

From Social Ties to Network Processes: Do Tie Definitions Matter?

Ed Fincham¹, Dragan Gašević², Abelardo Pardo³

Abstract

The widespread adoption of digital e-learning environments and other learning technology has provided researchers with ready access to large quantities of data. Much of this data comes from discussion forums and has been studied with analytical methods drawn from social network analysis. However, within this large body of research there exists considerable variation in the definition of what constitutes a social tie, and the consequences of this choice are rarely described or examined. This paper presents findings from two distinct learning environments regarding different social tie extraction methods and their influence on the structural and statistical properties of the induced networks, and the association between measures of centrality and academic performance. Our findings indicate that social tie definitions play an important role in shaping the results of our analyses. The primary purpose of this paper is to raise awareness of the consequences that such methodological choices may have, and to promote transparency in future research.

Notes for Practice

- Social network analysis has been one of the most commonly applied methods within learning analytics. However, many of the common constructs and tools these methodologies employ have not been subjected to robust validation. Such concerns pertain to construct validity: namely, does a metric actually measure what it purports to measure?
- In this study, we find that different social tie extraction methods influence the structural and statistical properties of the induced networks, as well as the associations between centrality measures and academic performance.
- Our results emphasize not only the importance of transparency in the choice of tie definition, but also the importance of providing a justification for that choice. Given the impact that tie definitions may have, we advise that practitioners investigate a number of options to ascertain the extent to which such methodological choices can bias their results.

Keywords

Social network analysis, discussion forum, tie extraction, construct validity, ERGM, MOOC, academic achievement

Submitted: 06.11.17 — **Accepted:** 06.02.18 — **Published:** 05.08.18

Corresponding author ¹Email: ed.fincham@ed.ac.uk Address: Informatics Forum, University of Edinburgh, 10 Crichton Street, Edinburgh, EH8 9AB, UK ORCID ID: 0000-0002-5558-8126

²Email: Dragan.Gasevic@monash.edu Address: Clayton Sch of Info Technology, Monash University, 19 Ancora Imparo Way, Learning and Teaching Building, 3800 Clayton, Australia ORCID ID: 0000-0001-9265-1908

³Email: Abelardo.Pardo@unisa.edu.au Address: Information Technology Engineering and the Environment Divisional Office, University of South Australia, GPO Box 2471, Adelaide, Australia ORCID ID: 0000-0002-6857-0582

1. Introduction

Research into learning analytics has garnered much attention for its potential impact upon a number of central issues in education. For instance, the identification of learning strategies (Jovanović, Gašević, Dawson, Pardo, & Mirriahi, 2017), the prediction of academic success (Gašević, Dawson, Rogers, & Gasevic, 2016), and the provision of personalized feedback at scale (Pardo, Jovanović, Dawson, & Gašević, in press), to name a few. While this field of research promises much, the spectre of validity looms large, and many of the most frequently relied upon measures have not been subjected to robust validation. A pertinent example of this is time-on-task, the estimation of which was rarely discussed in the literature. Instead, researchers would often opt for a heuristic approach such as limiting session activity to a defined time period (Ba-Omar, Petrounias, & Anwar, 2007; Munk & Drlík, 2011). However, there was little consideration of the consequences such

estimation heuristics had on the results of the final predictive model. To address this oversight, Kovanović et al. (2015) investigated how different time-on-task estimation methods affected predictive models of learner outcomes. Across diverse learning contexts, the authors found that estimation methods play an important role in shaping the final study results.

Concerns regarding validity ultimately relate to the extent to which a metric actually measures what it purports to measure. This is known as construct validity (Messick, 1995) and is highly relevant to learning analytics methods, particularly in the context of this study on social network analysis (SNA).

SNA has been one of the most commonly applied methods within learning analytics (Cela, Sicilia, & Sánchez, 2015; Dawson, Gašević, Siemens, & Joksimović, 2014). While SNA can offer insight into the types of relationships and interactions that occur between individuals, groups, and communities, little research has considered the validity of findings derived from common SNA methods. For instance, although Batool and Niazi (2014) assessed the construct validity of centrality metrics in complex networks, studies such as this are the exception rather than the rule. Existing research has paid little attention to the validity of a number of common SNA constructs and, in particular, the impact of various tie definitions on these constructs remains largely overlooked.

Following Messick (1995), validity may be deconstructed into a number of different aspects, including structural validity, or the fidelity of the scoring structure to the structure of the construct itself (Loevinger, 1957); generalizability, or extent to which score properties generalize to and across populations and settings (Cook & Campbell, 1979; Shulman, 1970); and external validity, such as supportive or dissuasive evidence arising from related constructs.

The present study investigates the construct validity of a number of social tie definitions in the context of online discussion forums. These definitions seek to represent the relationships formed between individuals on the basis of interactions and mutual participation within threads. That is, we assess how variations in tie definition result in different characterizations of these relationships. In doing so, we investigate three aspects of construct validity. Structural validity is assessed by comparing tie definitions on the basis of the structural features of the derived networks, and by using statistical models to compare the statistical properties of these networks. External validity is evaluated by investigating how measures of network centrality are associated with academic performance, and generalizability is assessed by pursuing the foregoing analysis in two distinct learning settings.

2. Literature Review

2.1. SNA and Discussion Forums

The analysis of discussion forums, particularly Massive Open Online Course (MOOC) discussion forums, has received considerable attention in recent years. In this body of research, SNA has proven to be a powerful tool in extracting patterns of connections between learners, exploring their relationship with learning, and generating understanding about the underlying relational structure of a community across a variety of contexts. In particular, the emergence of MOOCs has provided ample opportunity for the application of SNA methods (Gašević, Kovanović, Joksimović, & Siemens, 2014; De Laat & Prinsen, 2014). Given the increasing number of students enrolling in MOOCs (Jordan, 2015), SNA has become an increasingly adopted tool for visualizing and extracting interaction patterns from social learning activities (Dowell et al., 2015; Jiang, Fitzhugh, & Warschauer, 2014), as well as for investigating the association of network centrality with constructs such as academic performance (Joksimović et al., 2016; Schreurs, Teplovs, Ferguson, de Laat, & Buckingham Shum, 2013; Skrypnik, Joksimović, Kovanović, Gašević, & Dawson, 2015), sense of community (Dawson, 2008), social presence (Kovanović, Joksimović, Gašević, & Hatala, 2014), and creative potential (Dawson, Tan, & McWilliam, 2011).

There is considerable heterogeneity in how learners interact with the discussion forum. Gillani and colleagues (Gillani & Eynon, 2014; Gillani, Yasseri, Eynon, & Hjorth, 2014), for example, analyzed forum users on the basis of co-participation in the same threads, and found that the coherence of the network mainly depends on a small set of central users. Rather than a close-knit community, forum users may be more accurately characterized as a loosely connected crowd. Poquet and Dawson (2016) explicitly analyzed different user groups, and found that regular users form a denser, more centralized network as they have more opportunities to establish connections. Further work by Boroujeni, Hecking, Hoppe, and Dillenbourg (2017) confirmed that membership of these groups remains stable over time. However, there is more to discussion forums than structure alone; Wise and Cui (2018) distinguished between discussions that were related to course material and those that were not. They found that students who made both content and non-content related posts had a higher passing rate than those who only contributed to one type. Furthermore, those who contributed to content-related threads performed slightly better than those who did not.

The results of these analyses, however, have not always been consistent. For instance, Joksimović et al. (2016) investigated the factors that influence social connections in two instances of a MOOC, offered in English and Spanish, that taught students how to program. Ties were extracted on the basis of direct reply from an online discussion forum. In

examining the association between centrality degree, closeness, betweenness, and academic performance, the authors found weighted degree was significantly associated only in the English offering, while the effects of betweenness and closeness were only found in the Spanish offering. Furthermore, the authors found evidence of performance-based homophily, indicating that learners tend to talk to those in the same performance group as themselves. Another study, by Jiang et al. (2014), also investigated the associations between social centrality and academic performance. Their study was conducted on two MOOCs in algebra and finance, and ties were extracted via co-presence in a thread, that is, on the basis of shared activity and participation in the discussion. While degree and betweenness were positively correlated with academic performance in the algebra course, no significant correlation was found between any centrality measure and academic performance in the finance course. In further contradiction of Joksimović's findings, the authors found that students tend to talk to those in different performance groups than themselves.

The findings of these two studies are largely inconsistent; a discrepancy which may in part be attributed to methodological differences. For instance, Joksimović et al. (2016), hypothesize that the association between academic assessment and network centrality measures was moderated by the presence of Simmelian ties (Krackhardt, 1998, 1999). In lieu of such a hypothesis, Jiang and colleagues' (2014) methodology did not consider the presence of a Simmelian influence. In this case and others, researchers have used different methods to extract social ties yet the effects on those extractions are rarely studied. While in the case of these two studies, the effects of tie extraction methods are studied in connection with the association between network centrality and academic performance, the same methodological oversight may be found regarding other constructs and hypotheses. Moreover, not all research into networked learning has relied upon MOOCs (Cho, Gay, Davidson, & Ingrassia, 2007; Dado & Bodemer, 2017) and, in investigating how social tie extraction methods impact upon the structures and statistical properties of networks, there is scope for a comparison between networks extracted from MOOCs and other, more formal educational contexts.

2.2. Network Processes and Exponential Random Graph Models

Studies that apply SNA methods rely upon mathematical models to describe relationships between variables, to reveal important characteristics, and to identify processes within the social network (Carrington, Scott, & Wasserman, 2005; Goodreau, Kitts, & Morris, 2009). For instance, descriptive models enable us to identify whether or not reciprocity exists within a given network. However, to understand whether or not such processes occur more often than expected if ties were generated randomly, we must rely on statistical models (Goodreau et al., 2009). One commonly proposed method is the Exponential Random Graph Model (ERGM; Joksimović et al., 2016; Morris, Handcock, & Hunter, 2008; DuBois, Butts, & Smyth, 2013).

Introduced by Frank and Strauss (1986) and Wasserman and Pattison (1996), ERGMs belong to a family of probability models that allow for generalizable inferences over the structural foundations of social behavioural patterns within networks (Morris et al., 2008; Robins, Pattison, Kalish, & Lusher 2007). ERGMs treat network ties as random variables, and model the overall network structure through a set of local network processes, such as triadic closure, mutuality, or transitivity (Robins et al., 2007). The model assumes each tie within these processes is conditionally dependent, indicating that "empirical network ties do not form at random, but they self-organize into various patterns arising from underlying social processes" (Wang, Robins, Pattison, & Lazega, 2013, p. 3).

Though ERGMs have long been successfully applied in other fields, their application to the structural analysis of forum networks is relatively novel (e.g., Poquet & Dawson, 2016; Kellogg, Booth, & Oliver, 2014; Joksimović et al., 2016; Zhu et al., 2016). In general, these results have revealed a reciprocal tie effect within networks, and a lack of network centralization beyond a few influential nodes. For instance, Kellogg et al. (2014) used ERGMs to provide a more comprehensive understanding of the dynamic processes underpinning peer support learning in MOOCs. The authors used both descriptive and statistical methods and found a strong and significant reciprocity effect, indicating that students are more likely to aid their peers when there is prior evidence of reciprocity. In a more recent study, Joksimović et al. (2016) utilized ERGMs to determine whether network social dynamics, such as Simmelian ties, have an impact on the predictive power of network positions. The study found that incorporating both descriptive and statistical models allowed for more nuanced and contextually salient inferences about learning within a network. Poquet, Dawson, & Dowell (2017) found that different facilitation, or pedagogical, approaches mediated the extent of reciprocity. That is, while direct reciprocal ties were characteristic of non-facilitated forums, triadic reciprocal ties were more prominent in forums with a high degree of facilitation (that is, instructor involvement).

While statistical models such as ERGMs have facilitated much valuable research and provided considerable insight into network processes, the literature has neglected research into whether and, if so, to what extent network processes and statistical properties are influenced by variations in the tie definitions that underpin them.

2.3. Social Tie Definitions

Research into SNA and, in particular, SNA studies of MOOCs have relied upon a variety of definitions to construct social ties. While some authors (e.g., Joksimović et al., 2016; Kellogg et al., 2014) defined ties on the basis of direct replies, others (e.g., Gillani & Eynon, 2014; and Jiang et al., 2014) have relied on co-presence. As the basis of any SNA analysis, tie definition is crucial and each definition carries with it a set of assumptions about the nature of social interactions. In the literature, this theoretical oversight remains largely unaddressed and studies often establish a tie definition with no explanation or rationale. Even when one is provided, each decision often carries its own shortcomings. For instance, Gruzd & Haythornthwaite (2008) consider three potential tie definitions and note that each makes specific assumptions about the nature of social interactions that may not hold.

In a study of MOOC forums, Wise, Cui, and Jin (2017) investigated the impact of different tie definitions on social network structures and the resultant characteristics at the network, community, and individual node level. While their study found that network properties were characterized by a limited sensitivity to differences in tie definitions, their analysis was limited to descriptive statistics and did not consider the statistical properties of networks, such as the propensity for reciprocity or homophilic ties. Accordingly, there is scope for an investigation into how different tie definitions relate to differences in the statistical properties of networks and the interpretation of such social networks.

Social tie definitions can typically be classified into two distinct types: those that interpret a tie as created when an individual speaks to another, and those that extract ties on the basis of co-participation within a discussion. Perhaps the most prevalent and straightforward of the former is Direct Reply. Under this rubric, a tie is constructed when there is a reply relationship between two nodes in the same thread. For instance, between the starter of a thread and the author of a reply post addressed to it, or between the author of a reply post and the author of a reply to that reply. While this definition has been widely used (see Joksimović et al., 2016; Kellogg et al., 2014) there is no guarantee that users will opt for the correct location and level of post, nor that the platform itself will support a sufficiently rich reply structure. For instance, in one popular platform for online discussions, Piazza, only three levels of post are supported: post, reply, and reply to a reply. If a poster wishes to reply to a “reply to a reply” post, it is logged as another “reply to a reply” post. In building a network, a Direct Reply tie definition would link this new post to the “reply” post rather than the “reply to a reply” post to which it was originally intended to address (see Figure 1). Accordingly, the extent to which the reply structure thus derived reflects the actual relations among learners is open to question.




Figure 1: Limitations of Platform Supported Reply Structures

To address such concerns, Zhu and colleagues (2016) proposed the Star Reply tie definition. While Direct Reply considers multiple levels of replies and defines ties as connections across levels, Star Reply does away with reply structures and considers all posts within a thread as tied to the thread starter. The justification for this is that even if a reply post does not directly address the thread starter, it was made in the context of the discussion originated by the thread starter. While Star Reply emphasizes the thread starter, it fails to distinguish between different levels of replies and does not consider connections formed between posters within the same thread. To address this, Direct Star (Gruzd & Haythornthwaite, 2008) amalgamates the two and defines ties across different levels within a reply structure on the basis of Direct Reply, while also linking posts within a discussion back to the thread starter. However, the methods thus far identified strictly emphasize the act of speaking to another individual. Within a thread, a prospective poster may read much of the existing discussion before penning their own reply. Accordingly, defining ties on the basis of speaking contact alone overlooks the interactions between individuals who do not speak directly, but share an interest and an awareness of each other within the same thread.

Tie definitions of the second type — those based on co-participation within a discussion — seek to address this issue by creating a network of co-presence across nodes. Within such a network, a tie is defined as being present in the same part of the discussion; there is no necessity for direct interaction. Ties are thus created without regard for the reply structure present

in a discussion: connections are formed both between a parent node (post) and its children (replies to it), and between the children themselves. Accordingly, this type of tie definition represents online discussions not as strict hierarchies but as collective conversations.

Within this type of tie definition, a common approach is that of total co-presence, where any two nodes in the same thread are connected, regardless of post type (Gruzd & Haythornthwaite, 2008). While this is often used to map interaction, in the case of large threads it can prove problematic. In the case of small threads, linking all individuals as part of a collective discussion might be reasonable, however, this assumption becomes implausible when the number of replies is very large. One way to address this problem, proposed by Wise and colleagues (2017) is to set a cap on the reasonable number of posts in the same thread to create a measure of limited co-presence. Beyond this threshold, all posts within subthreads are connected to all other posts within that subthread, and the thread starter.

An alternative method of assessing co-participation that has been largely unexplored is that of viewing ties as contained within a moving window. Within a large thread, a prospective poster may only attend to recent posts in framing their reply and so the collective conversation of which they are part is defined as some subset of the overall thread. A moving window, defined as some number of posts, moves sequentially over a thread and, at every step, all posts within the window are connected. This approach ignores the hierarchical structure of a discussion and instead emphasizes each post as being part of a temporally defined, collective discussion. However, there is no *a priori* rationale for choosing one window size over another, and different sizes may lead to a variety of different conclusions.

While each tie definition carries with it a set of assumptions about the nature of social interaction, in the literature little heed was paid to this, and extraction methods may instead be chosen on the basis of expediency, such as whichever structure is readily permitted by the discussion platform. Furthermore, there has been little research into the impact that variations in tie definitions have, both on the statistical properties of networks and on the association between network centrality metrics and academic performance.

2.4. Study Framing

In this study, we examine the effects that different social tie definitions have on the structural and statistical properties of the derived networks. These range from network-level properties such as reciprocity to individual properties, such as the association between metrics of centrality and academic performance. To validate our results, our analysis is applied in two separate contexts: a blended learning environment and a MOOC. The importance of doing so is two-fold: on the one hand, it provides a glimpse into how differences in learning contexts may impact upon social interaction. On the other, the two contexts allow us to assess the construct validity of tie definitions by measuring their impact on the structural and statistical properties of the resultant networks. The paucity of existing research into whether network construction choices determine network properties motivates our two research questions:

RQ1: Do differences in tie formation mechanisms determine the statistical properties of networks across different learning contexts?

RQ2: Do differences in tie formation mechanisms affect the association between social centrality and academic performance across different learning contexts?

3. Method

3.1. Data Sources

This study analyzed forum discussions from two separate courses. The first dataset comes from a flipped classroom, first-year engineering course at an Australian higher education institution offered in 2016. The course, called Introduction to Computer Systems (ICS), lasted 13 weeks and, of a total enrollment of 477 students, 227 students participated in the discussion forum. The flipped classroom design was composed of two elements: a set of online resources intended to be completed in preparation for the plenary session (the lecture), and the re-framing of the plenary session to embrace an active learning design requiring student preparation and participation in collaborative problem-solving tasks (Lage, Platt, & Tregua, 2000; Pardo & Mirriahi, 2017).

The second dataset comes from a course called Code Yourself! (CY), which was delivered on the Coursera platform in 2015. The MOOC was designed to introduce teenagers to computer programming, while covering basic topics in software engineering and computational thinking. The course lasted five weeks and, of a total enrollment of 59,900 students, discussion forum data was available for 1,421 students. The content consisted of lecture videos, quizzes, and peer-assessed

programming projects. If students scored at least 50% in their coursework, they were deemed to have passed, while a distinction was awarded to students receiving a score of 75% or more.

Participation in the discussion forum was not required in either course. Forum activity in ICS consisted of 536 threads, comprised of 1,115 posts. Activity in CY, by contrast, consisted of 774 threads, comprised of 5,950 posts. Summary statistics of the two forums are provided in Table 1.

Table 1: Forum Summary Statistics

Statistic	ICS	CY
Thread Count	536	774
Post Count	1,115	5,950
Average Thread Length	2.08	7.69
Average Sub-thread Length	1.47	1.94
Average Posts per Student	4.91	4.01

The courses were selected to provide two disparate learning contexts for assessing the construct validity of typical SNA methods. In particular, the two courses exhibit drastic differences regarding structure: while ICS is a blended learning environment where the students are likely to have offline connections not captured by the discussion forum, CY is a MOOC where students are likely to interact solely through the discussion forum. This difference is particularly salient since ICS involves offline, collaborative problem solving. Furthermore, ICS is considerably longer, lasting 13 weeks compared to just five for CY. Pedagogy also differs: in ICS, instructors mediate and interact with students in the forum, with the intention of prompting in-depth discussion of the relevant concepts. By contrast, no such mediation exists in CY. These differences and others frame two different contexts and are essential for understanding and interpreting the different social relationships that arise within them.

3.2. Tie Extraction

Ties were extracted using the six tie definitions. Self-ties were excluded in all cases.

Direct Reply Ties (Figure 2.1): The author of each post was connected with the author of its parent post. Concretely, for each thread in the discussion forum and each post within each thread, if a post was either an instructor answer, a student answer, or a level-two post (that is, a reply to a thread starter), a tie was created from the level-two poster to the thread starter. However, if a post was classified as a level-three post (a reply to a level-two post), a tie was created from the level-three poster to the author of the parent level-two post (to which the post was directed).

Star Reply Ties (Figure 2.2): The author of each level-two and level-three post was connected with the author of the thread starting post. To be more concrete, for each thread in the discussion forum, the thread starter was identified and, for each post in the thread, a tie was created from the poster to the thread starter.

Direct Reply & Star Ties (Figure 2.3): Ties defined in both Direct Reply and Star Reply were included but the same tie was never counted more than once. Specifically, for each thread in the discussion forum and each post within each thread, if the post was an instructor answer, a student answer, or a level-two post, a tie was created from the poster to the thread starter. In the case of level-three post, a tie was created from the level-three poster to the level-two poster and, if they were not one and the same person, the thread starter too.

Total Co-presence (Figure 2.4): All authors in the same thread were connected with each other. In this case, ties are considered to be undirected.

Limited Co-presence (Figure 2.5): All users in small threads (<5 replies) were connected to each other with undirected ties; in larger threads users were connected to all other users in their sub-thread and the thread starter only. For threads of five or more posts, if a post was a level-two post (a reply to a thread starter) the level-two poster was connected to the thread starter. However, if level-three posts exist (posts replying to a level-two post), then each level-three post was linked to each other, the level-two post, and the thread starter.

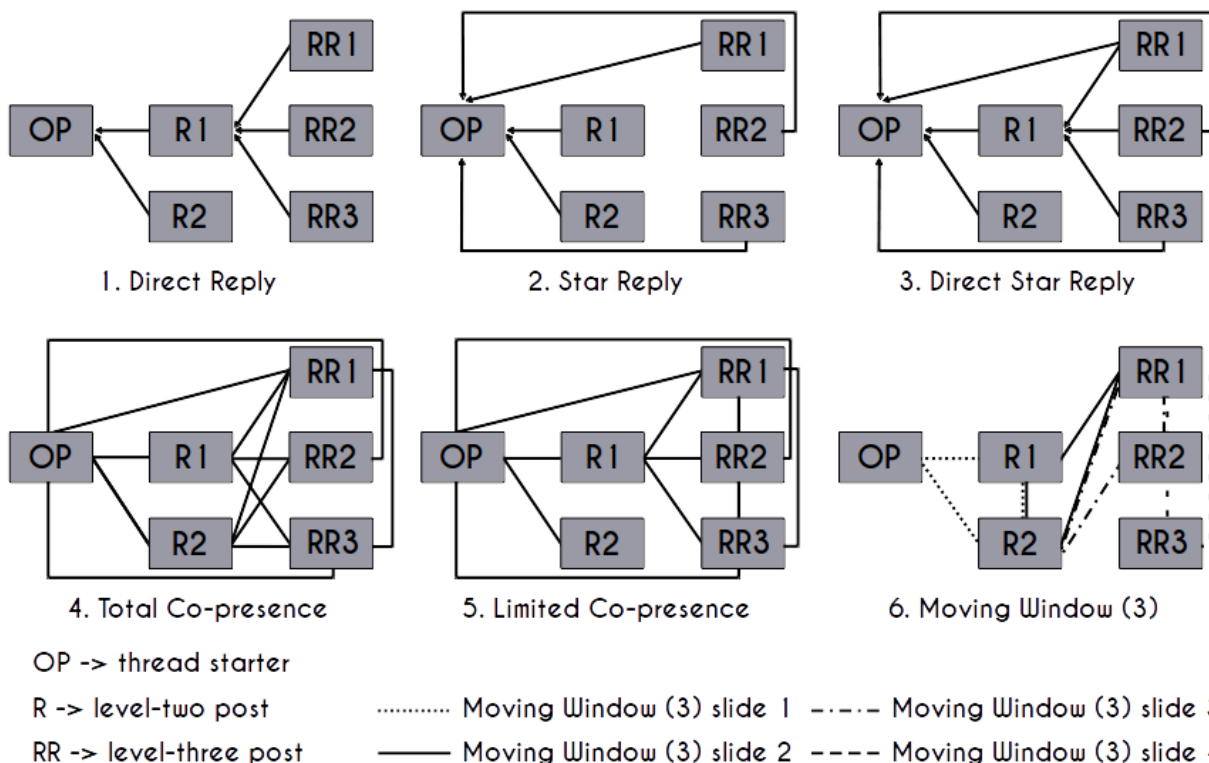


Figure 2: Six Social Tie Definitions

Moving Window (Figure 2.6): All posts within a moving window of size N were connected to each other. Concretely, for each thread in the discussion forum, if the number of posts within a thread was less than N, then a tie was created between each post. Otherwise, starting from the first post, the first N posts are selected, and an undirected tie was created between each post, then the window moved to the second post and the next N posts were selected and ties created. This process repeated until the window reached the end of the thread. In this study we investigated windows of sizes two through five.

Given our definitions, each type of social tie resulted in the number of connections as show in Table 2.

Table 2: ICS and CY Tie Counts

Tie Definition	ICS	CY
Direct Reply	567	3,386
Star Reply	564	2,641
Direct & Star Reply	588	3,761
Total Co-presence	3,757	126,225
Limited Co-presence	589	5,960
Moving Window (5)	969	9,884
Moving Window (4)	863	8,134
Moving Window (3)	727	6,152
Moving Window (2)	520	3,630

4. Analysis

4.1. Social Network Analysis

To address our research questions, networks were extracted for the two courses in accordance with the six tie definitions. Social network analysis was then conducted across all networks in two complementary phases: structural and statistical network analysis.

Our analysis of the networks' structural features relied on some of the most commonly used node-level SNA metrics to characterize centrality, including degree, closeness, betweenness, and eigenvector centrality. Degree centrality captures the local structure surrounding the node, and indicates the number of connections (in- and out-going for reply-based networks) a node has (Freeman, 1979). For this reason, degree is often considered a measure of popularity (Carrington et al., 2005). Closeness centrality measures the distance of a given node to all other nodes in a network (Freeman, 1979) and so can be viewed as a measure of each node's potential to connect with other nodes. Betweenness measures the number of shortest paths between all other nodes that a given node lies on, and so can be viewed as a metric of brokerage or the extent to which a node bridges distinct parts of the network. Finally, eigenvector centrality gives greater prominence to a node the more it is connected to other highly prominent nodes. Accordingly, it can be viewed as a ranked metric of influence.

Additionally, we investigated structural features at the network-level, including density, diameter, and average path length. Within a network, density measures the proportion of actual connections between nodes to all possible connections and so can be viewed as a measure of the extent to which all members of a network are connected to each other (Wasserman, 1994). Diameter measures the maximum eccentricity of any node in a network; that is, the maximum distance between any two nodes. Finally, average path length measures the average number of steps along the shortest paths for all possible pairs of network nodes.

For the statistical analysis of networks, ERGMs were used to reveal a variety of network statistics and investigate network formation processes. In particular, we aimed to investigate the effects of reciprocity, popularity, and transitivity. As a network statistic, reciprocity represents the tendency of students to form mutual ties and group together (Morris et al., 2008). In the context of our datasets and tie definitions, this would indicate whether or not students tend to continue interaction with their peers who replied to their posts. As this metric represents directed loops of length two, it only applies in the case of reply-based networks (that is, networks with directed edges). Popularity was modelled by the *geometrically weighted degree distribution* (*gwidegree* and *gwodegree* in reply-based networks; *gwdegree* in co-participation-based networks, which have undirected edges). *Gwidegree* is a statistic that geometrically discounts the value of incoming ties when the indegrees are summed in the statistic or, more intuitively, captures a popularity effect. *Gwodegree* considers the number of ties an individual sends in the network, and captures structures that result from highly active nodes. Transitivity refers to the extent to which the relation that ties two connected nodes in a network is transitive: that is, the extent to which the friend of my friend is also my friend. This statistic is represented by the *geometrically weighted edgewise shared partner distribution* (*gwesp*).

For each of the networks we consider a variety of models. In the case of reply-based networks (Direct Reply, Star Reply, Direct Star Reply), we examined a model for each of our statistics of interest (*reciprocity*, *gwidegree*, *gwodegree*, and *gwesp*). Similarly, for co-participation-based networks (Total Co-presence, Limited Co-presence, Moving Window 5, 4, 3, and 2), we examined a model for each of our statistics of interest (*gwdegree* and *gwesp*). Models were then analyzed on the basis of goodness-of-fit statistics. Networks were extracted using the *ergm* 3.8.0 (Hunter, Handcock, Butts, Goodreau, & Morris, 2008) R package.

4.2. Regression Analysis

To examine the association between academic performance and our node-level measures of centrality, and so answer our second research question, we conducted a regression analysis. In the case of reply-based networks, we examined six metrics and, in the case of co-participation-based networks, which do not distinguish between in- and out-going ties, we examined four. Since our dependent variable, the course outcome, was measured differently in our two datasets, two distinct approaches were required. In the case of ICS, the dependent variable was continuous, so a linear regression model was fitted. By contrast, the CY course outcome was categorical (obtained certificate). Accordingly, multinomial logistic regression, a method that explains the association between a categorical dependent variable and one or more continuous independent variables (Cramer, 2003), was adopted. To investigate this association, four models were fitted for each dataset. Each model included the dependent variable (course result), one of the centrality measures, and, in order to control for an activity effect, a variable representing an individual's forum post count. Multinomial logistic regressions were performed using the *mlogit* 0.2-4 R package (Croissant, 2013). In the case of ICS, the dependent variable was heavily skewed, and all independent variables across both courses appeared to following a power law distribution; they were therefore log transformed.

5. Results

5.1. Networks Formed by Six Tie Definitions

Descriptive statistics for both datasets reveal clear distinctions between reply-based tie definitions and those based on co-participation. Network centrality metrics were calculated by averaging over the node-level values for each network and centrality type, except for eigenvector centrality. Being more akin to a ranking measure, this network metric was calculated as the sum of differences between each node’s eigenvector centrality and the maximum eigenvector centrality, divided by the maximum possible such value.

Reply-based definitions (Direct Reply, Star Reply, and Direct Star Reply) produced networks of striking similarity regarding centrality metrics at both node and network level (Tables 3 and 4). Regarding co-participation-based definitions, the range of Moving Windows investigated exhibit clear trends across all centrality and network level metrics. As the window size decreased, the number of connections to each was, on average, attenuated, as was the overall density of the network. Furthermore, the distribution of influence across the networks, as measured by eigenvector (network) centrality increased. Compared to the other co-participation-based tie definitions, Total Co-Presence produced dramatically distinct networks in both datasets. By contrast, Limited Co-presence resulted in very similar networks to the Moving Window 3 definition across both courses.

Table 3: ICS Network Descriptive Statistics

Descriptives	DR	SR	DSR	TC	LC	MW(5)	MW(4)	MW(3)	MW(2)
Degree	2.66	2.65	2.76	35.28	5.53	9.10	8.10	6.83	4.88
Closeness	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Betweenness	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Eigenvector	207.31	207.53	207.20	131.59	190.95	188.66	190.49	192.33	195.14
Density	0.013	0.012	0.013	0.166	0.026	0.043	0.038	0.032	0.023
Diameter	13	13	11	6	6	6	6	6	6
Average Path Length	6.66	5.92	6.04	2.04	2.46	2.37	2.41	2.48	2.72

Table 4: CY Network Descriptive Statistics

Descriptives	DR	SR	DSR	TC	LC	MW(5)	MW(4)	MW(3)	MW(2)
Degree	2.38	1.86	2.65	117.66	8.39	13.91	11.45	8.66	5.11
Closeness	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Betweenness	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Eigenvector	1388.19	1380.01	1386.76	1036.18	1368.68	1347.94	1359.07	1370.76	1385.39
Density	0.00	0.00	0.00	0.13	0.01	0.01	0.01	0.01	0.0
Diameter	10	11	9	5	8	6	7	8	13
Average Path Length	4.31	3.98	4.18	2.11	2.923	2.93	3.06	3.29	3.90

5.2. Statistical Network Properties

In both courses under investigation, the three reply-based tie definitions produced networks with largely consistent statistical properties. Estimated coefficients are presented in Table 5. Across all reply-based definitions and datasets, *gwesp* was insignificant, indicating an absence of transitive ties. This consistency across two disparate learning environments is surprising: in the case of ICS, students interact in both the discussion forum and the face-to-face plenary sessions, and accordingly one might expect the derived networks to resemble those emerging from social media, where transitive ties are a *sine qua non* (Nick, Lee, Cunningham, & Brandes, 2013). Regarding reply-based networks, the effect of reciprocity was

significant in all networks across both courses, except for Star Reply in ICS. Across all tie definitions and courses, the effects of popularity and activity, as measured by *gwidegree* and *gwodegree*, were strong, negative, and highly significant, indicating an absence network structures characterized by highly popular or active agents.

In the case of co-participation-based networks, results across both courses were consistent in that all of the investigated network processes were predominantly absent, and the baseline model provided the best fit. The only exception to this pattern was for Moving Window 5 and Moving Window 4 in the CY course. Here the derived networks exhibited evidence of transitive ties. In future research, it is worth investigating to what extent this transitivity is induced by selective mixing, as increasing the likelihood of within category ties provides opportunities for completed triangles within categories, especially when groups are small, as the low density in Table 4 indicates (Goodreau et al., 2009).

Table 5: Reply-based Network Properties

Properties	<u>ICS</u>						<u>CY</u>					
	<u>DR</u>		<u>SR</u>		<u>DSR</u>		<u>DR</u>		<u>SR</u>		<u>DSR</u>	
	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.
<i>Baseline (Edges)</i>	-4.36*	0.04	-4.37*	0.04	-4.33*	0.04	-6.39*	0.02	-6.64*	0.02	-6.64*	0.02
<i>Reciprocity</i>	1.18*	0.31	0.72	0.39	1.10*	0.31	3.26*	0.15	2.73*	0.21	2.73*	0.21
<i>Gwidegree</i>	-3.19*	0.15	-3.21*	0.15	-3.26*	0.14	-2.06*	0.06	-1.38*	0.08	-1.38*	0.07
<i>Gwodegree</i>	-2.96*	0.05	-3.07*	0.14	-3.00*	0.15	-4.91*	0.07	-8.66*	0.09	-8.66*	0.09
<i>Gwesp</i>	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA

Note: * $p < .001$

Table 6a: ICS Co-participation-based Network Properties

Properties	<u>TC</u>		<u>LC</u>		<u>MW(5)</u>		<u>MW(4)</u>		<u>MW(3)</u>		<u>MW(2)</u>	
	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.
<i>Baseline (Edges)</i>	-1.61*	0.02	-3.62*	0.04	-3.10*	0.03	-3.23*	0.03	-3.40*	0.04	-3.75*	0.04
<i>Gwdegree</i>	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA
<i>Gwesp</i>	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA

Note: * $p < .001$

Table 6b: CY Co-participation-based Network Properties

Properties	<u>TC</u>		<u>LC</u>		<u>MW(5)</u>		<u>MW(4)</u>		<u>MW(3)</u>		<u>MW(2)</u>	
	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.
<i>Baseline (Edges)</i>	-1.94*	0.00	-5.13*	0.01	-4.62*	0.01	-4.81*	0.01	-5.09*	0.01	-5.62*	0.01
<i>Gwdegree</i>	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA
<i>Gwesp</i>	NA	NA	NA	NA	8.55*	0.16	3.30*	0.06	NA	NA	NA	NA

Note: * $p < .001$

5.3. Centrality and Academic Achievement

Assessing the association between node-level centrality measures, forum activity, and academic outcomes revealed further differences between the tie definitions. In the case of reply-based networks in the ICS dataset (Table 7a), no centrality metrics exhibited any significant association with course performance. However, for all centrality metrics except for in-degree, activity was significantly and positively associated with course performance (albeit with a small coefficient).

Table 7a: ICS Reply-based Linear Regression Results

	DR		SR		DSR	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
<i>In Degree</i>						
Degree	0.07	0.07	0.06	0.07	0.07	0.07
Activity	0.12	0.06	0.12	0.06	0.12	0.06
<i>Out Degree</i>						
Degree	0.01	0.06	-0.01	0.05	0.01	0.05
Activity	0.16**	0.06	0.17**	0.05	0.16**	0.06
<i>In Closeness</i>						
Closeness	507.25	1719.78	3929.00	5649.00	470.30	1702.00
Activity	0.16**	0.05	0.15**	0.05	0.16**	0.05
<i>Out Closeness</i>						
Closeness	1038.00	2168.00	1190.00	2360.00	1038.00	2150.00
Activity	0.16***	0.05	0.15**	0.05	0.16***	0.05
<i>Betweenness</i>						
Betweenness	0.23	0.84	0.80	1.91	0.34	0.90
Activity	0.16**	0.05	0.15**	0.05	0.16**	0.05
<i>Eigenvector</i>						
Eigenvector	0.98	0.71	0.88	0.75	1.01	0.72
Activity	0.14**	0.05	0.15**	0.05	0.14**	0.05

Note: * $p < .05$, ** $p < 0.01$, *** $p < .001$

In the case of co-participation-based networks in the ICS dataset (Tables 7b & 7c), no centrality metrics except closeness (Total Co-presence) and eigenvector (Moving Window 2) were significant. However, for all tie definitions except for Limited Co-presence (degree), and Moving Windows 4, 3, and 2 (eigenvector), activity was significantly and positively associated with course performance. In the case of Total Co-presence, closeness centrality was significantly and negatively associated with course performance, indicating that as the mean distance between nodes decreased, academic outcomes suffered. However, it should also be noted that these assessments across centrality measures effectively constitute multiple comparisons that have not been controlled for. Given the number of estimated parameters, the occurrence of some significant parameters at the 5% level is likely even under the null hypothesis.

For the CY dataset, in the case of reply-based networks (Table 8a), in-degree centrality was significantly associated with obtaining a certificate of distinction (Direct – $\chi^2(60) = 179.18, p < 0.001$; Star – $\chi^2(54) = 213.43, p < 0.001$; Direct Star – $\chi^2(70) = 178.77, p < 0.001$), but it did not have a significant impact upon the likelihood of obtaining a normal certificate. By contrast, for Direct Reply and Direct Star Reply, out-degree centrality increased the likelihood of obtaining a normal certificate (Direct – $\chi^2(64) = 168.85, p < 0.001$; Direct Star – $\chi^2(72) = 192.82, p < 0.001$) but not a certificate of distinction. Betweenness centrality was significantly and negatively associated with course performance across all reply-based networks (Direct – $\chi^2(756) = 1098.20, p < 0.001$; Star – $\chi^2(360) = 584.01, p < 0.001$; Direct Star – $\chi^2(696) = 1010.50, p < 0.001$). Specifically, increases in betweenness significantly reduced the likelihood of obtaining a certificate with distinction.

In the case of co-participation-based networks in the CY dataset (Tables 8b & 8c), nodes ranked higher by eigenvector centrality were significantly less likely to obtain either certificate in the cases of Total and Limited Co-presence (Total Co-presence – $\chi^2(1384) = 1953.20, p < 0.001$; Limited Co-presence – $\chi^2(1840) = 2265, p < 0.001$) but were significantly more likely to obtain either certificate across all other definitions (Moving Window 5 – $\chi^2(2634) = 2737.70, p = 0.078$; Moving Window 4 – $\chi^2(2654) = 2738.00, p = 0.125$; Moving Window 3 – $\chi^2(2682) = 2750.80, p = 0.173$; Moving Window 2 – $\chi^2(2610) = 2679.30, p = 0.167$). Increases in activity significantly increased the likelihood of obtaining both a distinction and a normal certificate for all metrics except for eigenvector centrality where influence was only significantly associated with a normal certificate for Moving Windows 4 and 3.

Table 7b: ICS Co-participation-based Linear Regression Results

	TC		LC		MW(5)	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
<i>Degree</i>						
Degree	-0.03	0.02	0.06	0.09	-0.03	0.05
Activity	0.19***	0.05	0.12	0.08	0.19**	0.06
<i>Closeness</i>						
Closeness	-257.61*	100.93	-154.02	162.27	-86.19	163.60
Activity	0.21***	0.05	0.19***	0.05	0.18**	0.05
<i>Betweenness</i>						
Betweenness	1.49	8.46	-1.37	1.97	2.48	7.89
Activity	0.16**	0.05	0.17***	0.05	0.15**	0.06
<i>Eigenvector</i>						
Eigenvector	-0.18	0.12	-0.15	0.77	0.13	0.84
Activity	0.18***	0.05	0.18*	0.07	0.16*	0.07

Note: * $p < .05$, ** $p < 0.01$, *** $p < .001$

Table 7c: ICS Co-participation-based Linear Regression Results

	MW(4)		MW(3)		MW(2)	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
<i>Degree</i>						
Degree	-0.04	0.06	-0.03	0.08	0.21	0.11
Activity	0.20**	0.06	0.19*	0.07	0.01	0.09
<i>Closeness</i>						
Closeness	-47.85	159.67	-7.23	154.14	179.98	140.54
Activity	0.17**	0.05	0.17**	0.05	0.13*	0.05
<i>Betweenness</i>						
Betweenness	5.26	8.11	5.77	7.17	4.45	4.96
Activity	0.14*	0.06	0.14*	0.06	0.14**	0.05
<i>Eigenvector</i>						
Eigenvector	0.44	1.03	0.94	1.06	2.18*	0.99
Activity	0.13	0.09	0.10	0.09	0.02	0.07

Note: * $p < .05$, ** $p < 0.01$, *** $p < .001$

Table 8a: CY Reply-based Multinomial Regression Results

	<u>DR</u>		<u>SR</u>		<u>DSR</u>	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
<i>In Degree</i>						
Distinct (deg)	0.29*	0.14	0.76***	0.17	0.31*	0.13
Normal (deg)	-0.10	0.21	0.30	0.24	-0.07	0.19
Distinct (act)	0.91***	0.12	0.79***	0.11	0.91***	0.11
Normal (act)	0.87***	0.17	0.74***	0.16	0.85***	0.17
<i>Out Degree</i>						
Distinct (deg)	0.07	0.11	0.12	0.10	0.11	0.10
Normal (deg)	0.33*	0.15	0.23	0.14	0.31*	0.15
Distinct (act)	1.02***	0.12	1.01***	0.11	0.98***	0.12
Normal (act)	0.53**	0.18	0.65***	0.16	0.54**	0.18
<i>In Closeness</i>						
Distinct (close)	934950.00	497760.00	2273700.00	1302200.00	93170.00	497690.00
Normal (close)	316290.00	758370.00	513790.00	1981100.00	316610.00	758270.00
Distinct (act)	1.04***	0.10	1.05***	0.09	1.04***	0.09
Normal (act)	7.86***	0.14	0.79***	0.14	0.79**	0.14
<i>Out Closeness</i>						
Distinct (close)	109090.00	89692.00	235150.00*	11355.00	109160.00	89657.00
Normal (close)	316240.00*	139270.00	528180.00**	168350.00	316060.00*	139220.00
Distinct (act)	1.00***	0.11	1.00***	0.10	1.00***	0.11
Normal (act)	0.61***	0.16	0.55***	0.16	0.61***	0.16
<i>Betweenness</i>						
Distinct (bet)	-45.15**	14.28	-46.12*	23.52	-45.44**	14.65
Normal (bet)	-113.42	74.36	-105.08	107.66	-146.91	93.71
Distinct (act)	1.18***	0.10	1.12***	0.10	1.18***	0.10
Normal (act)	0.98***	0.16	0.87***	0.15	1.01***	0.16
<i>Eigenvector</i>						
Distinct (eigen)	6.73**	2.46	3.57	2.19	6.41**	2.40
Normal (eigen)	-0.68	3.80	0.18	3.21	-0.20	3.66
Distinct (act)	0.87***	0.11	0.98***	0.11	0.89***	0.11
Normal (act)	0.87***	0.17	0.82***	0.16	0.85***	0.17

Note: * $p < .05$, ** $p < 0.01$, *** $p < .001$; Reference levels for each analysis was “None” — i.e., student did not obtain a certificate.

Table 8b: CY Co-participation-based Multinomial Regression Results

	<u>TC</u>		<u>LC</u>		<u>MW(5)</u>	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
<i>Degree</i>						
Distinct (deg)	-0.14***	0.04	0.00	0.09	0.05	0.10
Normal (deg)	-0.05	0.06	-0.13	0.14	0.13	0.16
Distinct (act)	1.14***	0.09	1.07***	0.12	1.03***	0.12
Normal (act)	0.81***	0.14	0.91***	0.18	0.70***	0.18
<i>Closeness</i>						
Distinct (close)	35347.39	27880.28	32963.69	28063.95	45736.00	29815.00
Normal (close)	85267.53	65480.71	78347.90	62951.64	91833.00	68420.00
Distinct (act)	1.06***	0.09	1.06***	0.09	1.06***	0.09
Normal (act)	0.77***	0.14	0.77***	0.14	0.77***	0.14
<i>Betweenness</i>						
Distinct (bet)	-34.36*	17.06	-21.25	12.37	-25.68**	8.62
Normal (bet)	-143.51	83.05	-52.66	53.16	-137.71*	60.41
Distinct (act)	1.13***	0.10	1.13***	0.10	1.16***	0.10
Normal (act)	0.99***	0.17	0.90***	0.16	1.13***	0.18
<i>Eigenvector</i>						
Distinct (eigen)	-1.67***	0.30	-3.72***	1.06	13.08***	2.26
Normal (eigen)	-1.25**	0.46	-6.26**	2.39	6.62*	3.10
Distinct (act)	1.18***	0.10	1.16***	0.10	0.31*	0.16
Normal (act)	0.87***	0.14	0.94***	0.14	0.52*	0.23

Note: * $p < .05$, ** $p < 0.01$, *** $p < .001$; Reference levels for each analysis was “None” — i.e., student did not obtain a certificate.

Table 8c: CY Co-participation-based Multinomial Regression Results

	<u>MW(4)</u>		<u>MW(3)</u>		<u>MW(2)</u>	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
<i>Degree</i>						
Distinct (deg)	0.09	0.11	0.19	0.14	0.25	0.18
Normal (deg)	0.13	0.18	0.14	0.21	0.07	0.28
Distinct (act)	1.01***	0.13	0.92***	0.14	0.88***	0.17
Normal (act)	0.70***	0.20	0.69**	0.21	0.75**	0.25
<i>Closeness</i>						
Distinct (close)	46786.00	30046.00	48880.00	30462.00	55103.00	31639.00
Normal (close)	90892.00	67927.00	93322.00	69012.00	104740.00	74237.00
Distinct (act)	1.05***	0.10	1.05***	0.09	1.05***	0.09
Normal (act)	0.77***	0.14	0.76***	0.13	0.76***	0.14
<i>Betweenness</i>						
Distinct (bet)	-23.60**	7.78	-20.05**	7.01	-16.36**	5.89
Normal (bet)	-147.05*	60.07	-154.93**	59.19	-129.18**	49.11
Distinct (act)	1.16***	0.10	1.15***	0.10	1.14***	0.10
Normal (act)	1.16***	0.18	1.19***	0.18	1.17***	0.18
<i>Eigenvector</i>						
Distinct (eigen)	17.53***	2.70	22.04***	3.22	19.96***	3.43
Normal (eigen)	8.65*	3.64	10.33*	4.34	10.57*	4.84
Distinct (act)	0.17	0.16	0.10	0.16	0.38*	0.15
Normal (act)	0.49*	0.23	0.49*	0.24	0.53*	0.22

Note: * $p < .05$, ** $p < 0.01$, *** $p < .001$; Reference levels for each analysis was “None” — i.e., student did not obtain a certificate.

6. Discussion

6.1. Structural Network Properties

Descriptive statistics (Section 4.1) for both datasets clearly partitioned social ties according to reply-based and co-participation-based definitions. While these two tie types produced distinct network structures, there remained some notable intra-type variations. In particular, Total Co-presence produced dramatically different networks with high values of degree centrality. This finding, in keeping with Wise et al. (2017), suggests that Total Co-presence should be used with caution due to the disproportionate influence it assigns to large threads. By contrast, in both datasets Limited Co-presence produced networks comparable to Moving Window, in particular Moving Window 3. There are a number of possible explanations for this similarity, such as the predominance of short threads and associated sub-threads in both datasets (see Table 1).

6.2. Statistical Network Properties

In addressing our first research question, there is evidence that across different learning contexts and pedagogies, variations in social tie formation mechanisms may produce different statistical properties in the derived networks. In the case of reply-based networks, we investigated the propensity of networks to form directed loops of length two (that is, reciprocal ties), a popularity and activity effect (whether or not the degree distribution, affected by popular or active agents, contributes to the likelihood of the observed network, captured by *gwidegree* and *gwodegree*), and the extent to which the friend of my friend is also my friend (that is, transitive ties, here captured by *gwesp*). Across all tie definitions, the results were broadly consistent for our statistics of interest, except for Star Reply in the ICS course, where reciprocal ties were notably absent.

In the case of co-participation-based networks, we investigated the propensity of networks to exhibit a popularity effect (whether or not the degree distribution, affected by popular agents, contributes to the likelihood of the observed network, captured by *gwdegree*), and the propensity for transitive ties to form (in the case of undirected edges, this represents the average probability that two neighbours of a vertex are themselves nearest neighbours). The co-participation-based tie definitions we investigated produce a number of transitive (closed) triangles within each thread and sub-thread. However, a transitive effect is only identified in the case of Moving Window 5 and 4. The predominance of “NA” values in Tables 6a and 6b is the result of model degeneracy.

The absence of reciprocal ties in Star Reply (ICS) may in part be explained by both the definition itself, whereby all ties are from responders to a thread starter, and the relatively low student count (on an absolute basis) and thread count of the ICS dataset compared to the CY dataset; since reciprocal ties are not created within Star Reply threads, a larger student and thread count provides greater opportunity for reciprocity to manifest itself. Beyond this exception, the estimated coefficients are broadly consistent on an intra-course basis. While the estimates in CY are quite high, these are in line with those found in existing studies (e.g., Joksimović et al., 2016; Kellogg et al., 2014; Lusher, Koskinen, & Robins, 2012). As Lusher et al. (2012) have argued, it appears that a strong reciprocal effect may be seen as a defining characteristic of interaction in online social networks in general. Furthermore, Lusher et al. (2012) identified such networks as “self-disclosing,” characterized by strong relations between nodes. In these networks, students self-disclose themselves to establish *social presence* (Cobb, 2009; Richardson, Maeda, Lv, & Caskurlu, 2017), creating comfortable learning and knowledge sharing environments. However, the low network level cohesion, evidenced by low density, indicates that students commonly interact with smaller groups (Scott, 2017).

Across all reply-based tie definitions and courses, the effects of popularity and activity (as measured by *gwdegree* and *wdegree*, respectively) were strong, negative, and significant. Such an effect could indicate that within the network, the distribution of popularity and activity were largely homogeneous, rather than being centralized on in- or out-degree. Regarding popularity, this result is consistent with existing studies (e.g., Joksimović et al., 2016; Kellogg et al., 2014; Lusher et al., 2012). In the case of reply-based CY definitions, where reciprocity and a negative popularity effect were particularly strong, the interpretation is quite intuitive: rather than be concentrated in a few individuals, the high propensity of students to engage with each other on a reciprocal basis distributes the effects of popularity over the population.

6.3. Centrality and Academic Achievement

Regarding our second research question, our results indicate that the choice of tie definition can affect the observed association between centrality and academic performance. For instance, for all reply-based networks in the CY dataset, in-degree centrality significantly improved the likelihood of obtaining a certificate of distinction. However, for the very same networks, betweenness centrality significantly decreased the likelihood of obtaining a certificate of distinction.

Given these findings, it is important to assess the underlying assumptions that give rise to such inconsistencies. For instance, while in-degree centrality was significantly and positively associated with obtaining a distinction in all CY reply-based networks, the relation was reverted in the case of Total Co-presence. This may be because the construct being measured differs between the two definitions. In the case of reply-based networks, in-degree centrality indicates the extent of social prominence. However, in the case of Total Co-presence, degree centrality measures the extent of shared interest. While shared interest increases with thread size, social prominence is diluted, which could account for the contrary associations. Similarly, in the CY dataset, eigenvector centrality rank across all tie definitions (except Star Reply) was significantly associated with obtaining a distinction, but the direction of the association depends upon the tie definition. While this association was positive for most definitions, it was negative for Total and Limited Co-presence. This may be because reply-based and, to a limited extent, Moving Window definitions represent the purposeful, direct exchange of information. Total and Limited Co-presence, by contrast, dilute this effect, place inflated importance on large threads, and so provide limited information for assessing influence. These findings give cause to reiterate the warning provided by Wise et al. (2017) that Total Co-presence and, to a lesser extent, Limited Co-presence should be used with caution.

While our results indicate that the choice of tie definitions can affect observed associations, it is important to emphasize that these comparisons have been made at the overall network level, not for specific individuals. It remains unclear to what extent individual centrality metrics are consistent across definitions. This is an interesting avenue of future research, and is an essential consideration when seeking to identify individuals with certain social status.

6.4. Learning and Pedagogical Context

While inconsistent associations between centrality metrics and performance may be in part explained by tie definition, they may also be attributable to differences in learning and pedagogical context. Regarding such contextual factors, the two courses analyzed exhibit important differences. For instance, while ICS is a blended learning environment where students

interact both inside and outside the discussion forum, CY is a MOOC where the forum is students' only point of contact. This could result in MOOC interactions being characterized more by Q&A than in-depth discussions. While a content analysis would have to be conducted to ascertain this in our dataset, the literature provides some evidence in favour of such a hypothesis: for instance, Gillani & Eynon (2014) found that forums harbour crowds, not communities of learners; networks were fragmented and became increasingly so over the duration of the course. Furthermore, although instructor mediation existed in ICS, no such mediation occurred in CY. Poquet et al. (2017) found that different facilitation strategies mediated a reciprocal effect whereby non-facilitated forums were characterized by direct reciprocal ties. Our findings replicate this result: a reciprocal effect was present across all reply-based tie definitions but was almost three times stronger in the case of CY.

Pedagogy may also have impacted student behaviour: for instance, ICS involved a collaborative problem-solving exercise in the plenary session, which could have led learners to participate differently in the discussion forum. Participation may also be affected by contextual factors. For instance, while students created far more posts in the CY forum, students, on average, posted a similar number of times in each course. However, a small proportion of MOOC users even participated in the forum, let alone consistently; not only did a far greater proportion of ICS participate, but they also had offline connections formed over a far longer course period (13 weeks compared to just five).

While our research questions did not directly address the impact of pedagogy and learning context, such factors likely played an important role and should be explicitly addressed before any conclusions or comparisons can be made between courses. However, our results also point to the importance of selecting an appropriate tie definition for a given research goal. For instance, reply-based tie definitions emphasize the purposeful, directed exchange of information between individuals and the derived networks may be useful for identifying roles or influence within a group. Co-participation-based definitions, such as Total and Limited Co-presence, instead treat all ties within a thread as homogenous and focus on identifying shared interest. Our more novel tie definition, Moving Window, has a number of appropriate applications, depending on the course context and window size. For instance, it may be useful within a collaborative learning context where posts within a thread are strongly related to recent posts within the same thread. However, such a structure is highly idealized: not only do learners not always read discussions chronologically (Hewitt, 2003), but in an asynchronous, many-to-many discussion forum, messages may refer to several others appearing far earlier in the chain (Gruzd & Haythornthwaite, 2008). That said, the Moving Window definition emphasizes the temporal structure of threads, a potentially important aspect overlooked by more conventional tie definitions.

7. Implications

This study investigated the construct validity of a number of social tie definitions. Such ties purport to characterize the relationships formed between individuals on the basis of interactions within an online discussion forum. Structural validity was assessed by comparing the definitions on the basis of the structural and statistical properties of the networks they induced from our datasets. Our tie definitions could be categorized into two types (reply-based relationships and co-participation-based relationships, respectively) and while we found broadly consistent structural and statistical properties within these two categories, across category comparisons revealed striking differences.

External validity was assessed by investigating how measures of network centrality were associated with academic performance. While we found that increased social centrality was predominantly associated with opportunities and improved academic outcomes for students, there were some notable exceptions including significant, negative associations. This would indicate that external validity cannot be assured, and that the choice of tie definition does matter.

We also assessed the extent to which our findings generalized by conducting our analyses in two distinct learning settings. Regarding the structural properties of networks, we found reply-based tie definitions produced strikingly similar node-level centrality measures, even in spite of the considerable differences in course context and scale. This was not reflected in the case of co-participation-based networks although, in both contexts, Total Co-presence produced vastly inflated figures. Cross-context similarities were also found within the networks' statistical properties: for reply-based tie definitions, both courses exhibited a significant, negative popularity effect counter-balanced by a significant, positive reciprocal effect. Regarding the association between centrality and academic performance, most metrics enjoyed consistent associations. However, there were deviations, particularly Total Co-presence which, compared to other definitions, in some cases exhibited the opposite association. Though, in many cases, our results generalized across the definitions under study, the departures from this consistency indicate that the validity of SNA methods cannot be assured, and researchers should proceed with caution.

Our results lend support to the argument that researchers should be transparent in their choice of tie definition and, moreover, provide justification for their choice. Given the impact that tie definitions can have, it is advised that researchers

try a number of different methods to ascertain the extent to which such methodological choices can bias their results. On the basis of this study, we recommend future SNA researchers pursue an exploratory comparison of Total Co-presence with a reply-based definition, as this could produce contrasting results and provide clarity on the internal validity of their chosen methods.

Declaration of Conflicting Interest

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) declared no financial support for the research, authorship, and/or publication of this article.

References

- Ba-Omar, H., Petrounias, I., & Anwar, F. (2007). A framework for using web usage mining to personalise e-learning. *Proceedings of the 7th IEEE International Conference on Advanced Learning Technologies (ICALT '07)*, 18–20 July 2007, Niigata, Japan (pp. 937–938). IEEE Computer Society. <http://dx.doi.org/10.1109/ICALT.2007.13>
- Batool, K., & Niazi, M. A. (2014). Towards a methodology for validation of centrality measures in complex networks. *PLOS ONE*, 9(4), e90283. <http://dx.doi.org/10.1371/journal.pone.0090283>
- Boroujeni, M. S., Hecking, T., Hoppe, H. U., & Dillenbourg, P. (2017). Dynamics of MOOC discussion forums. *Proceedings of the 7th International Conference on Learning Analytics and Knowledge (LAK '17)*, 13–17 March 2017, Vancouver, BC, Canada (pp. 128–137). New York: ACM. <http://dx.doi.org/10.1145/3027385.3027391>
- Carrington, P. J., Scott, J., & Wasserman, S. (2005). *Models and methods in social network analysis*. Structural Analysis in the Social Sciences #28. Cambridge, UK: Cambridge University Press.
- Cela, K., Sicilia, M., & Sánchez, S. (2015). Social network analysis in e-learning environments: A preliminary systematic review. *Educational Psychology Review*, 27(1), 219–246. <http://dx.doi.org/10.1007/s10648-014-9276-0>
- Cho, H., Gay, G., Davidson, B., & Ingraffea, A. (2007). Social networks, communication styles, and learning performance in a CSCL community. *Computers & Education*, 49(2), 309–329. <http://dx.doi.org/10.1016/j.compedu.2005.07.003>
- Cobb, S. C. (2009). Social presence and online learning: A current view from a research perspective. *Journal of Interactive Online Learning*, 8(3), 241–254.
- Cook, T. D., & Campbell, D. T. (1979). *Quasi-experimentation: Design and analysis issues for field settings*. Chicago, IL: Rand McNally.
- Cramer, J. S. (2003). The standard multinomial logit model. Chapter 7 in *Logit Models from Economics and Other Fields* (pp. 104–125), Cambridge, UK: Cambridge University Press. <https://doi.org/10.1017/CBO9780511615412.008>
- Croissant, Y., (2013). *mlogit: multinomial logit model*. R Package Version 0.2–4. <http://CRAN.R-project.org/package=mlogit>
- Dado, M., & Bodemer, D. (2017). A review of methodological applications of social network analysis in computer-supported collaborative learning. *Educational Research Review*, 22, 159–180. <http://dx.doi.org/10.1016/j.edurev.2017.08.005>
- Dawson, S. (2008). A study of the relationship between student social networks and sense of community. *Educational Technology & Society*, 11(3), 224–238.
- Dawson, S., Gašević, D., Siemens, G., & Joksimović, S. (2014). Current state and future trends: A citation network analysis of the learning analytics field. *Proceedings of the 4th International Conference on Learning Analytics and Knowledge (LAK '14)*, 24–28 March 2014, Indianapolis, IN, USA (pp. 231–240). New York: ACM. <http://dx.doi.org/10.1145/2567574.2567585>
- Dawson, S., Tan, J. P. L., & McWilliam, E. (2011). Measuring creative potential: Using social network analysis to monitor a learners' creative capacity. *Australasian Journal of Educational Technology*, 27(6), 924–942. <http://dx.doi.org/10.14742/ajet.921>
- De Laat, M., & Prinsen, F. (2014). Social learning analytics: Navigating the changing settings of higher education. *Journal of Research & Practice in Assessment*, 9(4), 51–60.
- Dowell, N., Skrypnik, O., Joksimović, S., Graesser, A. C., Dawson, S., Gašević, D., de Vries, P., Hennis, T., & Kovanović, V. (2015). Modeling learners' social centrality and performance through language and discourse. In O. C. Santos et al. (Eds.), *Proceedings of the 8th International Conference on Educational Data Mining (EDM2015)*, 26–29 June 2015, Madrid, Spain (pp. 250–257). International Educational Data Mining Society.
- DuBois, C., Butts, C., & Smyth, P. (2013). Stochastic blockmodeling of relational event dynamics. *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics (AISTATS 2013)*, 29 April–1 May 2013, Scottsdale, AZ, USA (pp. 238–246).

- Frank, O. D., & Strauss, D. (1986). Markov graphs. *Journal of the American Statistical Association*, 81(395), 832–842. <http://dx.doi.org/10.1080/01621459.1986.10478342>
- Freeman, L. C. (1979). Centrality in social networks conceptual clarification. *Social Networking*, 1(3), 215–239. [http://dx.doi.org/10.1016/0378-8733\(78\)90021-7](http://dx.doi.org/10.1016/0378-8733(78)90021-7)
- Gašević, D., Dawson, S., Rogers, T., & Gasevic, D. (2016). Learning analytics should not promote one size fits all: The effects of instructional conditions in predicting learning success. *The Internet and Higher Education*, 28, 68–84. <http://dx.doi.org/10.1016/j.iheduc.2015.10.002>
- Gašević, D., Kovanović, V., Joksimović, S., & Siemens, G. (2014). Where is research on massive open online courses headed? A data analysis of the MOOC research initiative. *The International Review of Research in Open and Distributed Learning*, 15(5). <http://dx.doi.org/10.19173/irrodl.v15i5.1954>
- Gillani, N., & Eynon, R. (2014). Communication patterns in massively open online courses. *The Internet and Higher Education*, 23, 18–26. <http://dx.doi.org/10.1016/j.iheduc.2014.05.004>
- Gillani, N., Yasseri, T., Eynon R., & Hjorth, I. (2014). Structural limitations of learning in a crowd: Communication vulnerability and information diffusion in MOOCs. *Scientific Reports*, 4, article 6447. <http://dx.doi.org/10.1038/srep06447>
- Goodreau, S., Kitts, J., & Morris, M. (2009). Birds of a feather, or friend of a friend? Using exponential random graph models to investigate adolescent social networks. *Demography*, 46(1), 103–125. <http://dx.doi.org/10.1353/dem.0.0045>
- Gruzd, A., & Haythornthwaite, C. (2008). Automated discovery and analysis of social networks from threaded discussions. *Proceedings of the 18th International Sunbelt Social Network Conference (INSNA 2008)*, 22–27 January 2008, St. Pete Beach, FL, USA. <https://pdfs.semanticscholar.org/e2d3/85cce48b77d567671dc3709774946ec3a309.pdf>
- Hewitt, J. (2003). How habitual online practices affect the development of asynchronous discussion. *Journal of Educational Computing Research*, 28(1), 31–45. <http://dx.doi.org/10.2190/PMG8-A05J-CUHI-DK14>
- Hunter, D. R., Handcock, M. S., Butts, C. T., Goodreau, S. M., & Morris, M. (2008). ergm: A package to fit, simulate and diagnose exponential-family models for networks. *Journal of Statistical Software*, 24(3), 1–29. <http://dx.doi.org/10.18637/jss.v024.i03>
- Jiang, S., Fitzhugh, S. M., & Warschauer, M. (2014). Social positioning and performance in MOOCs. Workshop on Graph-Based Educational Data Mining (G-EDM 2014), 4–7 July 2014, London, UK. *CEUR Workshop Proceedings*, Vol. 1183, paper #8. http://ceur-ws.org/Vol-1183/gedm_paper08.pdf
- Joksimović, S., Dawson, S., Manataki, A., Kovanović, V., Gašević, D., & Friss de Kereki, I. (2016). Translating network position into performance: Importance of centrality in different network configurations. *Proceedings of the 6th International Conference on Learning Analytics and Knowledge (LAK '16)*, 25–29 April 2016, Edinburgh, UK (pp. 314–323). New York: ACM. <http://dx.doi.org/10.1145/2883851.2883928>
- Jovanović, J., Gašević, D., Dawson, S., Pardo, A., & Mirriahi, N. (2017). Learning analytics to unveil learning strategies in a flipped classroom. *The Internet and Higher Education*, 33, 74–85. <http://dx.doi.org/10.1016/j.iheduc.2017.02.001>
- Jordan, K. (2015). Synthesising MOOC completion rates. MoocMoocher. <https://moocmoocher.wordpress.com/2013/02/13/synthesising-mooc-completion-rates/>
- Kellogg, S., Booth, S., & Oliver, K. (2014). A social network perspective on peer supported learning in MOOCs for educators. *The International Review of Research in Open and Distributed Learning*, 15(5). <http://dx.doi.org/10.19173/irrodl.v15i5.1852>
- Kovanović, V., Joksimović, S., Gašević, D., & Hatala, M. (2014). What is the source of social capital? The association between social network position and social presence in communities of inquiry. Workshop on Graph-Based Educational Data Mining (G-EDM 2014), 4–7 July 2014, London, UK. *CEUR Workshop Proceedings*, Vol. 1183, paper #3. http://ceur-ws.org/Vol-1183/gedm_paper03.pdf
- Kovanović, V., Gašević, D., Dawson, S., Joksimović, S., Baker, R. S., & Hatala, M. (2015). Does time-on-task estimation matter? Implications for the validity of learning analytics findings. *Journal of Learning Analytics*, 2(3), 81–110. <http://dx.doi.org/10.18608/jla.2015.23.6>
- Krackhardt, D. (1998). Super strong and sticky. *Power and Influence in Organisations*, 21. <http://dx.doi.org/10.4135/9781483345291.n2>
- Krackhardt, D. (1999). The ties that torture: Simmelian tie analysis in organizations. *Research in the Sociology of Organizations*, 16, 183–210.
- Lage, M. J., Platt, G. J., & Tregua, M. (2000). Inverting the classroom: A gateway to creating an inclusive learning environment. *The Journal of Economic Education*, 31(1), 30–43. <http://dx.doi.org/10.2307/1183338>
- Loevinger, J. (1957). Objective tests as instruments of psychological theory. *Psychological Reports*, 3, 635–694. <http://dx.doi.org/10.2466/pr0.1957.3.3.635>
- Lusher, D., Koskinen, J., & Robins, G. (2012). *Exponential random graph models for social networks: Theory, methods, and applications*. Cambridge, UK: Cambridge University Press.

- Messick, S. (1995). Validity of psychological assessment: Validation of inferences from persons' responses and performances as scientific inquiry into score meaning. *American Psychologist*, 50(9), 741–749. <http://dx.doi.org/10.1037/0003-066X.50.9.741>
- Morris, M., Handcock, M. S., & Hunter, D. R. (2008). Specification of exponential-family random graph models: Terms and computational aspects. *Journal of Statistical Software*, 24(4). <http://dx.doi.org/10.18637/jss.v024.i04>
- Munk, M., & Drlík, M. (2011). Impact of different pre-processing tasks on effective identification of users' behavioral patterns in web-based educational system. *Procedia Computer Science*, 4, 1640–1649. <http://dx.doi.org/10.1016/j.procs.2011.04.177>
- Nick, B., Lee, C.-K., Cunningham, P., & Brandes, U. (2013). Simmelian backbones: Amplifying hidden homophily in Facebook networks. *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM '13)*, 25–28 August 2013, Niagara, ON, Canada (pp. 525–532). New York: ACM. <http://dx.doi.org/10.1145/2492517.2492569>
- Pardo, A., Jovanović, J., Dawson, S., & Gašević, D. (in press). Using learning analytics to scale the provision of personalised feedback. *British Journal of Educational Technology*, 1–11. <http://dx.doi.org/10.1111/bjet.12592>
- Pardo, A., & Mirriahi, N. (2017). Design, deployment and evaluation of a flipped learning first-year engineering course. In C. Reidsema, L. Kavanagh, R. Hadgraft, & N. Smith (Eds.), *The Flipped Classroom: Practice and Practices in Higher Education* (pp. 177–191). Springer. http://dx.doi.org/10.1007/978-981-10-3413-8_11
- Poquet, O., & Dawson, S. (2016). Untangling MOOC learner networks. *Proceedings of the 6th International Conference on Learning Analytics and Knowledge (LAK '16)*, 25–29 April 2016, Edinburgh, UK (pp. 208–212). New York: ACM. <http://dx.doi.org/10.1145/2883851.2883919>
- Poquet, O., Dawson, S., & Dowell, N. (2017). How effective is your facilitation? Group-level analytics of MOOC forums. *Proceedings of the 7th International Conference on Learning Analytics and Knowledge (LAK '17)*, 13–17 March 2017, Vancouver, BC, Canada (pp. 208–217). New York: ACM. <http://dx.doi.org/10.1145/3027385.3027404>
- Richardson, J. C., Maeda, Y., Lv, J., & Caskurlu, S. (2017). Social presence in relation to students' satisfaction and learning in the online environment: A meta-analysis. *Computers in Human Behavior*, 71, 402–417. <http://dx.doi.org/10.1016/j.chb.2017.02.001>
- Robins, G., Pattison, P., Kalish, Y., & Lusher, D. (2007). An introduction to exponential random graph (p^*) models for social networks. *Social Networks*, 29(2), 173–191. <http://dx.doi.org/10.1016/j.socnet.2006.08.002>
- Schreurs, B., Teplovs, C., Ferguson, R., de Laat, M., & Buckingham Shum, S. (2013). Visualizing social learning ties by type and topic: Rationale and concept demonstrator. *Proceedings of the 3rd International Conference on Learning Analytics and Knowledge (LAK '13)*, 8–12 April 2013, Leuven, Belgium (pp. 33–37). New York: ACM. <http://dx.doi.org/10.1145/2460296.2460305>
- Scott, J. (2017). *Social Network Analysis*, 4th ed. London: SAGE Publications.
- Shulman, L. S. (1970). Reconstruction of educational research. *Review of Educational Research*, 40, 371–396. <http://dx.doi.org/10.2307/1169372>
- Skrypnik, O., Joksimović, S., Kovanović, V., Gašević, D., & Dawson, S. (2015). Roles of course facilitators, learners, and technology in the flow of information of a cMOOC. *International Review of Research in Online and Distance Learning*, 16(3), 188–217. <http://dx.doi.org/10.19173/irrodl.v16i3.2170>
- Wang, P., Robins, G., Pattison, P., & Lazega, E. (2013). Exponential random graph models for multilevel networks. *Social Networks*, 35(1), 96–115. <http://dx.doi.org/10.1016/j.socnet.2013.01.004>
- Wasserman, S., & Pattison, P. (1996). Logit models and logistic regressions for social networks: An introduction to Markov graphs and p . *Psychometrika*, 61(3), 401–425. <http://dx.doi.org/10.1007/BF02294547>
- Wasserman, S. (1994). *Social network analysis: Methods and applications*. New York: Cambridge University Press.
- Wise, A. F., & Cui, Y. (2018). Unpacking the relationship between discussion forum participation and learning in MOOCs: Content is key. *Proceedings of the 8th International Conference on Learning Analytics and Knowledge (LAK '18)*, 5–9 March 2018, Sydney, NSW, Australia (pp. 330–339). New York: ACM. <http://dx.doi.org/10.1145/3170358.3170403>
- Wise, A. F., Cui, Y., & Jin, W. Q. (2017). Honing in on social learning networks in MOOC forums: Examining critical network definition decisions. *Proceedings of the 7th International Conference on Learning Analytics and Knowledge (LAK '17)*, 13–17 March 2017, Vancouver, BC, Canada (pp. 383–392). New York: ACM. <http://dx.doi.org/10.1145/3027385.3027446>
- Zhu, M., Bergner, Y., Zhang, Y., Baker, R., Wang, Y., & Paquette, L. (2016). Longitudinal engagement, performance, and social connectivity: A MOOC case study using exponential random graph models. *Proceedings of the 6th International Conference on Learning Analytics and Knowledge (LAK '16)*, 25–29 April 2016, Edinburgh, UK (pp. 223–230). New York: ACM. <http://dx.doi.org/10.1145/2883851.2883934>