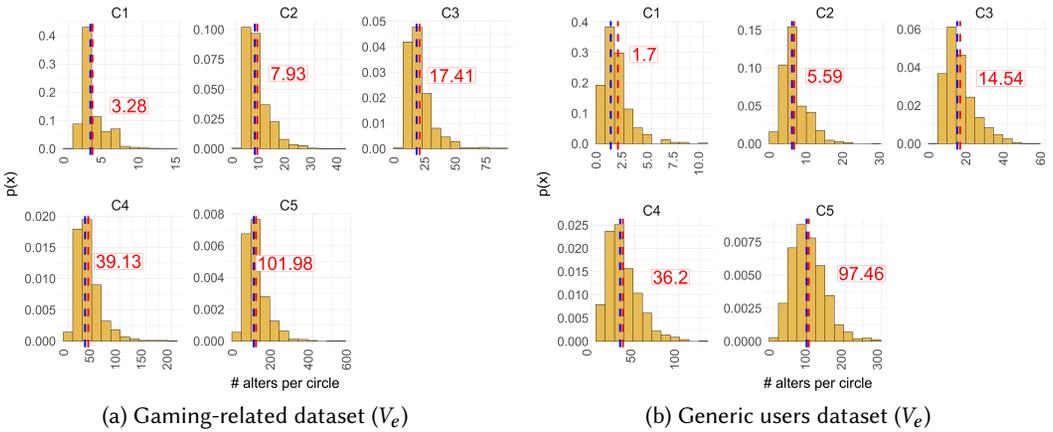


Fig. 3. Average alter distribution per each circle of gamer egos and generic user egos. For the sake of clarify, egos that have optimal circle number equal to five have been used. Red and blue lines represent the mean and median values, respectively, of the distribution.

shows that the number of social circles of both types of egos is compatible with the findings of previous studies which establish that human ego network can be layered into 5 layers/circles. For further understanding of the ego network structures, we examine the alter distribution through the layers. In previous studies [15], it has been shown that in OSN the external layers are slightly smaller than the reference model discussed in Section 2.1. In Figure 3, we show the distribution of the alter number for the egos with optimal circle number equal to 5 (the case generally studied in the related literature). These results, both for gamers and generic users, are compatible with the typical findings from OSNs analysis in the related literature.

#### 5.4 Evaluation of unsupervised link prediction

As shown in Figure 2, different nodes can have a different number of social circles, but typically the majority has approximately five circles. For this reason, in our circle-based slicing we will consider circles from C1 to C5, then we group together into the ACTIVE circle all the circles beyond C5 (for those nodes that have more than five circles). Note that, since social layers are concentric and the

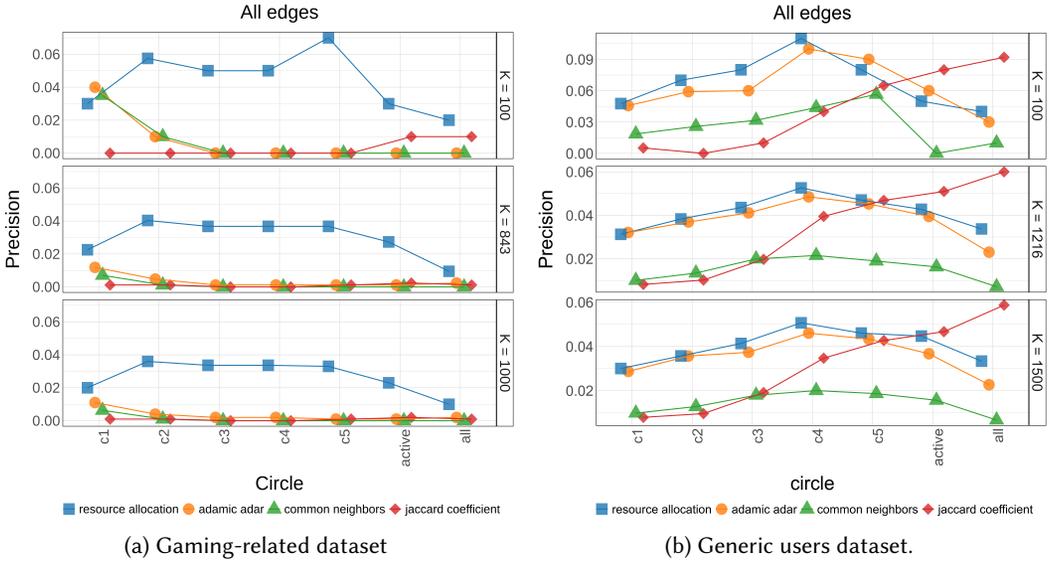


Fig. 4. Precision in unsupervised settings - ALLEdges.

$i$ -th also include the  $(i - 1)$ -th up to the first one, for egos with less than five circles the predictions in, e.g. C5, are simply based on the last non-empty social layer (which includes all the innermost ones). Recall that the alters in the ACTIVE circle are all those with a *significant* relationship with the ego, i.e. they interact with the ego at a frequency of at least one contact per year [54]. We denote with ALL the situation in which all relationships are considered, which implies considering both the nodes in the ACTIVE circle (significant bonds) and those outside (acquaintances). This corresponds to the baseline from the related literature: all relationships are treated as equal, without factoring in the role of social circles. As anticipated in Section 4, we will consider two scenarios. In the first one, denoted as ALLEDGES, we consider the full graph  $\mathcal{G}$ , slicing it based on the social circles both in the gaming-related dataset and the generic dataset. In the second one (DOMAINSPECIFICEDGES), we perform both social-based and category-based slicing, by only retaining edges between nodes that are domain-specific (i.e., gamers and games, in our case) with the gaming-related dataset. Note that credible intervals are not provided here, because in the unsupervised case there is no sampling of negatives and the links to be predicted are, deterministically, those that have actually appeared between the first and second download.

We start our discussion from the ALLEDGES setting within the gaming-related dataset. The precision<sup>5</sup> for varying  $K$  (number of new links recommended) is shown in Figure 4a. The maximum precision across all policies decreases when we increase  $K$ : the more the links recommended, the more the mistakes (as it is generally the case in the related literature). RA clearly outperforms the others in the gaming-related dataset, and this is consistent with previous findings in the related literature [34]. RA seems to benefit significantly from circle-awareness. Specifically, ignoring the outer circle gives a clear advantage to RA, which at least doubles its precision for all  $K$  values. Thus, the mechanism whereby high-degree penalization and social-awareness work hand in hand is very effective from a link prediction standpoint. It is also interesting to note that this advantage seems

<sup>5</sup>Note that the generally low precision values are due to the high class imbalance in the dataset. The benchmark random predictor (the always positive one), in the same settings, would achieve an even worse performance, approximately equal to  $2 * 10^{-4}$ .

Table 3. ALLEDGES scenario with gaming-related dataset: AUC ( $\times 10$ ) for the precision-recall curve of unsupervised link prediction. We highlight in **bold** the social layer in which each policy performs at its best. We underline the best two AUC overall.

	AUC ( $\times 10$ )						
	C1	C2	C3	C4	C5	ACTIVE	ALL
RA	0.01179	0.03162	0.04699	<u><b>0.05364</b></u>	<u>0.05162</u>	0.03641	0.01916
AA	<b>0.02148</b>	0.00749	0.01023	0.01169	0.01270	0.01372	0.01122
CN	<b>0.01914</b>	0.00540	0.00764	0.00902	0.01000	0.01144	0.00924
JC	0.00332	0.00448	0.00648	0.00813	0.00927	<b>0.01007</b>	0.00969

lost when only the innermost circle is considered. However, leveraging only the innermost circle C1 means using very little information (only a few links per ego, as can be seen in Figure 3). This can be a crucial advantage in case of resource limitations (either in terms of storage space, or in terms of computational time, since computing the similarity is faster on smaller neighborhoods, as we discuss in Section 6). In such cases, we can have a good prediction (still better than that at ALL) by using only C1. Thus, the predictive power of using only the most intimate relationships unexpectedly outperforms baselines relying on *all* relationships. The drastically different performance of AA with respect to RA in Figure 4a tells us that the way we penalise high-degree nodes impacts on the effect that social circle information has on link prediction effectiveness. However, it is interesting to note that while in general RA is better than AA, this is not the case when resource constraints are considered. When only the most intimate links are kept in the ego networks, AA performs at its best and its precision is not only equivalent to that of RA, but also comparable to the precision of RA in the baseline (ALL). In this case, then, strong degree penalization seems to be less important than strong intimacy. The Common Neighbors policy performs quite similarly to AA, even if, in CN, all neighbors contribute the same to the similarity score, regardless of their degree. This is evidence that here AA is working in a regime in which all nodes weight approximately the same (i.e., its logarithmic degree penalization is not enough). Given the similarity with AA, the same considerations we made for AA hold for CN. Finally, we highlight the poor precision of the Jaccard-based strategy, for all  $K$ s and for all circles in the gaming-related dataset. Jaccard similarity gives more weight to neighborhoods that are very similar to each other. This means that not only the overlapping part is considered, but also the number of nodes that are not in common (union of neighbors). As testified by the plots, this restriction does not give an advantage for link prediction in the gaming-related dataset. With respect to social-circles awareness, Jaccard seems to suffer from using only the intimate relationships, while using the ACTIVE layer yields the same precision as the the baseline ALL while saving computational resources.

We report in Table 3 the Area Under the prediction-recall Curve. Recall that the AUC is not dependent on  $K$ , since it is obtained by exploring the whole range of similarities for making recommendations. When considering the relation between precision and recall through the AUC, we observe that the best overall results (values underlined in the table) are achieved by RA that leverages C4 and C5. Looking at the layers that, for each policy, provide the best AUC, we note that the ALL case (which is the baseline for similarity-based policies from the related literature) is never the best performing. From the circle-awareness standpoint, there seem to exist two classes of policies, those that benefit most from using the outermost circles and those that benefit from the innermost ones.

We now focus on the ALLEDGES setting within the generic users dataset (Figure 4b). Again we observe that the maximum precision across all policies decreases when we increase  $K$ . In this

case, though, the baseline results (i.e., those without social-awareness) are different: while, for the gamers network, RA was already the best predictor in the ALL case, here we observe a generalised advantage of the JC approach for all  $K$  values. Given this different starting point, let us study how the prediction policies react to circle-awareness. RA preserves the characteristics we observed for the gamers, and still very much benefits from social-awareness: all the circle-based scenarios have better or comparable precision than the baseline ALL. Note that, as mentioned in the gaming-related dataset results, even if the prediction precision is similar – as it is the case in Figure 4b for C1 and ALL with large  $K$  – using the most intimate links in C1 gives us the advantage of less computation time and memory usage against traditional methods (see Section 6 for a thorough discussion). Turning our attention from RA to AA, Figure 4b shows that AA performs significantly better compared to what we observed for the gaming-related dataset: here, AA substantially mimics the performance of RA, both in terms of achieved precision and in terms of social circles in which it performs best (specifically, around C4). Interestingly, AA was mimicking instead the performance of CN in the gaming-related dataset, with peak performance in C1 and a generally poor precision both in the other social circles and in the baseline ALL. The fact that this difference is due to the different structural properties of the generic users graph with respect to the gamers one follows directly from the definitions of RA, AA, and CN in Section 4.1. Indeed, all the three policies are based on the principle that the more the common neighbors, the better for link prediction. Then, in CN all common neighbors weight the same, while in AA and RA the common neighbors with high degree weight less (much less in RA than in AA, since in RA the penalization is linear while in AA it is logarithmic) than those with small degree. The fact that the performance of AA is approaching that of RA implies that the degree of the common neighbors is not very high (hence the logarithmic penalization of AA and the linear one of RA are in a regime in which they yield similar scores). Taking into account the degree of the common neighbors still provides an advantage, though, as highlighted by the worse performance of CN with respect to AA and RA. As a final remark, note that while JC is performing extremely well with the generic users dataset in the baseline ALL, its performance rapidly deteriorates as we incorporate circle information. This tendency was already present in the gamers datasets (Figure 4a) but was somewhat masked by the generally poor precision achieved by JC regardless of the circle/baseline considered. It is interesting to speculate on why circle-awareness does not help JC in general. To this aim, let us consider the definition of the Jaccard similarity in Section 4.1. JC captures the fraction of common neighbors out of all neighbors for a pair of users  $i, j$ . Thus, its very nature requires “useless” neighbors (i.e. those that are not in common) to be present, in order to differentiate between high and low neighborhood overlap. However, the slicing based on social-circles effectively remove “useless” nodes, hence impairing the discriminating capabilities of the JC score.

We report the Area Under the prediction-recall Curve for the generic uses dataset in Table 4. Again, recall that the AUC captures the trade-off between precision and recall and that it is not dependent on the specific  $K$  value considered, as it spans the whole  $K$  range. The best overall precision-recall performance is achieved by RA when it leverages social information in C5. The second best AUC is provided by JC in the baseline case (no circle-awareness). AA and CN provides the highest AUC when leveraging C5 information. Overall, from the results in Table 4, we can thus conclude that, even though the advantage of circle awareness is less evident in the generic users dataset, it still outperforms non-circle-related approaches.

We now consider the `DOMAINSPECIFICEGES` scenario, which is only relevant to the gaming-related dataset. Recall that, in this case, for any social-based slicing, we also remove all nodes (and associated edges) that are not domain-specific. All the nodes left after this additional, category-based slicing are either gamers or games, i.e., they belong to the specific community for which we are making link recommendations. Making a link prediction is now much easier, because the network

Table 4. ALLEDGES scenario with generic users dataset: AUC ( $\times 10$ ) for the precision-recall curve of unsupervised link prediction. We highlight in **bold** the social layer in which each policy performs at its best. We underline the best two AUC overall.

	AUC ( $\times 10$ )						
	C1	C2	C3	C4	C5	ACTIVE	ALL
RA	0.03028	0.05755	0.10289	0.14233	<b><u>0.15142</u></b>	0.13941	0.09313
AA	0.02730	0.04987	0.07930	0.11554	<b><u>0.12075</u></b>	0.10057	0.06273
CN	0.01106	0.01914	0.03093	0.04678	<b><u>0.04714</u></b>	0.04171	0.03088
JC	0.00927	0.01737	0.03573	0.07824	0.10735	0.12498	<b><u>0.14390</u></b>

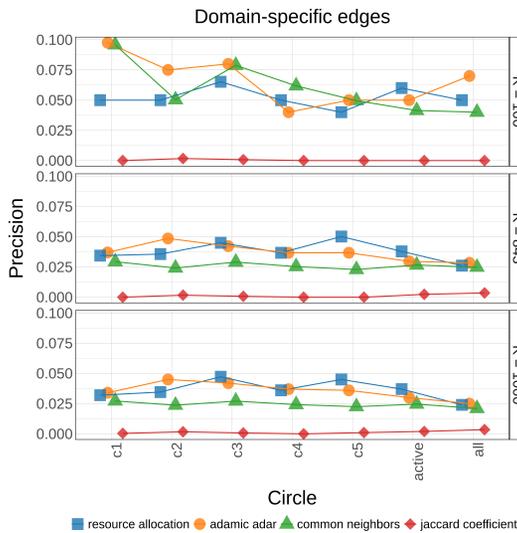


Fig. 5. Precision in unsupervised settings - DOMAINEDGES (only relevant to the gaming-related dataset) .

has been pruned by nodes potentially irrelevant or misleading. Indeed, as shown in Figure 5, the precision of AA and CN significantly improves with respect to the ALLEDGES scenario in Figure 4a. This improvement is also present when no social-circle information is used (ALL circle in Figure 4a). This implies that when social-circle awareness is not used for prediction, leveraging the category information provides a significant advantage. Vice versa, social-circles awareness seem to make information on the category unnecessary. Since the extraction of user categories typically requires a domain-specific classifier or manual labelling, being able to skip this phase without affecting the link prediction quality would be very important. The precision of the Jaccard-based approach remains very low also when considering domain-specific edges only. We conclude this part on the DOMAINSPECIFICEDGES scenario by analysing the AUC of the precision-recall curve, reported in Table 5. Table 5 shows that, also in this case, the best performance is achieved by RA in the outermost circles (specifically, ACTIVE and C5). We observe again the duality of behaviours with respect to the gain from social-circles awareness: RA and AA achieve their best AUC in ACTIVE, CN and JC in C1. The baseline ALL is never optimal (neither from an absolute standpoint nor for specific policies). This further confirms the advantage of using social circle awareness in general.

Table 5. DOMAINEDGES scenario with gaming-related dataset: AUC ( $\times 10$ ) for the precision-recall curve of unsupervised link prediction. We highlight in **bold** the social layer in which each policy performs at its best. We underline the best two AUC overall.

	AUC ( $\times 10$ )						
	C1	C2	C3	C4	C5	ACTIVE	ALL
RA	0.3331	0.5457	0.7275	0.7361	<u>0.8175</u>	<b>0.8807</b>	0.6605
AA	0.6564	0.5754	0.7638	0.6762	0.6834	<b>0.7870</b>	0.5745
CN	<b>0.6523</b>	0.5643	0.5381	0.4289	0.4131	0.5562	0.4028
JC	<b>0.1296</b>	0.1270	0.1353	0.1205	0.1117	0.1053	0.0950

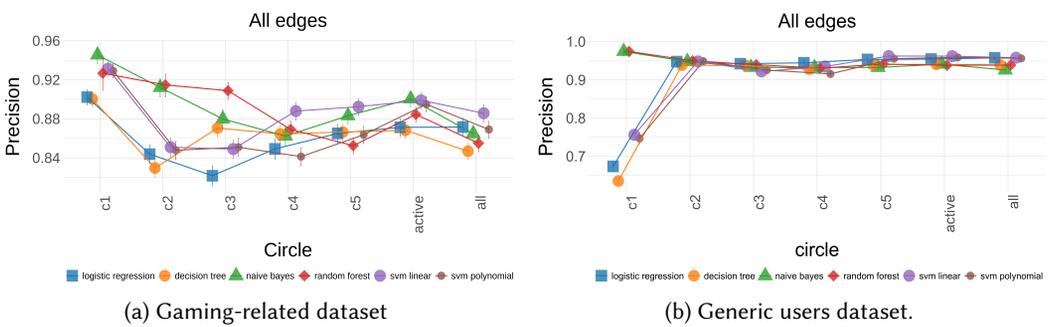


Fig. 6. Precision (with credible intervals) - SUPERVISED ALL EDGES on the undersampled datasets.

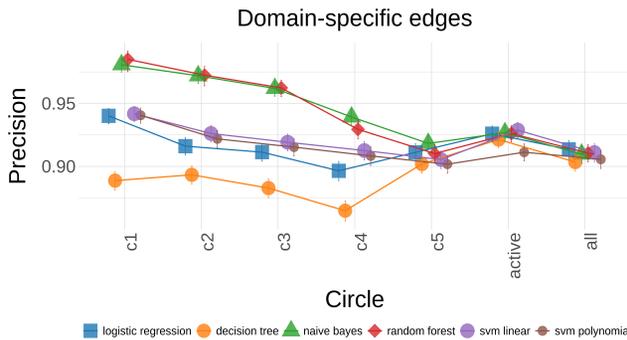


Fig. 7. (Gaming-related dataset) Precision (with credible intervals) - SUPERVISED DOMAIN EDGES on the undersampled dataset.

### 5.5 Evaluation of supervised link prediction

In this section we assess the advantages brought about by including a supervised classifier in the link prediction approach. As discussed in Section 4.2, our classifier uses as features the similarity-based metrics introduced in Section 4.1. The supervised approaches that we test are logistic regression, decision trees, naïve Bayes, Random Forest, SVM. Recall (from Section 5.1) that with supervised learning we investigate the performance both on the original dataset and on the undersampled one. Also, with supervised learning we do not carry out a top-K analysis but we allow the algorithms to freely classify all the edges in the test set.

Figures 6a, 6b, and 7 show the results, for the ALLEdGES and the DOMAINEdGES case respectively, when the negatives are undersampled to make the positive and negative classes balanced (the negatives are then split 90%-10% between train and test). Undersampling, as discussed in Section 5.1, is a common approach to deal with class imbalance in supervised settings, though not without drawbacks. For this reason, later in the section we also investigate the link prediction performance when no undersampling is carried out. Figure 6a shows the precision of the supervised link prediction approaches in the ALLEdGES scenario of gaming-related dataset. We can identify two important results. First, all approaches that rely on C1 information provide better precision than with any other layer. In particular, using all edges (case ALL) is never better than using C1 alone. Second, all circle-based approaches that leverage information on the outermost social layers perform at least as well as the baseline (ALL). When considering the generic users dataset (Figure 6b), we observe that the peak performance in C1 is confirmed for Random Forest and naïve Bayes. However, for the other learning strategies, the precision drastically drops in C1. For generic users, we can conclude that all the social circles but the first one provide approximately the same precision as the baseline ALL. Even with no performance improvement with respect to the ALL case, circle-awareness allows us to save computational resources, and this is an important advantage. Finally, please note that, for both datasets, the precision achieved with supervised strategies is much higher than that obtained with unsupervised ones (consistently with the results in the related literature [46]).

In the DOMAINSPECIFICEdGES scenario for the gaming-related dataset (Figure 7), again naïve Bayes and Random Forest achieve the highest precision when leveraging only C1 relationships. Differently from the previous case, the precision of decision trees in C1 is lower than with other layers. The remaining learning strategies still gain from using C1 information, but the margin is smaller in this case. Differently from the unsupervised case, supervised learning seems able to overcome the difficulties of predicting new links in the ALLEdGES scenario, hence the relative advantage of filtering out links not domain-specific is partially lost.

As already mentioned, the precision achieved with supervised learning is generally much higher than that observed with unsupervised algorithms. In order to investigate whether this might be due to an artificial bias introduced by the undersampling of negatives, we have also run the same link prediction experiment using the complete set of negatives (again splitting 90%-10% between train and test). Figures 8a, 8b, and 9 show the precision obtained in this case. Only logistic regression, decision trees, and naïve Bayes are shown, since the others did not scale well to new size of the problem. The precision values achieved in this case are significantly smaller than those obtained with the undersampled dataset (thus confirming the limitations of this approach). Still, they improve over the ones obtained with the standalone similarity measures. Again, it is confirmed that the C1 layer is the most useful for making precise predictions in the gaming-related dataset (Figures 8a). This result is very encouraging, since C1 is the most lightweight in terms of resources consumed. In the generic users dataset (Figures 8b), the role played by the social circles is reversed: using C5 or ACTIVE, we are able to perform predictions that are at least as good as the baseline, however the precision drops as we move towards the innermost layers. Again, even though the precision scores of the baseline and the social-aware method are similar, still the social-aware approach provides an advantage in terms of computational resources, as explained in Section 6.

In summary, supervised learning consistently yields better predictions than their unsupervised counterparts. This may be due to more flexible identification of boundaries between positives and negatives (with respect to the simple threshold-based approach of similarity-based heuristics) or to a smart combination of the similarity-based heuristics. Regardless, social circles again prove very effective in boosting the precision of link prediction strategies, especially in the gaming-related dataset. In particular, supervised strategies seem able to effectively exploit the innermost layer C1 in the domain-dependent dataset and the active network layer in domain-independent dataset

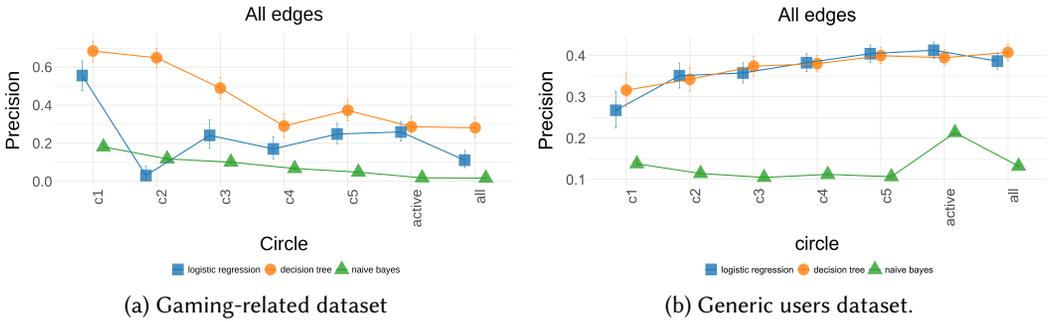


Fig. 8. Precision (with credible intervals) - SUPERVISED ALLEDES on the full datasets.

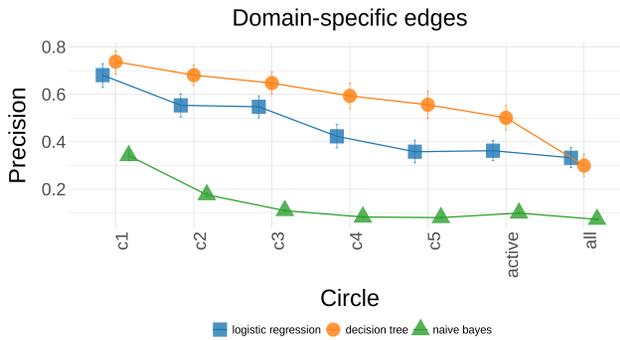


Fig. 9. Precision (with credible intervals) - SUPERVISED DOMAINEDGES on the full gaming-related dataset.

much better than the unsupervised cases. This is an important result, since leveraging C1, or any layer of the ego network, minimises the resources needed to make predictions.

We conclude the section analysing the F1 score, which is the harmonic mean of precision and recall. We only show the results obtained using the full negative set, as they provide a more reliable evaluation. For both the ALLEDES and the DOMAINEDGES case (Figures 10a, 10b, and 11), we observe a different ranking with respect to the corresponding precision results, with the highest F1 being achieved by the naive Bayes approach. Still, since, as discussed in [47], precision is much more important than recall in link prediction, we argue that decision trees should be preferred to naive Bayes. If we only look at the role played by the social circles, we observe that leveraging social circles is always better or as good as the ALL baselines. However, when aiming at striking a balance between precision and recall, C1 is not the best performing circle anymore in the gaming-related dataset, probably due to its high selectivity for the most intimate relationships. In the generic users dataset, instead, the ACTIVE layer is confirmed as the one providing the best prediction performance while consuming less resources.

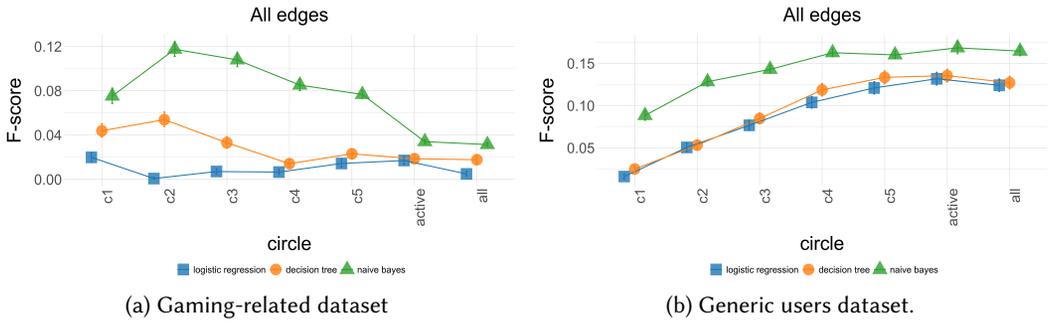


Fig. 10. F1 score (with credible intervals) - SUPERVISED ALLEDES on the full datasets.

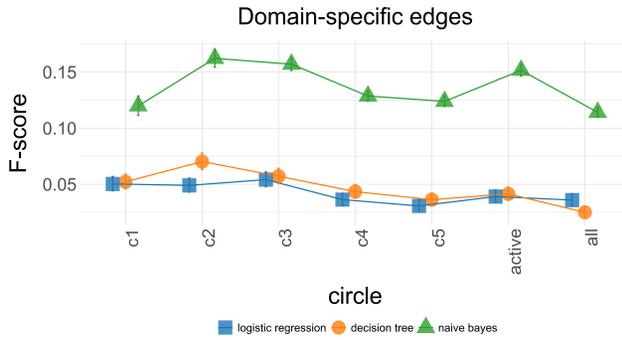


Fig. 11. F1 score (with credible intervals) - SUPERVISED DOMAINEdges on the full gaming-related dataset.

## 5.6 Evaluation of network embedding-based link prediction

Until now, we have studied the social-aware link prediction method by embedding the intimacy levels of Dunbar’s ego network model into existing unsupervised and supervised approaches. The ego network model is a way of capturing the local information of relationships of the individuals. Since the aim of the study is to understand the contribution of the social circles to the performance of the link prediction methods, we selected prediction methods where the feature calculation is not a black-box, and the link prediction method is not an end-to-end method. All the methods discussed so far were based on explicit graph features. In the literature, there are newly suggested link prediction methods based on *latent features* of the graphs, where the feature vectors are low-dimensional vector representations, produced by approaches such as graph representation learning [9] and graph neural networks [48]. For a detailed description and classification of the methods, the reader may refer to [9, 28, 41, 48]. In this section, our goal is to carry out a preliminary analysis, where the results we obtained using explicit features are compared against link prediction approaches based on self-learning of latent graph features. To this aim, we selected one of the most popular graph embedding methods, *node2vec* [23].

*node2vec* is a framework that maps the nodes of a graph to low-dimensional feature vectors while preserving nodes’ neighborhoods [23]. It proposes a biased random walk based on breadth-first sampling (BFS) and depth-first sampling (DFS) with transition probabilities  $p$  and  $q$ . BFS ( $p$ ) captures the embeddings related to structural equivalence (e.g., being a hub node) while DFS ( $q$ ) captures embeddings related to communities based on homophily. There exist other four parameters for *node2vec*:  $d$  the dimension/length of the feature vector,  $l$  the length of random walk,  $k$  the size of the neighborhood, and  $r$  the number of walks per node. A feature vector is produced per node in the network by combining the Skip-gram architecture of *word2vec* with a flexible neighborhood definition obtained with biased random walks. For the link prediction task, these node embeddings are transformed into edge embeddings by applying binary operators like, e.g., Hadamard. Then, logistic regression is applied on the edge embeddings to train a binary classification model on positive and negative data (existing and missing links).

Given that *node2vec* is designed to capture latent features, in this study, we apply *node2vec* method as is, without explicitly adding information on the social circles. Since the aim is to predict the possible future links between the egos whose social circles can be computed, as described in Section 4, we learn the node embeddings of the vertex set  $V_e$  by using all edges  $E_e$ . For the extraction of node embeddings, we try the different combinations of  $p$  and  $q$  values ( $p, q \in \{0.25, 0.5, 1, 2, 4\}$ ) while keeping other parameters of random walk fixed ( $d = 128, l = 80, k = 10, r = 10$ ) as mentioned in [23]. Then, we calculate the edge embeddings of all existing and missing links in  $V_e \times V_e$ . Given two feature vectors,  $f(i)$  and  $f(j)$ , of egos  $i$  and  $j$  where  $i, j \in V_e$ , the edge embedding of the edge  $e_{ij}$  is calculated as:  $f(i) \circ f(j)$  where  $\circ$  is a binary operator (average, Hadamard, weighted-L1, weighted-L2). We are using Hadamard operator, which is shown to be the most effective one in [23].

Figure 12 shows the precision results of the undersampled generic users dataset and DOMAINEDGES scenario of undersampled gaming-related dataset for different  $(p, q)$  pairs. Recall that we are using *node2vec* approach as a baseline method where all edges are included without any slicing (this corresponds to the case ALL in the previous section). We are using microaveraging of the precision scores of  $k$ -fold cross-validation as described in Section 5.1. In order to observe the effects of  $p$  and  $q$  on the link prediction results, we didn’t apply  $k$ -fold cross validation within our training sets to optimize  $(p, q)$  pairs. Therefore, the results we show here are actually overestimated in favor of *node2vec* approach against the other methods we mentioned in Sections 5.4 and 5.5. As we can see from the figures, the precision values in the grid of parameters  $(p, q)$

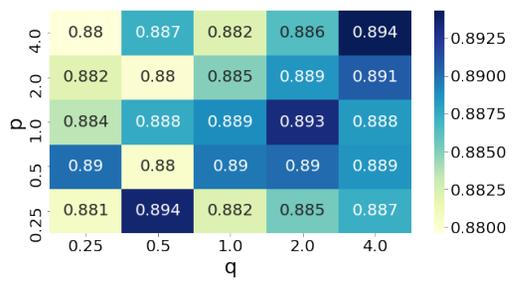
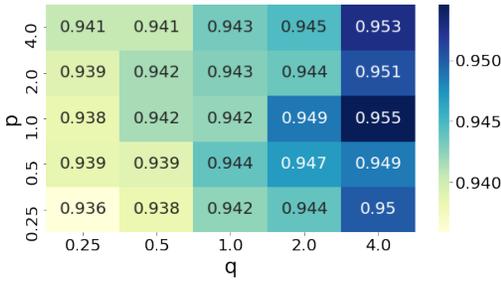
are quite close to each other in all cases. However, it seems that  $q$  plays a more important role than  $p$ , as a high  $q$  yields very good precision regardless of the  $p$  value. Since a high  $q$  biases the random walk to stay closer to the current node (BFS), homophily seems to be more important than structural equivalence in our datasets.

Let us now compare the results in Figure 12 against those obtained in Section 5.5<sup>6</sup>. The maximum precision score achieved with *node2vec* for the generic users dataset is 0.955, as shown in Figure 12a, that is approximately the same with other baseline methods visualized as ALL category in Figure 6b. The only social circle playing a major role in Figure 6b is C1. In C1, Random Forest and naïve Bayes perform better than other strategies (including *node2vec*) while the performance drops drastically for logistic regression and decision tree. In C1, Random Forest 0.974 and naïve Bayes both yield a precision of 0.974, hence they perform better than any other strategy, including *node2vec*. In the DOMAINEDGES scenario of the gaming-related dataset, the best precision score of *node2vec* is 0.894, which is slightly smaller than other baseline methods shown in Figure 7, but the difference is not significant. When leveraging social circles, Random Forest and naïve Bayes (in C1, they yield a precision value of 0.985 and 0.981 respectively) performs better than any other strategies, including baseline methods and *node2vec*, using the innermost circles (C1, C2, C3). Using C1, C2, C3, any strategy but decision trees perform better than *node2vec*. To conclude this analysis, we also consider the ALLEDGES scenario with the undersampled gaming-related dataset (not reported in Figure 12). In this case, the best precision achieved by *node2vec* is 0.91. If we compare this result with the ones obtained in Figure 6a for the supervised strategies, we see that *node2vec* outperforms all of them when they leverage circle C3 or higher. In C1, naïve Bayes, Random Forest, and SVM (with a precision value of 0.945, 0.927, and 0.931 respectively) perform better than *node2vec*, in C2 the performance of naïve Bayes, Random Forest (with a precision value of 0.912 and 0.915, respectively), and *node2vec* are comparable.

Overall, the latent features learnt by *node2vec* seem to be very effective for link prediction, as we have shown that *node2vec* approach performs as well as other baseline methods that use all edge information without any circle-based slicing. On the other hand, the circle-aware approaches are able to perform as well as *node2vec* approach, in some cases even better. This is especially true when the innermost circles are used, where social-aware supervised learning achieves maximum precision while allowing to save computational resources, which is a plus considering the scale of social network data.

---

<sup>6</sup>Please note that this comparison is fair, since we are studying machine learning-based algorithms trained on explicit features against a machine learning-based algorithm trained on latent features. Vice versa, we do not consider the unsupervised approaches from Section 5.4 as they do not leverage machine learning techniques.



(a) (Generic users dataset) - NODE2VEC ALLEdGES on the undersampled dataset.

(b) (Gaming-related dataset) NODE2VEC DOMAINEdGES on the undersampled dataset.

Fig. 12. Precision scores for  $(p, q) \in \{0.25, 0.5, 1, 2, 4\}$

**Algorithm 1:** Find the similarity between nodes

**Data:**  $\mathcal{G} = (V_e, E)$ ;  $\Gamma(i)$  = neighborhood of node  $i$ ;  $C_x$  = the social circle to be considered.

**Result:** Similarity  $sim_{C_x}(i, j)$  for each node pair  $i, j$  in  $V_e$ .

```

1 begin
2   for  $i \in V_e$  do
3     // Compute the ego network and slice the graph according to  $C_x$ 
4      $AssignToCircles(\Gamma(i))$ 
5      $\Gamma_{C_x}(i) = Slice(\Gamma(i), C_x)$ 
6   for  $i \in V_e$  do
7     for  $j \in V_e$  do
8       // Compute the similarity
9       if CN then
10         $sim_{C_x}(i, j) = |\Gamma_{C_x}(i) \cap \Gamma_{C_x}(j)|$ 
11       if  $\mathcal{J}C$  then
12         $sim_{C_x}(i, j) = \frac{|\Gamma_{C_x}(i) \cap \Gamma_{C_x}(j)|}{|\Gamma_{C_x}(i) \cup \Gamma_{C_x}(j)|}$ 
13       if AA or RA then
14         $sim_{C_x}(i, j) = 0$ 
15        for  $z \in \Gamma_{C_x}(i) \cap \Gamma_{C_x}(j)$  do
16           $sim_{C_x}(i, j) = sim_{C_x}(i, j) + penalised(\Gamma(z))$ 

```

## 6 COMPLEXITY ANALYSIS

In Section 5, we showed that ego-aware link prediction performs better than traditional methods in the domain-specific dataset (gaming-related), and performs at least as well as traditional methods in the domain-independent dataset of generic users. We repeatedly highlighted the fact that, even when the performance is comparable, leveraging social circles information reduces the complexity (both temporal and spatial), hence the computational resources needed to solve the link prediction problem. In this section, we formally prove that this is the case.

First, we focus on the different time complexity of the baseline and ego-aware strategies. To this aim, let us start from the unsupervised case (discussed in Section 4.1). Algorithm 1 illustrates the steps towards computing the similarity between a pair of nodes depending on the specific similarity considered. Please note that the weighted graph discussed in Section 3 is assumed as input. With respect to the baselines (traditional methods which are represented as ALL in the results) in which no circle information is used, our approach requires two additional steps: (i) computing the ego networks of the nodes for which predictions should be made (line 3) and (ii) pruning the ego networks based on the social circle we want to leverage (line 4). The code from line 5 on is common to both the baselines and the circle-aware approaches, since the specific slicing considered is controlled with  $C_x$  (and  $x = ALL$  for the baseline). In order to estimate the time complexity in the average case, we have to derive the time complexity of computing the ego network ( $\mathcal{O}(\text{ego})$ ) and that of computing the similarity ( $\mathcal{O}(\text{social similarity})$ ), as illustrated in Equation 7. On the other hand, only the similarity complexity is needed for the baseline case (Equation 8).

$$\mathcal{O}(\text{social}) = \mathcal{O}(n)\mathcal{O}(\text{ego}) + \mathcal{O}(n^2) * \mathcal{O}(\text{social similarity}) \quad (7)$$

$$\mathcal{O}(\text{baseline}) = \mathcal{O}(n^2) * \mathcal{O}(\text{baseline similarity}) \quad (8)$$

Clearly, the complexity of the social and baseline similarities depends on the specific approach (CN, JC, AA, RA) considered. We focus first on CN. The complexity  $\mathcal{O}(\text{similarity CN})$  simply corresponds to an intersection between sets (see line 7 of Algorithm 1), hence it is given by Equation 9 for the social-aware case, and by Equation 10 for the baseline. The complexity of a set intersection by using hashing-based dictionaries is  $\mathcal{O}(\min(|\text{Set}_1|, |\text{Set}_2|))$  [6].

$$\mathcal{O}(\text{social similarity}) = \mathcal{O}(\min(|\Gamma_{C_x}(i)|, |\Gamma_{C_x}(j)|)) \sim \mathcal{O}(|\Gamma_{C_x}(i)|) \quad (9)$$

$$\mathcal{O}(\text{baseline similarity}) = \mathcal{O}(\min(|\Gamma(i)|, |\Gamma(j)|)) \sim \mathcal{O}(|\Gamma(i)|) \quad (10)$$

While the above formulas have been derived for CN, it is easy to see that also JC, AA, and RA share the same complexity. Indeed, JC leverages the set union in addition to the set intersection, operation that is still linear in the size of the smaller set. AA and RA are the analogous of CN but with different weights, hence, again, they are linear in the size of the smaller set.

Let us now estimate the complexity  $\mathcal{O}(\text{ego})$  of computing ego networks. Line 3 in Algorithm 1 entails a) sorting the neighbors of node  $i$  ( $\Gamma(i)$ ) based on contact frequencies, b) extracting the active part (corresponding to relationship with contact frequency larger than once per year, as discussed in Section 2.1), and c) running a unidimensional clustering algorithm on the remaining contact frequencies. The complexity of step (a) is  $\mathcal{O}(|\Gamma(i)| * \log(|\Gamma(i)|))$  (e.g., assuming standard merge sort is used [11]), the complexity of step (ii) is  $\mathcal{O}(\log(|\Gamma(i)|))$  which is equivalent to the complexity of a search algorithm on a sorted input (e.g., assuming binary search is used [11]). After step (b), for each ego  $i$ , we are not working anymore on its full neighborhood but only on the active part  $\Gamma_{\text{active}}(i)$ , which is smaller than  $\Gamma(i)$  since it only includes strong relationships. The complexity of step (c) is mainly dependent on the chosen clustering algorithm (note that the classification of alters in circles is robust against the different clustering methods [15]). For example, DBSCAN is quite efficient and it runs in  $\mathcal{O}(|\Gamma_{\text{active}}(i)| \log(|\Gamma_{\text{active}}(i)|))$  [45]. The time complexity of k-means is linear [45] in the number of points to cluster (i.e.,  $\mathcal{O}(|\Gamma_{\text{active}}(i)|)$ ) but it needs additional computations for the selection of the best  $k$  (e.g.,  $\mathcal{O}(k|\Gamma_{\text{active}}(i)|)$  for the partition coefficient method,  $\mathcal{O}(|\Gamma_{\text{active}}(i)|^2)$  for the Silhouette method). Mean Shift runs in  $\mathcal{O}(|\Gamma_{\text{active}}(i)| \log(|\Gamma_{\text{active}}(i)|))$ , like DBSCAN. The last step to carry out is the slicing in line 4, whose complexity is equivalent to that of a search in  $\Gamma_{\text{active}}(i)$ , hence it is logarithmic in its size. In summary, the complexity of computing an ego network is given by the following (assuming DBSCAN or Mean Shift are used for clustering):

$$\begin{aligned} \mathcal{O}(\text{ego}) &= \mathcal{O}(|\Gamma(i)| * \log(|\Gamma(i)|)) + \\ &\quad + \mathcal{O}(\log(|\Gamma(i)|) + |\Gamma_{\text{active}}(i)| \log(|\Gamma_{\text{active}}(i)|) + \log(|\Gamma_{\text{active}}(i)|)) \\ &= \mathcal{O}(|\Gamma(i)| * \log(|\Gamma(i)|)), \end{aligned} \quad (11)$$

where we have leveraged  $|\Gamma(i)| \gg |\Gamma_{\text{active}}(i)|$ . We can now substitute Equations 9, 10, and 11 into Equations 7 and 8 above, thus obtaining:

$$\begin{aligned} \mathcal{O}(\text{social CN}) &= \mathcal{O}(n * |\Gamma(i)| * \log(|\Gamma(i)|) + \mathcal{O}(n^2) * |\Gamma_{C_x}(i)|) \\ &= \mathcal{O}(n^2) * |\Gamma_{C_x}(i)| \end{aligned} \quad (12)$$

$$\mathcal{O}(\text{baseline CN}) = \mathcal{O}(n^2) * |\Gamma(i)|. \quad (13)$$

The right-hand side of Equation 12 stems from the fact that social networks are sparse, hence  $|\Gamma(i)| \ll n$ . Since the number of neighbors in any circle of the ego network is (significantly) smaller than all neighbors in the baseline (because, by definition, social circles only retain the strongest relationships), we have that the time complexity of the ego-aware system is always (significantly) smaller than the baseline.

In order to complete the analysis of the time complexity for the link prediction algorithms considered in this work, we now turn our attention to the supervised case (discussed in Section 4.2).

The complexity of supervised learning can be decomposed into the complexity for extracting the relevant features and the complexity of solving the learning problem on this features. It is easy to see that the latter component is the same regardless of the social circle considered, since the features described in Section 4.2, once computed, are simply four numbers associated with each user  $i$ . Hence, the different complexity between baseline and ego-aware approaches only lies in the computation of the features. Since, as we have discussed in the previous paragraph, computing the ego-aware similarity has smaller time complexity than computing the baseline one, we have again that the ego-aware approach is computationally more convenient also in the supervised case.

The derivation of the space complexity is more straightforward. We will discuss it briefly for the unsupervised case, as for the supervised one it suffices to multiply by 4 (i.e., the number of features considered) the complexity obtained for the unsupervised scenario. For both the baseline and the ego-aware case, we start with an occupancy of  $\mathcal{O}(|V_e| + |E|)$ . However, in the social-aware case, after the slicing we move to  $\mathcal{O}(|V_e| + |E_{C_x}|)$ , where  $|E_{C_x}|$  denotes the set of edges belonging to the circle  $C_x$  (of any node) based on which the slicing is performed. The size of  $E_{C_x}$  is much smaller than  $|E|$ , because it only contains meaningful social relationships. As discussed earlier, the circle size per ego can be as small as 1-2 people (in the case of C1), hence it is clear that circle-based slicing allows us to work with smaller graphs and lower space complexities.

## 7 CONCLUSION

In this paper, we have studied the performance of popular link prediction algorithms enriched with the knowledge on the social structures in the network. Specifically, relying on very well-established models from anthropology, we have considered the social circles in individual ego networks, using the circle as a proxy of intimacy. We have selected four benchmark heuristics (Common Neighbors, Jaccard's Coefficient, Adamic-Adar, and Resource Allocation) and we have modified them to include awareness of the social circles. We have tested them in an unsupervised setting, in which only the potential links with the highest scores are recommended, and in a supervised setting, in which the above heuristics are fed to standard machine learning algorithms (we tested logistic regression, decision trees, naïve Bayes, Random Forest, SVM) that take care of the labelling. In order to evaluate the predictive performance of these solutions, we have used two longitudinal Twitter datasets, one comprising a community of gamers and one containing generic users. Our results show that social-circles-based link prediction is generally very effective. Specifically, in the vast majority of cases, regardless of the prediction approach (unsupervised or supervised), the specific heuristic or learning algorithm, and the metric (precision, AUC, F1 score) considered, leveraging social circles information outperforms the corresponding baseline in which circles are ignored. In case of resource (storage, CPU, time, etc.) limitations, using only information about the innermost social circles guarantees the same performance achieved when using the whole network. In addition, using social circles information seems to provide the same performance as using additional classifiers on nodes, which might be impractical or costly to set up. Finally, we have carried out a preliminary comparison of social-aware supervised link prediction against graph-embedding-based link prediction. When used for link prediction, explicit features that leverage social circle information are able to beat latent features extracted using graph embedding methods. This generally happens when the innermost circles are used, thus combining resource savings with top prediction performance.

## ACKNOWLEDGMENTS

This work was partially funded by the SoBigData++, HumaneAI-Net, MARVEL, and OK-INSAlD projects. The SoBigData++ project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 871042. The HumaneAI-Net project

has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No (under definition). The MARVEL project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 957337. The OK-INSAID project has received funding from the Italian PON-MISE program under grant agreement ARS01 00917.

## REFERENCES

- [1] Lada A Adamic and Eytan Adar. 2003. Friends and neighbors on the Web. *Social Networks* 25, 3 (jul 2003), 211–230.
- [2] Luca Maria Aiello and Nicola Barbieri. 2017. Evolution of Ego-Networks in Social Media with Link Recommendations. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining (Cambridge, United Kingdom) (WSDM '17)*. Association for Computing Machinery, New York, NY, USA, 111–120. <https://doi.org/10.1145/3018661.3018733>
- [3] Sinan Aral and Marshall Van Alstyne. 2011. The diversity-bandwidth trade-off. *Amer. J. Sociology* 117, 1 (2011), 90–171.
- [4] Valerio Arnaboldi, Marco Conti, Massimiliano La Gala, Andrea Passarella, and Fabio Pezzoni. 2014. Information Diffusion in OSNs: The Impact of Nodes’ Sociality. In *Proceedings of the 29th Annual ACM Symposium on Applied Computing (Gyeongju, Republic of Korea) (SAC '14)*. Association for Computing Machinery, New York, NY, USA, 616–621. <https://doi.org/10.1145/2554850.2555000>
- [5] Valerio Arnaboldi, Marco Conti, Andrea Passarella, and Fabio Pezzoni. 2013. Ego networks in Twitter: an experimental analysis. In *2013 Proceedings IEEE INFOCOM*. IEEE, IEEE, New York, NY, USA, 3459–3464.
- [6] Philip Bille, Anna Pagh, and Rasmus Pagh. 2007. Fast Evaluation of Union-Intersection Expressions. In *Algorithms and Computation*, Takeshi Tokuyama (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 739–750.
- [7] Catherine A Bliss, Morgan R Frank, Christopher M Danforth, and Peter Sheridan Dodds. 2014. An evolutionary algorithm approach to link prediction in dynamic social networks. *Journal of Computational Science* 5, 5 (2014), 750–764.
- [8] Chiara Boldrini, Mustafa Toprak, Marco Conti, and Andrea Passarella. 2018. Twitter and the press: an ego-centred analysis. In *Companion Proceedings of the The Web Conference '18*. 1471–1478.
- [9] Hongyun Cai, Vincent W Zheng, and Kevin Chen-Chuan Chang. 2018. A comprehensive survey of graph embedding: Problems, techniques, and applications. *IEEE Transactions on Knowledge and Data Engineering* 30, 9 (2018), 1616–1637.
- [10] Lei Cai, Jundong Li, Jie Wang, and Shuiwang Ji. 2020. Line graph neural networks for link prediction. *arXiv preprint arXiv:2010.10046* (2020).
- [11] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. 2009. *Introduction to algorithms*. MIT press, Cambridge, MA, USA.
- [12] William Cukierski, Benjamin Hamner, and Bo Yang. 2011. Graph-based features for supervised link prediction. In *Proceedings of the International Joint Conference on Neural Networks*. 1237–1244.
- [13] Sam De Winter, Tim Decuyper, Sandra Mitrović, Bart Baesens, and Jochen De Weerd. 2018. Combining temporal aspects of dynamic networks with Node2Vec for a more efficient dynamic link prediction. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, 1234–1241.
- [14] RI Dunbar. 1998. The social brain hypothesis. *Evolutionary Anthropology* 9, 10 (1998), 178–190.
- [15] R.I.M. Dunbar, Valerio Arnaboldi, Marco Conti, and Andrea Passarella. 2015. The structure of online social networks mirrors those in the offline world. *Social Networks* 43 (2015), 39–47.
- [16] Martin Everett and Stephen P Borgatti. 2005. Ego network betweenness. *Social networks* 27, 1 (2005), 31–38.
- [17] George Forman and Martin Scholz. 2010. Apples-to-apples in cross-validation studies. *ACM SIGKDD Explorations Newsletter* 12, 1 (nov 2010), 49. <https://doi.org/10.1145/1882471.1882479>
- [18] Dario Garcia-Gasulla, Ulises Cortés, Eduard Ayguadé, and Jesús Labarta. 2015. Evaluating Link Prediction on Large Graphs. In *Frontiers in Artificial Intelligence and Applications*. <https://doi.org/10.3233/978-1-61499-578-4-90>
- [19] Global State of Digital. 2019. *The Global State of Digital in 2019 Report*. Technical Report. We are social and Hootsuite.
- [20] Bruno Gonçalves, Nicola Perra, and Alessandro Vespignani. 2011. Modeling users’ activity on twitter networks: Validation of dunbar’s number. *PLoS one* 6, 8 (2011), e22656.
- [21] Leo A. Goodman. 1961. Snowball Sampling. *The Annals of Mathematical Statistics* 32, 1 (1961), 148–170.
- [22] Cyril Goutte and Eric Gaussier. 2005. A Probabilistic Interpretation of Precision, Recall and F-Score, with Implication for Evaluation. In *Advances in Information Retrieval. ECIR 2005*. Springer, Berlin, Heidelberg, 345–359.
- [23] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. 855–864.
- [24] Jan O Haerter, Bjørn Jamtveit, and Joachim Mathiesen. 2012. Communication dynamics in finite capacity social networks. *Physical review letters* 109, 16 (2012), 168701.

- [25] Mohammad Al Hasan, Vineet Chaoji, Saeed Salem, Mohammed Zaki, and New York. 2006. Link Prediction using Supervised Learning. In *SDM06: workshop on link analysis, counter-terrorism and security*.
- [26] Russell A Hill and Robin IM Dunbar. 2003. Social network size in humans. *Human nature* 14, 1 (2003), 53–72.
- [27] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [28] Ajay Kumar, Shashank Sheshar Singh, Kuldeep Singh, and Bhaskar Biswas. 2020. Link prediction techniques, applications, and performance: A survey. *Physica A: Statistical Mechanics and its Applications* 553 (2020), 124289.
- [29] David Liben-Nowell. 2005. *An algorithmic approach to social networks*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [30] David Liben-Nowell and Jon Kleinberg. 2007. The link-prediction problem for social networks. *JASIST* (2007). <https://doi.org/10.1002/asi.20591>
- [31] Ryan N Lichtenwalter, Jake T Lussier, and Nitesh V Chawla. 2010. New perspectives and methods in link prediction. In *Proceedings of ACM SIGKDD*. 243–252. <https://doi.org/10.1145/1835804.1835837>
- [32] Nan Lin, Karen S Cook, and Ronald S Burt. 2001. *Social capital: Theory and research*. Transaction Publishers, Piscataway, New Jersey, USA.
- [33] Weiping Liu and Linyuan Lü. 2010. Link prediction based on local random walk. *EPL (Europhysics Letters)* 89, 5 (2010), 58007.
- [34] Linyuan Lü and Tao Zhou. 2010. Link prediction in weighted networks: The role of weak ties. *EPL* 89, 1 (2010). <https://doi.org/10.1209/0295-5075/89/18001>
- [35] Linyuan Lü and Tao Zhou. 2011. Link prediction in complex networks: A survey. *Physica A: Statistical Mechanics and its Applications* 390, 6 (mar 2011), 1150–1170. <https://doi.org/10.1016/J.PHYSA.2010.11.027>
- [36] Víctor Martínez, Fernando Berzal, and Juan-Carlos Cubero. 2016. A Survey of Link Prediction in Complex Networks. *Comput. Surveys* 49, 4 (dec 2016), 1–33. <https://doi.org/10.1145/3012704>
- [37] Christopher McCarty. 2002. Structure in personal networks. *Journal of social structure* 3, 1 (2002), 20.
- [38] Miller McPherson, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a feather: Homophily in social networks. *Annual review of sociology* 27, 1 (2001), 415–444.
- [39] Aditya Krishna Menon and Charles Elkan. 2011. Link prediction via matrix factorization. In *Joint european conference on machine learning and knowledge discovery in databases*. Springer, 437–452.
- [40] Giovanna Miritello, Esteban Moro, Rubén Lara, Rocío Martínez-López, John Belchamber, Sam G.B. Roberts, and Robin I.M. Dunbar. 2013. Time as a limited resource: Communication strategy in mobile phone networks. *Social Networks* 35, 1 (jan 2013), 89–95. <https://doi.org/10.1016/j.socnet.2013.01.003>
- [41] Ece C Mutlu, Toktam Oghaz, Amirarsalan Rajabi, and Ivan Garibay. 2020. Review on Learning and Extracting Graph Features for Link Prediction. *Machine Learning and Knowledge Extraction* 2, 4 (2020), 672–704.
- [42] Joshua O’Madadhain, Jon Hutchins, and Padhraic Smyth. 2005. Prediction and ranking algorithms for event-based network data. *ACM SIGKDD Explorations Newsletter* 7, 2 (2005), 23–30. <https://doi.org/10.1145/1117454.1117458>
- [43] Sucheta Soundarajan and John Hopcroft. 2012. Using community information to improve the precision of link prediction methods. In *WWW’12 - Proceedings of the 21st Annual Conference on World Wide Web Companion*. 607–608.
- [44] Alistair G Sutcliffe, Di Wang, and Robin IM Dunbar. 2015. Modelling the role of trust in social relationships. *ACM Transactions on Internet Technology (TOIT)* 15, 4 (2015), 16.
- [45] Pang-Ning Tan, Michael Steinbach, Anuj Karpatne, and Vipin Kumar. 2019. *Introduction to Data Mining* (2nd edition ed.). Pearson, New Delhi, India.
- [46] Dashun Wang, Dino Pedreschi, Chaoming Song, Fosca Giannotti, and Albert-Laszlo Barabasi. 2011. Human mobility, social ties, and link prediction. In *Proceedings of the ACM SIGKDD’11*. ACM Press, New York, New York, USA, 1100.
- [47] Yu Wang and Jihong Li. 2016. Credible Intervals for Precision and Recall Based on a K-Fold Cross-Validated Beta Distribution. *Neural Computation* 28, 8 (aug 2016), 1694–1722.
- [48] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* (2020).
- [49] Yang Yang, Ryan N Lichtenwalter, and Nitesh V Chawla. 2015. Evaluating link prediction methods. *Knowl Inf Syst* 45 (2015), 751–782. <https://doi.org/10.1007/s10115-014-0789-0>
- [50] Muhan Zhang and Yixin Chen. 2017. Weisfeiler-lehman neural machine for link prediction. In *Proceedings of ACM SIGKDD*, Vol. Part F1296. 575–583.
- [51] Muhan Zhang and Yixin Chen. 2018. Link prediction based on graph neural networks. In *Advances in Neural Information Processing Systems*, Vol. 2018-Decem. 5165–5175. arXiv:1802.09691
- [52] Peng Zhang, Xiang Wang, Futian Wang, An Zeng, and Jinghua Xiao. 2016. Measuring the robustness of link prediction algorithms under noisy environment. *Scientific Reports* 6 (2016). <https://doi.org/10.1038/srep18881>
- [53] Tao Zhou, Linyuan Lü, and Yi-Cheng Zhang. 2009. Predicting missing links via local information. *Eur. Phys. J. B* 71 (2009), 623–630. <https://doi.org/10.1140/epjb/e2009-00335-8>

- [54] W-X Zhou, D Sornette, R a Hill, and R I M Dunbar. 2005. Discrete hierarchical organization of social group sizes. *Proceedings. Biological sciences / The Royal Society* 272, 1561 (2005), 439–444.
- [55] Boyao Zhu and Yongxiang Xia. 2016. Link Prediction in Weighted Networks: A Weighted Mutual Information Model. *PLOS ONE* 11, 2 (feb 2016), e0148265. <https://doi.org/10.1371/journal.pone.0148265>