Multilayer Brokerage in Geo-Social Networks

Desislava Hristova
Computer Laboratory
University of Cambridge, UK
desislava.hristova@cl.cam.ac.uk

Pietro Panzarasa
School of Business and Management
Queen Mary University of London, UK
p.panzarasa@qmul.ac.uk

Cecilia Mascolo
Computer Laboratory
University of Cambridge, UK
cecilia.mascolo@cl.cam.ac.uk

Abstract

Open network structures and brokerage positions have long been seen as playing a crucial role in sustaining social capital and competitive advantage. The degree to which individuals intermediate between otherwise disconnected others can differ across online and offline social networks. For example, users may broker online between two others who then exchange offline the information received through social media. Yet network studies of social capital have often neglected the interplay between online and offline interactions, and have concentrated primarily on a single layer. Here, we propose a geo-social multilayer approach to brokerage that casts light on the integrated online and offline foundations of social capital. Drawing on a data set of 37,722 Foursquare users in London, we extend the notion of brokerage by examining users' positions in an online social network and their offline mobility patterns through check-ins. We find that social and geographic brokerage positions are distinct and asymmetric across the social and co-location networks. Users may appear to be brokers online when in fact their ability to intermediate would be mitigated if their offline positions were also taken into account. On the other hand, users who appear to have little brokerage power offline may be active brokers within networks that combine both online and offline interactions.

Introduction

The sociological tradition that places emphasis on gaps in social structures can be traced back to the late 1960s and early 1970s, when a group of sociologists began to develop the general idea that it is advantageous to forge connections to multiple, otherwise disconnected, individuals or groups (Cook and Emerson 1978; Freeman 1977; White 1970). One of the most celebrated theoretical endeavours in the social sciences that draws on this tradition is Granovetter's influential study of the bridging role of weak ties (Granovetter 1973). The broader the access to weaker ties, the closer an actor is likely to be placed to discontinuities in the social structure, which in turn enables the actor to be connected to various social circles of contacts and to be exposed to novel and diverse sources of information.

By placing emphasis on the correlation between tie weakness and the flow of information, Granovetter set the stage for a conception of social capital based on discontinuities in social structures and brokerage opportunities. The idea that social capital can originate from brokerage opportunities stemming from structural holes has been further explored by Burt, especially in organisational domains (Burt 1992). Burt defines a structural hole as the "separation between non-redundant contacts", "a buffer" that enables the two contacts to "provide network benefits that are in some degree additive rather than overlapping" (Burt 1992, p. 18). Burt further identifies two sources of the social capital originating from structural holes: information benefits and control benefits. The less an actor's contacts are already connected with each other, the more likely the actor is to obtain non-redundant information and to reap control advantages by trading-off contacts' requests against each other. From this vantage point, an actor's social capital is seen as a function of the brokerage opportunities arising from the structural holes the actor's social relationships span (Burt 1992; Lin 2001; Stovel and Shaw 2012).

Brokerage has been studied in various contexts within the boundaries of a single network - e.g., organisational relations (Burt 2005), online social network friendship (Lou and Tang 2013), or mobile communication (Eagle, Macy, and Claxton 2010). Brokerage, however, is not limited to one type of context. The same individuals may be engaged in different types of social relationships, and as a result may benefit from various brokerage positions that affect each other in complex ways. For instance, people can be brokers both online and offline. Recent studies have suggested that online social networking is directly linked to bridging and bonding social capital (Ellison et al. 2014; Vitak, Ellison, and Steinfield 2011), where social media sites such as Facebook have been shown to have a significant role in maintaining distant and near-by contacts, and in sustaining social capital (Burke and Kraut 2013). However, while an individual may seem to be embedded in a socially cohesive neighbourhood online (offline), there may be opportunities for brokerage arising from the non-redundant contacts the individual has offline (online). *Despite the growing availability and ubiquity of online social media, it is still unclear whether there is a trade-off between the brokerage positions that individuals occupy online and offline.*

In this work, we take a step in this direction, and propose a multi-relational perspective on brokerage using a multilayer network approach to social capital, where both online social network structure and physical co-location are taken into account to detect brokerage positions. In particular, we are interested in analysing the brokerage positions that the same individuals hold in both layers. Inspired by Burt's seminal work on structural holes and social capital (Burt 1992), we propose a multilayer approach to social capital, according to which brokerage opportunities are regarded as emerging from the interplay between an individual's online and offline contacts. In this way, an individual's opportunities that would otherwise remain undetected within a single layer in isolation can be unmasked when multiple layers are allowed to contribute to the individual's social capital.

The main contributions of this work can be summarised as follows:

- We propose a multilayer perspective, which allows for flexibility in the analysis of the structural properties of networks by representing each form of social interaction as a distinct layer of a social system. We apply this multilayer approach to the study of brokerage in online and offline social networks.

- Drawing on a data set of 37,722 users of the popular location-based social network Foursquare along with their check-ins and contacts, we detect online and offline brokerage positions using structural measures of social capital.

- Through the use of a number of measures of brokerage, we find that most users maintain either social or geographic brokerage but not both, thus suggesting a trade-off between the online and offline sources of social capital.

- We extend the notion of structural holes beyond a single layer to allow for brokerage opportunities across networks, and find that brokerage positions change when both the online and the co-location layers are taken into account.

- We explore the diversity of brokers' neighbourhoods in terms of their contacts' preference for categories of places visited, and find that, in qualitative agreement with extant literature, brokers in online and offline networks (and across) tend to maintain more diverse neighbourhoods than individuals with low brokerage potential.

The remainder of this work is structured as follows. We first review the relevant literature on social capital, online and offline brokerage, and multilayer networks, followed by a description of the methodology and the data set. We present our multilayer-network study of geo-social brokerage resulting from the interplay between online and offline structural positions, and examine the heterogeneity of brokers' local neighbourhoods. We conclude with a discussion of the implications and limitations of our findings, and suggest avenues for future work on brokerage in multilayer networks.

Related work

Social capital. Social capital refers to the value embedded in social networks in the form of connections which can potentially offer support (strong ties) and opportunities (weak ties), and more generally yield expected returns in the marketplace, including the community, the economic, financial, and political markets (Lin 2001). Individuals who maintain high social capital in the form of information brokerage are known to have more diverse neighbourhoods both in terms of novelty (Burt 1992) and geographic dispersion (Eagle, Macy, and Claxton 2010), which provides them with advantageous resources as compared to their peers. The competitive advantage of actors in social networks has been formalised as a function of the structural holes that the actors span between otherwise disconnected pairs of others. Spanning holes provides actors with access to negotiation and mediation power (Burt 1992). Brokerage and structural holes have been studied in a variety of contexts including online social networks (Lou and Tang 2013). However, relatively little work has so far examined the interplay between online and offline relationships and their combined contribution to the generation of social capital.

Online and offline brokerage. Most research efforts in the area of online and offline social capital have focused on establishing the role of the Internet and social media in the accumulation of social capital (Vitak, Ellison, and Steinfield 2011; Ellison, Steinfield, and Lampe 2011; Ellison et al. 2014). Notably, the authors in (Wellman et al. 2001) find that the usage of the Internet supplements social capital but does not increase or decrease it. More recent findings show that the use of social networking sites such as Facebook may in fact increase offline social capital by converting latent ties into online weak ties (Ellison, Steinfield, and Lampe 2011) and by allowing for a larger number of online heterogenous weak ties (Steinfield et al. 2009). Users of Facebook and other social media have been associated with an increased social capital when compared to non-users (Lampe, Vitak, and Ellison 2013; Steinfield et al. 2009). Our work builds on these findings, and further explores the structural properties of brokerage in two parallel online and offline networks.

Multilayer networks. The multilayer nature of systems has been explored in a variety of contexts from global air transportation (Cardillo et al. 2013) to massive online multiplayer games (Szell, Lambiotte, and Thurner 2010). A comprehensive review of multiplex (or multilayer) network models can be found in (Kivelä et al. 2014). Most works in the field are mainly theoretical with limited applications to social networks. Nevertheless, multilayer networks are useful representations of multi-relational data sets where more than one type of connection exist between the same pairs of nodes. Recently, multilayer networks have been applied to the study of social tie strength in multilayer interactions (Hristova, Musolesi, and Mascolo 2014). In a similar fashion, we propose a multilayer perspective to study brokerage opportunities individuals can exploit within and across networks.

Measuring geo-social brokerage

The brokerage potential of a node in a social network provides the node with an advantageous position compared to other nodes. In particular, nodes with high brokerage potential are typically characterised as being positioned near structural holes separating otherwise disconnected pairs of nodes (Burt 1992). In this section, we extend the notion of structural hole to a multilayer context so as to enable brokerage opportunities to arise from an individual's combined online and offline social network.

Multilayer network approach

Social networks are typically represented as single layers, where nodes are connected by one type of relationship such as friendship or collaboration. While single-layer networks may be sufficient in many cases, they do not realistically capture the online and offline interaction between people, which is becoming increasingly ubiquitous. We therefore propose a multilayer geo-social approach to the study of social relationships so as to identify the geographic layer of physical co-presence as well as the online social interaction layer. This will enable us to study the brokerage potential of people within and across layers.

We represent the structures of the online social network and of the geographic co-location network as a single multilayer network where the geographic and social layers can be considered as autonomous networks that share the same set of nodes, as illustrated in Figure 1. More formally, a multilayer network can be referred to as an M-layer graph $\mathcal{M} = \{G^1, ..., G^\alpha, ..., G^M\}$, where a link between nodes i and j on layer α is denoted as ℓ_{ij}^α. The same pair of nodes are allowed to be connected on both layers. Although here we focus on the two-layer scenario, our multilayer study can easily be further extended to networks of any number M of layers.

Measures of brokerage

We shall now outline the traditional measures of brokerage, and then extend them to the multilayer context.

Effective size. We use the measure of effective size as defined in Burt's seminal work on structural holes (Burt 1992). Burt defines the effective size of a node's neighbourhood as the non-redundant portion of it. Formally, in the most general case of a directed weighted network we define $z_{ij} + z_{ji}$ as the sum of the weights of the two links connecting node i and node j, while $\sum_m (z_{im} + z_{mi})$ is the total strength of node i. The effective size S_i of node i can be expressed as:

$$S_i = \sum_j \left[1 - \sum_q p_{iq} m_{jq} \right], \quad q \neq i, j, \qquad (1)$$

where

$$p_{iq} = \frac{z_{iq}}{\sum_j z_{ij}}, \quad i \neq j \qquad (2)$$

is the normalized weight of the link between nodes i and q over i's local neighborhood in the case of an undirected weighted network, and where
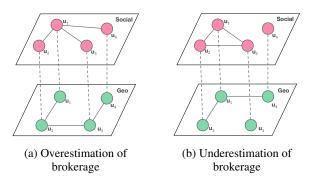

(a) Overestimation of   (b) Underestimation of
brokerage   brokerage

Figure 1: Geo-social networks

$$m_{jq} = \frac{z_{jq}}{\max_k(z_{jk})}, \quad j \neq k \qquad (3)$$

is the marginal strength of node j's link to node q (i.e., the weight z_{jq} of the link connecting nodes j and q divided by the weight of the strongest link node j has with any of its contacts). In turn, m_{jq} can be reduced to z_{jq} in the case of unweighted and undirected networks in which the maximum value a link can take is one. According to Equation 1, we have: $1 \leq S_i \leq k_i \; \forall i$, except for the case in which node i is an isolate (where S_i is set equal to zero). Thus, the effective size of node i's neighbourhood ranges from the minimum value of one, when all pairs of node i's contacts are connected with each other, to the maximum value equal to the node's degree k_i (i.e., the number of links incident upon node i) when there is no link connecting any pair of i's contacts. What the effective size essentially measures is therefore the number of non-redundant connections a node has to otherwise disconnected others. Such connections contribute to the structural hole in the node's local neighbourhood.

Efficiency. The efficiency E_i of node i's local neighbourhood refers to the proportion of i's neighbourhood that is non-redundant (Burt 1992). This can easily be derived by dividing the effective size by the degree of the node:

$$E_i = \frac{S_i}{k_i}. \qquad (4)$$

Efficiency thus helps shed light on the degree to which an individual's effort to expand social capital is directed toward novel social circles that provide exposure to non-overlapping and diverse sources of information (Burt 1992).

Multilayer brokerage. When individuals belong to multilayer networks, such as the one in Figure 1, one may over- or underestimate their brokerage potential if only a single layer is analysed (e.g., only online communication). For example, in Figure 1a node u_1 brokers between nodes u_2, u_3 and u_4 in the online social network layer. There are three links connecting node u_1 to three non-redundant contacts. The degree of node u_1 is equal to the effective size of u_1's local neighbourhood (i.e., three), and the efficiency is therefore one. This places node u_1 in an advantageous brokerage position in which it can bridge three structural holes and intermediate

between otherwise disconnected contacts. There is no feasible channel available to the other three nodes for communicating directly and exchanging information with one another. In this sense, node u_1 is needed to secure communication among nodes u_2, u_3, and u_4. However, when the geographic layer is taken into account in addition to the social layer, new links between node u_1's contacts become apparent: a link between node u_2 and node u_3, and a link between node u_3 and node u_4. This increases the redundancy of node u_1's contacts, thus reducing the effective size and efficiency of u_1's overall (i.e., geo-social) local neighbourhood to $5/3$ and $5/9$, respectively. An opposite problem of underestimation is illustrated in Figure 1b, where it appears that node u_1 has a fully redundant local neighbourhood in the online social layer, where u_1's contacts (i.e., u_2 and u_3) are connected with each other. In this case, the effective size of u_1's local neighbourhood is one, while the efficiency is $1/2$. However, when the geographic layer is taken into account, we find that node u_1 has an additional non-redundant contact, u_4, that has no connection with node u_2 and node u_3. This therefore increases the effective size of node u_1's overall local neighbourhood from one (in the social layer only) to $7/3$. Correspondingly, efficiency increases from $1/2$ to $7/9$. The two examples in Figure 1a and Figure 1b thus clearly suggest that *the analysis of brokerage opportunities can be biased by problems of over- and underestimation when only a single network layer is taken into consideration*.

To overcome this, we extend the notion of brokerage using a multilayer graph, where geographic and social links are both regarded as feasible communication channels that can jointly provide nodes with opportunities for brokering between others. This means that we must allow for cross-layer triads and triangles. The most straightforward way to achieve this is to calculate the union of the two layers, such that $G^{\alpha \cup \beta} = (V^{\alpha \cup \beta}, E^{\alpha \cup \beta})$. This reduces the problem to a single graph to which the existing measures of brokerage can be applied. Since we are interested in the structure of the resulting combined network, we do not assign different weights to each layer, even though in other problems this may be methodologically appropriate.

Data set

Our data set consists of the check-ins and links connecting $37,722$ active users of the location-based social network Foursquare in London, UK. We have downloaded $549,797$ check-ins, each representing a visit made by a user to a certain venue at a certain time and date. These check-ins have been made to $43,584$ venues, and have been posted to Twitter by the users in the period between December 2010 and September 2011, with their respective social networks downloaded at the end of that period. First, we build a *social network* from the reciprocal Twitter following between all Foursquare users who have shared their check-ins on Twitter. This procedure is often regarded as appropriate for identifying friendship, and is consistent with the construction of other undirected networks (Kwak et al. 2010). Because we want to detect structural holes spanned by users both in the online and in the co-location networks, we focus only on the
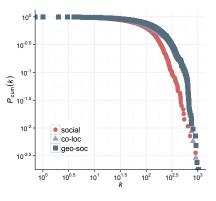

Figure 2: Cumulative degree distributions $P_{cum}(k)$ of the single-layer networks and the geo-social network

Foursquare users of Twitter.

The *co-location network* is built on top of the social network, with the same nodes and using the check-ins posted to Twitter by the Foursquare users. Two nodes in this network are connected if they were co-located, i.e., they happened to be at the same place and at the same time, which reflects the potential for exchanging information offline. We construct the network by using the timestamp of user's check-ins to venues, where if two users have checked-in to the same venue within a 1-hour window, we place a link between them in the co-location network. We only establish a link between two users when they were co-located more than once in order to minimise false positives. In this way, we use a proxy for offline interaction between Foursquare users, in agreement with previous studies of co-location from location-based services (Cranshaw et al. 2010).

Finally, we build the *geo-social multilayer network* as described in the previous section, by taking the set of links that are produced by the union of the social and co-location networks. The multilayer network is undirected, and links in this network are unweighted. Table 1 outlines the following network properties of the two single-layer networks and the combined multilayer network: the number of nodes, V; the number of links, L; the average degree, $< k >$; and the average (local) clustering coefficient (Watts and Strogatz 1998), $< C >$. The cumulative degree distributions of all three networks are shown in Figure 2.

network	$|V|$	$|L|$	$< k >$	$< C >$
social	36,926	176,164	9	0.15
co-location	8,059	112,367	27	0.51
geo-social	37,722	287,661	15	0.2

Table 1: Network properties

In what follows, we shall describe the results of our analysis of brokerage within and across the two networks.

Results

In this section, we aim to shed light on the interplay between online social network brokerage and co-location network

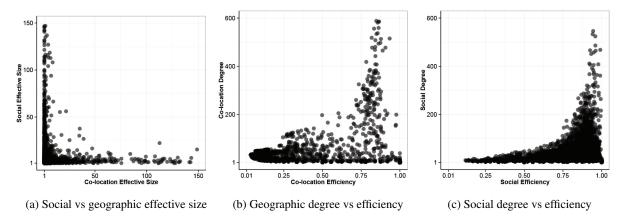| (a) Social vs geographic effective size | (b) Geographic degree vs efficiency | (c) Social degree vs efficiency |

Figure 3: Effective size and efficiency of nodes' local neighbourhoods

brokerage in the Foursquare network of users in London. To this end, we first investigate the changes in the effective size and efficiency of users' local neighbourhoods from one layer to the other, and then measure users' brokerage opportunities in the combined multilayer geo-social network.

Social vs geographic brokerage

We use Equation 1 to capture the brokerage positions of all nodes that belong to both the co-location (geographic) and social network layers. Nodes with degree equal to zero in either layer are therefore removed from the analysis. In Figure 3a, we pit the effective size of nodes in the geographic co-location layer against the nodes' effective size in the social network layer (up to a size of 150 for the sake of visibility). The majority of nodes are associated with a high brokerage potential only in one of the two layers, but not across layers. This suggests that users may seem to intermediate between others when evaluated within a single layer, when in fact their opportunities for brokerage are much fewer when the two layers are combined.

The social network degree and the co-location network degree of nodes bear no correlation. Since a node's opportunity for brokerage greatly depends on the node's degree, it would be unreasonable to expect a correlation between the effective sizes of the node's local neighbourhoods in both layers. This suggests that there is a trade-off between being physically co-located with many others in the co-location network and having many friends online. Correspondingly, there is also a trade-off between brokerage positions online and offline, as indicated by our findings. *Individuals that hold prominent brokerage positions either in the social or in the geographic layer of the network are not likely to hold an equally prominent position in the combined two-layer network.* This also implies that *the assessment of brokerage confined within a single layer in isolation may be positively or negatively biased.*

Figures 3b and 3c show the efficiency of each node's neighbourhood as compared to the node's degree. Although the vast majority of users have a low degree and low efficiency, it is clear from the figures that nodes with a degree higher than 200 tend to have an efficiency between 0.75 and 1.00. This thus suggests that as nodes increase the number of their contacts, they also optimise the efficiency of their local neighbourhoods by including non-redundant contacts, thus expanding the opportunities for brokerage (Burt 1992). As shown in Figure 3a, *high-degree nodes appear to be able to secure a higher efficiency of their networks than low-degree nodes.* This finding is consistent with other related studies that have documented an inverse relationship between the degree and the local clustering of nodes (Vázquez, Pastor-Satorras, and Vespignani 2002; Ravasz and Barabási 2003). Because nodes with high degrees tend to have a lower clustering coefficient than nodes with low degrees, the former nodes are also expected to leverage on greater brokerage opportunities than the latter ones (Latora, Nicosia, and Panzarasa 2013).

When the effective size of nodes' neighbourhoods is evaluated only in the co-location network, the nodes with high brokerage potential do not correspond to the nodes with high potential in the social network layer. Therefore, we propose a combined geo-social approach to brokerage, which unifies both interaction layers into an integrated source of brokerage opportunities. We shall expand on this approach and its applications in the following sections.

Geo-social brokerage

Geo-social brokerage is here regarded as arising from the union of the co-location network and the online social network. Information can indeed be transmitted both online and offline, and whether individuals can benefit from having access to novel and non-overlapping sources of information can only be investigated by analysing the structural positions individuals hold in the combined online and offline network. The union of the two layers has an effect on the degree of the individual user. In particular, for each node the degree in the multilayer network cannot be smaller than the lower degree the node has in either single-layer network, and cannot be higher than the sum of the two degrees the node has in the two single-layer networks. On the one hand, the degree in the multilayer network takes on the minimum value when all the contacts a node has in one layer are the same as the contacts in the other layer. On the other, it takes on its maximum

value when all contacts in both layers are unique. The implications in terms of opportunities for brokerage are straightforward: in the former case, the node has fewer opportunities than in the latter as there are fewer contacts among whom the node can intermediate. However, not all opportunities for brokerage will translate into actual structural holes. This will depend on the variation in links among contacts resulting from the inclusion of an additional layer. When one network layer is combined with another, some of a node's contacts that were unconnected in the former layer may be connected in the latter, thus mitigating the potential of the node to intermediate between contacts. The effective size of the node's neighbourhood in the combined network will ultimately depend on the interplay between variation in number of contacts and variation in number of links between contacts (Latora, Nicosia, and Panzarasa 2013).

Figure 4 shows the distribution of change in brokerage induced by the geo-social network with respect to each of the single-layer networks. For each node, this change is here measured as the difference between the effective size and efficiency of the node's neighbourhood in the composite geo-social network and the node's effective size and efficiency in the single-layer network. When a node has a degree (and effective size) equal to zero in either layer, the node's efficiency is set equal to zero in that layer. Figure 4a shows changes in effective size only within the range $(-50, 50)$. As suggested by the figure, there is an improvement of brokerage potential in the geo-social network over brokerage in the co-location network. When the social layer is also accounted for in the analysis of a node's brokerage position, additional structural holes emerge in the node's neighbourhood, thus amplifying the node's opportunities to intermediate among disconnected others. However, while the majority of nodes can also improve their brokerage positions when the co-location layer is added to the social layer, nonetheless there are some who suffer from a decrease in structural holes. Thus, the co-location network may contribute toward increasing the number of a node's unique contacts, but at the same time may also add new links among some of the node's contacts that would appear as unconnected in the social layer. These mixed effects of the geo-social network on brokerage are even more pronounced when assessed in terms of variation in efficiency. As indicated by Figure 4b, while most nodes seem to secure a more efficient neighbourhood when one layer is combined with the other, there are some who suffer from a loss of efficiency, especially when the co-location layer is added to the social one.

Overall, these results suggest that brokerage may be over- or underestimated when assessed in one single layer. *On the one hand, in the online social network layer many users may appear to be brokers but their influence and intermediation power may be overrated. On the other, many users who appear to have little brokerage power offline may be active brokers when the offline connections are combined with the online ones.*

Neighbourhood heterogeneity

One of the main benefits of brokerage lies in the advantage that individuals acquire through novel recombina-

(a) Variation in effective size
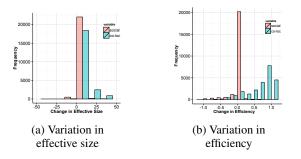
(b) Variation in efficiency

Figure 4: Variation in effective size and efficiency between the multilayer geo-social network and single-layer network

tions of non-redundant information (Burt 1992; Lin 2001; Stovel and Shaw 2012). Brokers are therefore expected to leverage on the diversity of their contacts in order to intensify the competitive advantage they can derive from their structural positions. We therefore test whether in our data set brokerage positions are associated with heterogeneity of local neighbourhoods. In particular, we examine whether this association between brokerage and heterogeneity can be detected in each of the single-layer networks, and whether there is a variation in the association when the two layers are combined.

To this end, for each node we compute the average cosine similarity between all unconnected pairs of the node's contacts. In turn, similarity between unconnected contacts is here assessed by using the frequency distribution vector of the categories of places that each user has visited. This vector is assumed to be representative of the user's personal preference for categories of places visited. There are nine top-level categories of places in Foursquare: "Professional & Other Places", "Shops & Services", "Travel & Transport", "Food", "Nightlife Spots", "Arts & Entertainment", "Colleges & Universities", "Outdoors & Recreation", and "Residences". We build a nine-cell numeric vector, in which each cell is representative of a category and contains the frequency of visits made to that category by each user within the period covered by our data set. For each node, and for each pair of users in the node's neighbourhood, we then compare the two corresponding vectors, and obtain an average score of similarity associated with the focal node's neighbourhood.

Figure 5 shows, for each node, the relationship between the effective size of the node's neighbourhood and the node's average neighbourhood similarity. In agreement with the literature (Burt 1992; Granovetter 1973), Figure 5a suggests that users with many opportunities for brokerage in their social networks tend to belong to a heterogeneous neighbourhood, while the majority of users that are embedded in socially cohesive networks (i.e., with low effective size) tend to be connected with similar others. Findings on the co-location layer are mixed, even though users with the highest brokerage potential are still associated with relatively high diversity in their local neighbourhood (Figure 5b). Finally, when both layers are combined into the geo-social network, neighbourhood heterogeneity is associated with brokerage

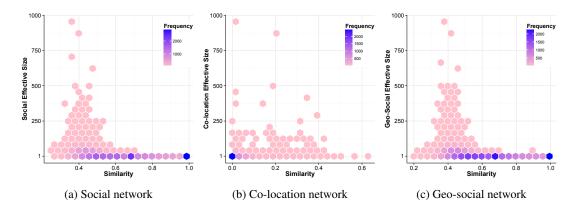| (a) Social network | (b) Co-location network | (c) Geo-social network |

Figure 5: Effective size of nodes in different layers as compared to their average neighbourhood cosine similarity

opportunities (Figure 5c). All users with an effective size of more than 250 have an average neighbourhood similarity of no more than 0.6. Thus, the geo-social network retains and reinforces the pattern observed in the other two layers.

In conclusion, the analysis of multilayer brokerage strengthens the study of social capital by uncovering sources of information benefits that would otherwise remain hidden in a single-layer network. For instance, findings indicate that the potential for novel recombinations of non-overlapping information would remain undetected if only the co-location layer were taken into account. Only when this layer is combined with the online social network layer can the heterogeneity of a node's unconnected contacts be fully captured and the value the node can extract from structural holes be properly assessed.

Discussion and conclusions

In this work we have analysed the interplay between online and offline brokerage for a group of Foursquare users in London. Through the construction of a multilayer network, we were able to integrate the connections that users forged both in the online and the offline layers, and we then detected the opportunities for brokerage originating from the two combined layers. Results indicated that there is a trade-off between brokerage positions online and offline. While most users acted as brokers either online or offline, very few users could secure substantial brokerage positions both online and offline. Moreover, those who succeeded in expanding their social circles in one layer were also capable of maximising the efficiency of their networks by securing connections to non-redundant others. However, such brokerage positions assessed within a single layer were found to be subject to problems of over- and underestimation. Individuals who appeared to benefit from many brokerage opportunities when assessed only within the online layer were in fact embedded in socially cohesive networks when the co-location layer was also taken into account. Conversely, individuals with relatively limited potential for brokering between others within the offline layer could in fact benefit from many brokerage opportunities when the online layer was combined with the offline one. Thus the two layers interact and jointly augment or mitigate the social capital that

individuals can extract from the relationships they forge both online and offline.

We additionally studied the heterogeneity of preferences expressed by users' contacts, and investigated the association between such heterogeneity and the brokerage positions users held. In agreement with the literature on brokerage and social capital (Burt 1992; Lin 2001; Stovel and Shaw 2012), findings indicated that users with many brokerage opportunities tend to belong to highly heterogeneous neighbuorhoods. They can therefore secure a vision advantage through exposure to a diverse set of contacts from whom they can extract and recombine non-overlapping information. However, when only the offline layer is taken into account, this association between information diversity and brokerage remains partly hidden. It is only when the online social layer is combined with the offline layer that such opportunities for novel recombinations of information can be fully captured, and thus the information advantage that individuals can extract from structural holes surrounding them can be fully detected.

Our findings have important implications for both research and practice. First, our analysis sheds light on the benefits that can be derived from a multilayer network approach to social capital. Most studies of social capital have traditionally been concerned with only one network layer, and as a result most network measures of both closed and open structures have captured nodes' positions originating only from the connections that belong to one layer. However, our study has suggested that layers can interact in complex ways, and can jointly affect the sources of social capital on which individuals will ultimately leverage. Single-layer structures that may appear as closed and rich in third-party relationships may instead transform into open structures, rich in brokerage opportunities, when allowed to combine with other single-layer structures. While there are scenarios in which social capital can be successfully studied within unrelated single network layers, our multilayer approach to brokerage has suggested that social capital can be over- or underestimated when the same individuals can exchange information in multiple contexts and or through multiple forms of communication. Second, our study has practical implications for the design and management of online

social networks. For instance, our findings call for a multi-layer approach to brokerage when intermediary roles are to be created and exploited in online social media in order to control and optimise processes of information diffusion and sharing. Targeting individuals who appear to be brokers in one layer may ultimately be the wrong strategy when these individuals can forge connections also in another layer in which what seemed to be a brokerage opportunity may vanish. Similarly, the potential for innovation that may appear to be hindered by the overlapping information detected in one layer can be greatly enhanced when that layer is allowed to interact with another one.

Our study is not without its limitations, which in turn suggest avenues for future investigation. While we measured brokerage through the effective size and efficiency of a node's neighbourhood, alternative measures could be applied to detect open structures. These include constraint (Burt 1992) and Simmelian brokerage (Latora, Nicosia, and Panzarasa 2013). Different measures of similarity between users could also be used to capture the information benefits of brokerage. These could be based on users' geographic location, demographic attributes, as well as mental attitudes inferred from the content of messages posted online. Finally, our study was based on a data set from a popular location-based social network within a large but limited geographic region. Our analysis could be corroborated and further extended if replicated on a larger-scale data set and across various layers of online and offline networks. Despite these limitations, however, our study paves the way for a better understanding of sources of social capital that originate within and across multiple layers.

References

Burke, M., and Kraut, R. 2013. Using Facebook after losing a job: Differential benefits of strong and weak ties. In *Proceedings of CSCW*.

Burt, R. S. 1992. *Structural Holes: The Social Structure of Competition*. Harvard University Press, Cambridge MA.

Burt, R. S. 2005. *Brokerage and Closure: An Introduction to Social Capital*. Oxford University Press, New York and Oxford.

Cardillo, A.; Zanin, M.; Gomez-Gardenes, J.; Romance, M.; Garcia del Amo, A. J.; and Boccaletti, S. 2013. Modeling the multi-layer nature of the European air transport network: Resilience and passengers re-scheduling under random failures. *The European Physical Journal Special Topics* 215(1):23–33.

Cook, K., and Emerson, R. 1978. Power, equity and commitment in exchange networks. *American Sociological Review* 43:712–739.

Cranshaw, J.; Toch, E.; Hong, J.; Kittur, A.; and Sadeh, N. 2010. Bridging the gap between physical location and online social networks. In *Proceedings of UbiComp*.

Eagle, N.; Macy, M.; and Claxton, R. 2010. Network diversity and economic development. *Science* 328(5981):1029–1031.

Ellison, N. B.; Vitak, J.; Gray, R.; and Lampe, C. 2014. Cultivating social resources on social network sites: Facebook relationship maintenance behaviors and their role in social capital processes. *Journal of Computer-Mediated Communication* 19(4):855–870.

Ellison, N. B.; Steinfield, C.; and Lampe, C. 2011. Connection strategies: Social capital implications of Facebook-enabled communication practices. *New Media & Society* 20(10):1–20.

Freeman, L. C. 1977. A set of measures of centrality based on betweenness. *Sociometry* 40:35–41.

Granovetter, M. 1973. The strength of weak ties. *American Journal of Sociology* 78:1360–1380.

Hristova, D.; Musolesi, M.; and Mascolo, C. 2014. Keep your friends close and your Facebook friends closer: A multiplex network approach to the analysis of offline and online social ties. In *Proceedings of ICWSM*.

Kivelä, M.; Arenas, A.; Barthelemy, M.; Gleeson, J. P.; Moreno, Y.; and Porter, M. 2014. Multilayer networks. *Journal of Complex Networks* 2(3):203–271.

Kwak, H.; Lee, C.; Park, H.; and Moon, S. 2010. What is Twitter, a social network or a news media? In *Proceedings of WWW*.

Lampe, C.; Vitak, J.; and Ellison, N. 2013. Users and nonusers: Interactions between levels of adoption and social capital. In *Proceedings of CSCW*.

Latora, V.; Nicosia, V.; and Panzarasa, P. 2013. Social cohesion, structural holes, and a tale of two measures. *Journal of Statistical Physics* 151(3-4):745–764.

Lin, N. 2001. *Social Capital. A Theory of Social Structure and Action*. Cambridge University Press, New York.

Lou, T., and Tang, J. 2013. Mining structural hole spanners through information diffusion in social networks. In *Proceedings of WWW*.

Ravasz, E., and Barabási, A.-L. 2003. Hierarchical organization in complex networks. *Physical Review E* 67.

Steinfield, C.; DiMicco, J. M.; Ellison, N. B.; and Lampe, C. 2009. Bowling online: Social networking and social capital within the organization. In *Proceedings of the Fourth International Conference on Communities and Technologies*, 245–254.

Stovel, K., and Shaw, L. 2012. Brokerage. *Annual Review of Sociology* 38:139–158.

Szell, M.; Lambiotte, R.; and Thurner, S. 2010. Multirelational organization of large-scale social networks in an online world. *Proceedings of the National Academy of Sciences of the United States of America* 107(31):13636–13641.

Vázquez, A.; Pastor-Satorras, R.; and Vespignani, A. 2002. Large-scale topological and dynamical properties of the Internet. *Physical Review E* 65.

Vitak, J.; Ellison, N. B.; and Steinfield, C. 2011. The ties that bond: Re-examining the relationship between facebook use and bonding social capital. In *Proceedings of the 44th Hawaii International Conference on System Sciences*, 1–10.

Watts, D., and Strogatz, S. 1998. Collective dynamics of small-world networks. *Nature* 393:440.

Wellman, B.; Haase, A. Q.; Witte, J.; and Hampton, K. 2001. Does the internet increase, decrease, or supplement social capital? social networks, participation, and community commitment. *American Behavioral Scientist* 45(3):436–455.

White, H. C. 1970. *Chains of Opportunity*. Harvard University Press, Cambridge, MA.