Longitudinal Engagement, Performance, and Social Connectivity: a MOOC Case Study Using Exponential Random Graph Models

Mengxiao Zhu
Yoav Bergner
Yan Zhang
Educational Testing Service
Princeton, NJ 08541
mzhu, ybergner, yzhang002 @ets.org

Ryan Baker
Yuan Wang
Teachers College
Columbia University
New York, NY 10027
ryanshaunbaker @gmail.com
elle.wang@columbia.edu

Luc Paquette
University of Illinois, U-C
Champaign, IL 61820
lpaq@illinois.edu

ABSTRACT
This paper explores a longitudinal approach to combining engagement, performance and social connectivity data from a MOOC using the framework of exponential random graph models (ERGMs). The idea is to model the social network in the discussion forum in a given week not only using performance (assignment scores) and overall engagement (lecture and discussion views) covariates within that week, but also on the same person-level covariates from adjacent previous and subsequent weeks. We find that over all eight weekly sessions, the social networks constructed from the forum interactions are relatively sparse and lack the tendency for preferential attachment. By analyzing data from the second week, we also find that individuals with higher performance scores from current, previous, and future weeks tend to be more connected in the social network. Engagement with lectures had significant but sometimes puzzling effects on social connectivity. However, the relationships between social connectivity, performance, and engagement weakened over time, and results were not stable across weeks.

Keywords
MOOC, network analysis, forum participation, exponential random graph model, ERGM, learning.

1. INTRODUCTION
By its design, the field of learning analytics takes a fine-grained approach to using data in service of understanding learning processes and supporting outcomes. Learning processes are of course contextual, dynamic, social, and highly variable in myriad ways. Thus, from a data-driven standpoint, the integration of multiple streams of data evolving over time (contextual, cognitive, affective, social, etc.) is the ultimate objective. Such integration is immensely challenging, requiring interdisciplinary efforts and high quality data.

This paper explores a longitudinal approach to combining engagement, performance and social connectivity data from a MOOC using the framework of exponential random graph models (ERGMs) [32]. The idea is to model the social network in the discussion forum in a given week not only using performance (quiz scores) and overall engagement (lecture and discussion views) covariates within that week, but also using person-level covariates from adjacent previous and subsequent weeks.

Our models, described in more detail below, address the following questions. First, do engagement with learning content and performance on assessments provide additional information—beyond the endogenous mechanisms of a random network structure—about a learner's connectedness to others in the social network? Second, do time-lagged effects, such as joining the discussion in response to struggling with the previous assignment, relate to performance and engagement counts from the previous week? Addition, rather than substitution, means that present-week effects are controlled for, in the sense of multiple linear regression predictors, in the estimation of past-week effects. Finally we add future-week effects as well. None of these are causal models, and there is clearly no causal mechanism for future week scores to influence social connectedness. A significant positive interaction within this framework, again controlling for present and past scores, would suggest that future outcomes may be related to social connectedness in the present.

We note that the approach here contains many important simplifications. Links or edges between networked participants are dichotomous, that is, they are not weighted by frequency of communications or characterized by other measures of type or strength. An edge either exists or does not exist. Related to this point, but distinct, the content of the forum posts [10] is completely ignored in this analysis. Finally, we acknowledge that our engagement and performance variables are relatively simple distillations of these aspects of the learning experience. The caveats above mean that results presented here should be taken as preliminary and provisional. Our intention, however, is to convey a framework for integrated analysis which can be improved through refinement. Augmenting the model with linguistic content is part of ongoing work.

The organization of the paper is as follows. Section 2 reviews related work. Section 3 describes our dataset. Section 4 introduces the methods for building social networks and for analyzing the network data. Section 5 provides the results and Section 6 the related discussions.

LAK '16, April 25-29, 2016, Edinburgh, United Kingdom
© 2016 ACM. ISBN 978-1-4503-4190-5/16/04…$15.00

2. RELATED WORK

2.1 Forums as Content Resources

Discussion forums in a MOOC may be alternately viewed as a (dynamic) content resource and as a social resource. These characterizations are not mutually exclusive, but they emphasize different aspects and suggest different analytic approaches. Insofar as students predominantly view posts authored by others, for example as a means of finding help with homework or information about course policies, the forums are a content resource. Of course, there would be no content if no one posted to the forum, but it is possible to count posts in addition to views as simply two dimensions of forum usage.

A number of studies have examined the relationship between performance (grades), attitudes, and discussion forum usage in terms of view and post counts [24, 26, 31]. The methods involved in such analyses are typically correlational using aggregated values over the duration course. Nevertheless operationalization of both predictor and outcome variables poses significant challenges, especially when applying similar methods to massive open online courses (MOOCs), where the population of users is so diverse [6].

2.1.1 Longitudinal effects

Week-by-week analysis of forum content, coded according to the Collaborative Learning Conversation Skill Taxonomy [38], was presented in [39], in order to study the evolution of content over time. Models incorporating student change from early to late course stages in their (volumetric) use of the forums for homework help were examined in [6].

2.2 Forums as Social Resources

The allure of discussion forums as a focus of study on online learning lies, at least partly, in connection to theories of learning such as social cognitive theory [4], social constructivism [42], and connectivism [35]. Insofar as forums are viewed as a social resource, the quantity of views or posts or content tags may take a back seat to the relational aspect of replying to or receiving a reply from another learner. Analyses that emphasize the social aspect of discussion forums thus commonly apply some type of social network analysis (SNA) to the communication network of learners as defined through their forum interactions.

Information sharing and the flow of content in social networks of learners were studied in both distance and on-campus learning over a decade ago [3, 9, 18], before the proliferation of learning management systems and MOOC platforms [11]. As discussed in a recent review of SNA applications in discussion forums [30], the most common methods since then have been visualization of networks; characterization of users, for example, by centrality; and community detection. Recent MOOC applications have examined super-poster behavior [21], structure of networks in a teacher-oriented MOOC [25], affinity characteristics of communities [7], impact of community affiliation on performance [8], as well as the distinction between communities and crowds [16].

2.2.1 Generative network models

While SNA has often been used to extract descriptive statistics about networks and individuals, it is also possible to model the network itself as an outcome of a random generating process. Such approaches allow one to test hypotheses of reciprocity or homophily in a network, as was done using exponential random graph models (ERGMs) in [2] and [25]. Alternate methods, similar in spirit, include latent space models [19] and stochastic blockmodels [36]. Hierarchical versions of these were applied to teacher networks in [40].

2.3 Nonsocial networks

It should be noted that as a general method for learning analytics, SNA has also been applied to networks where the entities are not persons but rather concepts or states in a state space navigated by learners. The use of the word "social" in this

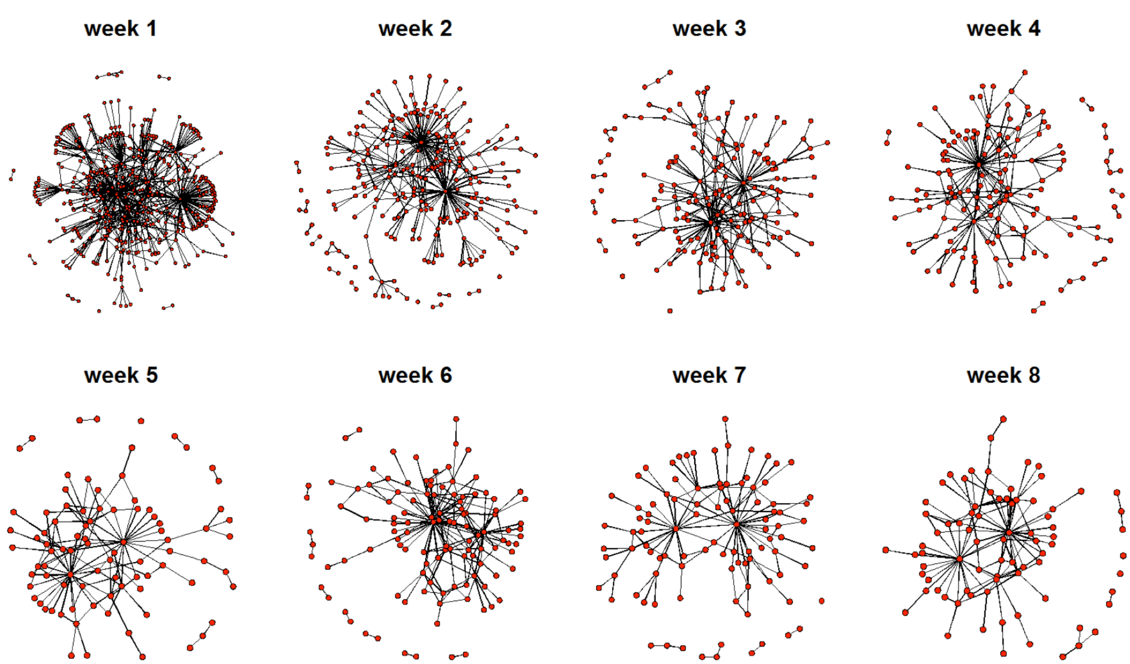

Figure 1. Weekly forum interaction social networks.

context, although habitual because of SNA, has the potential to create confusion. Community detection applied to a state-space graph in a logic tutoring system, for example, was used to identify student pitfalls or misconceptions [13]. Network analysis was also applied to eye-tracking movements in mathematics problem solving in [48].

3. DATA SET

The dataset used in this study was collected from a Coursera course “Big Data in Education” (BDE) offered by Teachers College, Columbia University in Fall 2013 [44]. In this course, students learned and applied educational data mining and learning analytics methods. The course spanned eight weekly sessions, each consisting of five to seven lecture videos, readings, and a quiz, which requires the students to do data analysis and submit the results. All weekly assignments were automatically graded and accounted for 100% of the final course grade (70% was the threshold for earning a certificate). Students were encouraged to use the online forums throughout the duration of the course, and this was the principal way in which instructors and TAs interacted with learners.

BDE had over 48,000 enrollees during the duration of the course, with a small portion *actively* participating. A total of 1,380 students completed at least one assignment, and 638 students received a certificate. This study concerns the relationship between forum participation and other learning activities. Thus, we restrict our attention to the users who posted or commented at least once in the forum. This yielded a sample of 770 individuals, including students, the instructor, and teaching assistants. Among this sample, 440 students (57%) completed at least one assignment, and 155 (20%) earned a certificate.

Following the organization of the course, and in order to enable comparisons between weeks, we constructed separate datasets for each week unit. The cutoff times for assigning events to weeks were determined by the release of new units by the instructor, rather than by strict calendar intervals. Variables extracted from the log data included the number of posts initiated/commented/replied, the number of posts viewed, the number of video lectures downloaded or viewed, and scores on the weekly quizzes. Although the quizzes allowed multiple attempts, we computed scores based on whether answers were correct on the first attempt. As found by [5], this type of scoring resulted in score distributions with more variance and reduced ceiling effects than scoring based on eventual correctness.

4. METHODS

4.1 Constructing social networks from the forum posts

As is typical, a forum thread in BDE was initiated by an original poster (OP). Participants could then “reply” to the OP or “comment” on previous replies. Although the MOOC data logs distinguish these types of actions, students did not consistently use them differently. In other words, sometimes a reply was really a comment to a specific prior reply, and sometimes a comment was really an open reply to the OP. Thus we chose not to distinguish these two actions after all.

The choice arises as to whether a reply should connect its poster to everyone who posted previously or only to the OP. This has been handled differently in different studies, and depending on whether the links are directed or undirected [7, 27]. All individuals involved in a post are frequently connected to each other in collaboration networks or co-authorship networks [1]. The resulting structure will be a fully-connected clique [47]. However, we did not adopt this method with the consideration that these relations, especially those among the repliers or commenters the forum posts, might be much weaker than the relations among researchers who co-authored a paper together.

To construct our social networks, we connected only the OP and the subsequent repliers. As a result, each thread is represented by a star in the network, with the OP as the node in the middle and all others in the periphery. However, there were multiple threads in the network for each week. Eight undirected social networks were thus constructed. The nodes in the networks are individuals and the links indicating the replying or commenting relations.

The visualizations of the weekly forum interaction social networks are shown in Figure 1 (generated using the *gplot* function in the R package Statnet [17]). All networks show a relatively big connected component and several smaller components in the periphery. The visualizations show the network shrinking over time, which is confirmed by basic summary statistics. We plot the size, density, size of the largest connected component, and the maximum degree of the nodes of the networks in Figure 2. These basic network statistics [45] are defined as follows. The size of a network is defined as the total number of nodes in the network. The density of the undirected network is the number of links divided by the maximum possible number, which for an undirected graph with n nodes is

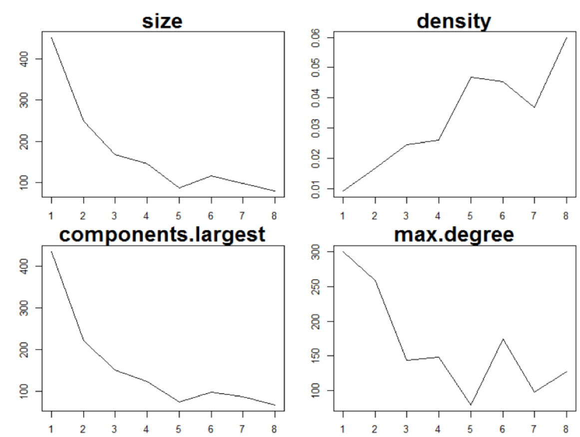

Figure 2. Basic network statistics for forum interaction social networks in BDE.

$n(n-1)/2$. The connected component is defined as the subset of a network, in which any two nodes are connected directly or through other nodes, and the largest connected component is the one with the largest number of nodes. The degree of a node is the count of number of links connected to that node, and the maximum degree of the nodes in a network is the degree of the most connected node.

The size of the networks decreased from the maximum of 450 in the first week to 80 in the eighth week. For each network, the largest component consisted of 85% or more of the nodes and thus had the same longitudinal trend as the full network. The maximum node degree for each network was achieved by either the instructor or the head community TA, again following the same decreasing trend. On the other hand, the densities of the networks increased over time, showing that forum participants tend to have closer relations as time went on. Week 5 seemed to depart a little in each of these trends, possibly due to the timing of the Thanksgiving holiday in the U.S.

4.2 Exponential Random Graph Models (ERGMs)

To relate network measures to individual attributes, one might consider using a standard regression approach. However, when individuals are embedded in the same network, their network attributes (edge/link statistics) are interdependent. These dependencies violate the independent observation assumptions of regression models. Rather than limit ourselves to correlations between descriptive network measures and person (node) attributes, we study the relation between network structures and node attributes in a framework that explicitly accounts for the dependencies among nodes and links. Exponential random graph models (ERGMs) [32], also known as p^* models, were developed specifically for analyzing social network data and testing hypotheses on both network structures and on the interactions of node attributes with network structures. ERGMs have been widely used in applications ranging from the structure of adolescent peer influence [41] to social connections in virtual worlds [22].

The basic idea of ERGMs is to build a stochastic model that captures the generative features of an observed social network. By analogy with regression models, dependent variables are the nodes and links in the observed network, and the independent variables are summary statistics for various network structural features. Node attribute effects are analogous to interaction terms, that is, of node attributes and structural features. Some examples of the structural features include links, triangles, stars, or links between nodes sharing a particular attribute. The summary statistic for links is just the density of a network.

Mathematically, the general ERGMs are a class of stochastic models that share the following general form [46].

$$\mathrm{P(Y = y)} = \frac{1}{\mathrm{k(\theta)}} \exp\left(\theta^{\mathrm{T}} \mathrm{g(y)}\right)$$

where Y is a random variable representing the network and y is the specific observed network. The state space of Y is the collection of all possible networks with the same number of nodes as the observed network, y. $\mathrm{g(y)}$ is a vector of network statistics, θ is the vector of corresponding coefficients, and $k(\theta)$ is a normalizing constant. It is calculated by summing up $\exp\left(\theta^{\mathrm{T}} \mathrm{g(y)}\right)$ over the space of all possible networks.

General ERGMs do not impose any dependence assumptions. Rather, one constructs different models by specifying these dependence assumptions. For instance, in the simplest case, it can be assumed that all links in the network are equivalent and thus that the probability of an observed network depends only on the number (equivalently, density) of links in the network. This is known as the Bernoulli Model or the Erdős–Rényi Model [14]. It is not particularly realistic for most social networks.

More complex dependency assumptions include dyadic independence models (p_1) [20], for which reciprocated edges are included as a structural feature or, further still, p_2 models, which add conditional dependence on node-level attributes [12, 28]. In Markov random graphs [15], it is assumed that an edge between two nodes *i* and *j* depends on any other possible edges involving *i or j*. Curved exponential models have introduced nonlinear functions of the θ parameters to better capture the structural features for social networks [23, 34]. Simulation and estimation methods were also developed for ERGMs with the Markov dependence assumption [29, 33, 46].

The output results of ERGMs (θ coefficient estimates) may also be interpreted by analogy with regression models. A significant positive value for a coefficient corresponding to a structural feature, for example, triangles, indicates that this feature occurs more than would be expected by chance. For node attributes (covariates), the interpretation of the coefficient is a bit different. In that case, the conditional log-odds of an edge connecting two nodes *i* and *j* is understood to be increased by the product of the coefficient and the *sum of covariate values for the two nodes*.

4.3 Variables from the MOOC data

Social Networks

For this study, we are interested in the relationship between MOOC forum social links and course participants' behavior and performance in current and adjacent weeks. As described in Section 4.1, there are eight social networks, one for each week.

Nodal attributes

Each individual has four attributes extracted from each week's data, including assignment score (*score*, the only performance attribute), the number of posts initiated/commented/replied (*posts*), the number of posts viewed (*post views*), and the number of video lectures downloaded or played (*lectures*). Counts of posts, post views, and lectures constitute our simplified engagement measures. We are also interested in how these attribute measures from adjacent weeks are related to the current week's social networks. Thus, the dataset for week 4 includes the social network from week 4's forum data and engagement and performance attributes from weeks 3, 4, and 5.

5. RESULTS AND DISCUSSION

5.1 Hierarchical analysis method for ERGM on Week 2

With the social networks and variables as described in Section 4.3, we build a hierarchical sequence of ERGMs. We use the R package Statnet [17] for the analysis. The variables, blocks, and coefficient estimates using the Week 2 data are shown in Table 1.

Model 1 serves as the baseline model, which contains two structural features of the forum network, *edge* and *alternating k-stars* [23]. The first feature accounts for the density of the network, and the second captures the tendency for hubs. As shown in Table 1, the negative value for *edge* (-4.09) indicates that the density of the network is lower than would occur by chance. A negative value for *alternating k-stars* (-0.99) indicates that hub frequencies are also lower than would occur by chance. The phenomenon in which higher degree nodes attract more links is known as preferential attachment [1]. In our dataset, even though we do see some nodes with high degree, the overall generating mechanism does not lean towards adding more links to the high degree nodes (after controlling the density of the network).

For Model 2, we add four nodal attributes, two related to performance and two related to engagement. *Score* captures the tendency for individuals with high assignment scores to be more active in the social network. *Score difference* captures the tendency for two linked individuals to have similar assignment scores. With this term, we hope to test the hypothesis that social interactions may be more likely between similarly performing students. Our engagement attributes account for increasing connectedness by individuals who view many posts or lectures.

The results from Model 2 are a significant positive effect for assignment scores on connectedness in the forum network and a significant negative effect for lecture views/downloads. The score difference and post view effects were not significant. The negative effect of lecture engagement is certainly a bit puzzling, especially if one thinks of this in terms of high lecture viewing being associated with low connectedness. However, the mathematically equivalent converse relation—high connectedness associated with low lecture viewing—is actually plausible if "expert" individuals who already know the content do a lot of replying.

Table 1: ERGMs for Week 2

	Model 1	**Model 2**	**Model 3**	**Model 4**
Effect	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***
density (edge)	-4.09* (0.09)	-4.06* (0.15)	-4.12* (0.30)	-3.95* (0.30)
tendency for hubs (alt. k-stars)	-0.99* (0.19)	-0.93* (0.19)	-0.76* (0.20)	-0.63* (0.21)
score (current week)		0.19* (0.10)	0.05 (0.11)	-0.43* (0.18)
score difference (current week)		0.18 (0.13)	0.18 (0.13)	0.17 (0.13)
post reads (current week)		-0.14 (0.34)	-1.36* (0.47)	-1.61* (0.49)
lectures (current week)		-0.49* (0.20)	-0.62* (0.28)	-1.26* (0.37)
Score (previous week)			0.28* (0.13)	0.32* (0.13)
posts reads (previous week)			2.01* (0.50)	2.39* (0.53)
post (previous week)			-1.20* (0.42)	-1.21* (0.42)
lectures (previous week)			-0.38 (0.64)	-1.05 (0.67)
score (next week)				0.38* (0.15)
lectures (next week)				0.62* (0.24)

* indicates $p < 0.05$; † indicates $p < 0.1$. Model convergence information available upon request.

In Model 3, we added attributes from the previous week (week 1 for the week 2 model). It turned out that individuals with higher assignment scores from the previous week tended to have more links in the social network in the current week. Individuals who viewed more posts from the previous week also tended to have more links in the current week, while individuals who posted more in the previous week tended to have fewer links in the current week's social network (controlling for other effects). The lecture views/downloads from the previous week did not have a significant effect on the current week's relations.

Finally, in Model 4 (the full model with all variables), we explore whether forum social relations correlate with future learning behavior. The effects for both assignment scores and the lecture views/downloads were positive and significant, indicating that individuals with more social links in the current week tend to have higher assignment scores and more lecture views/downloads in the next week.

5.2 Beyond Week 2

Results from Week 2 offer many interesting findings, especially the between-week interactions of social connectedness and learner performance and engagement attributes. We ran the full set of models on data from week 2 through week 7 (weeks 1 and 8 were included where appropriate given the inclusion of previous or following week counts). Results are reproduced in the Appendix for continuity.

The structural feature effects from the baseline Model 1 are relatively stable and consistent over all weeks. However, in more complex models, the effects were not consistent with the findings in the Week 2 network. For example, in Model 2, the effects of score difference were generally not significant, except in Week 8 where connected individuals did tend to have similar scores. The negative effects of lecture engagement (views/downloads) held true for Weeks 2, 3, and 4, but not beyond that and pointed in the opposite direction for Week 1. Overall, Week 1 exhibited exceptions on several effects, which may not be too surprising given first week effects in any course.

Similar observations concerning fading or inconsistent effects held true in Models 3 and 4 for network statistics with individual attributes from the previous week and subsequent week. The effect of previous week score, significantly positive in Week 2, was not significant in most other weeks except for Weeks 6 and 7, and in the latter case, the sign had changed. Subsequent week score was consistently a positive effect for Weeks 2 and 3, though nonsignificant afterwards. And future lecture engagement was significantly positively twice and negative once (following Week 5).

6. CONCLUSIONS AND FUTURE WORK

Summarizing the results across all four models for Week 2, our major findings are as follows. First, the forum social network is relatively sparse and does not tend to have the preferential attachment feature observed in a lot of social networks. Second, individuals' assignment scores, from current, previous, and future weeks, are all positively related to being more active in social network. Third, social connectedness was negatively correlated with lecture engagement from this week, but positively correlated with lecture engagement in the following week. As with regression models, results from the ERGMs do not enable one to draw causal conclusions, but they do provide evidence about associations. Another finding of this study is that it seems that the significant effects in earlier weeks of the course do not persist later in the course. Sometimes, the effect directions even reversed.

This study has several limitations, which suggest directions for future work. As discussed in the introduction, the meaning of social connectedness here remains fairly simple. Links are unweighted and reflect only the process of replying (or commenting) to an original poster. Content analysis, preferably automated for scalability, could significantly improve the definition of the network itself. Steps toward including natural language processing of posts in defining the network are part of ongoing work. However, even with a simple criterion for network links, a bipartite network [43] could be a useful alternative way to model the post-reply structure of a discussion forum. In such a network, forum threads and individuals would be included as two different types of nodes and directed links would represent the contributions of posters to the threads.

In analyzing our dataset over the weekly sessions, we did find different and even inconsistent results between weeks. The approach of meta-analysis [37] or hierarchical models [40] might be useful in aggregating results and finding the stable trend. On the other hand, it would be interesting to analyze and compare results from other datasets to see which, if any, of the effects in this study generalize to other cases. BDE was a traditional MOOC as opposed to a connectivist or social constructivist MOOC. It is possible that the interplay between social connectedness, engagement, and performance is more pronounced and/or consistent in courses that emphasize social construction of understanding.

7. ACKNOWLEDGMENTS

This research was supported by the National Science Foundation (DRL-1418219).

8. REFERENCES

[1] Albert, R. and Barabási, A.-L. 2002. Statistical mechanics of complex networks. *Reviews of Modern Physics*. 74, (2002), 47–97.

[2] Aviv, R. et al. 2004. Mechanisms and Architectures of Online Learning Communities. *Proceedings of IEEE ICALT 2004*. (2004).

[3] Aviv, R. et al. 2003. Network analysis of knowledge construction in asynchronous learning networks. *Journal of Asynchronous Learning Networks*. 7, 3 (2003), 1–23.

[4] Bandura, A. 1986. *Social foundations of thought and action: A social cognitive theory. Prentice-Hall series in social learning theory*. Prentice Hall.

[5] Bergner, Y. et al. 2015. Estimation of Ability from Homework Items When There Are Missing and / or Multiple Attempts. *Proceedings of the Fifth International Conference on Learning Analytics And Knowledge - LAK '15*. (2015).

[6] Bergner, Y.B. et al. 2015. Methodological Challenges in the Analysis of MOOC Data for Exploring the Relationship between Discussion Forum Views and Learning Outcomes. *Proceedings of 8th International Conference on Educational Data Mining*. (2015).

[7] Brown, R. et al. 2015. Communities of Performance & Communities of Preference. *Proceedings of the 2nd International Workshop on Graph-Based Educational Data Mining*. (2015).

[8] Brown, R. et al. 2015. Good Communities and Bad Communities : Does membership affect performance ? *Proceedings of the 8th International Conference on Educational Data Mining*. (2015), 612–614.

[9] Cho, H. et al. 2002. Social Network Analysis of Information Sharing Networks in a CSCL Community. *Proceedings of the Conference on Computer Support for Collaborative Learning*. (2002), 43–50.

[10] Crossley, S. et al. 2015. Language to Completion: Success in an Educational Data Mining Massive Open Online Class. *Proceedings of the 8th International Conference on Educational Data Mining*. (2015), 388–391.

[11] Dahlstrom, E. et al. 2014. The Current Ecosystem of Learning Management Systems in Higher Education: Student, Faculty, and IT Perspectives. *EDUCAUSE Research Report*. (2014), 27.

[12] Duijn, M.A.J. et al. 2004. p2: a random effects model with covariates for directed graphs. *Statistica Neerlandica*. 58, 2 (May 2004), 234–254.

[13] Eagle, M. et al. 2012. Interaction Networks: Generating High Level Hints Based on Network Community Clustering. *Proceedings of the 5th International Conference on Educational Data Mining*. (2012), 164–167.

[14] Erdős, P. and Rényi, A. 1959. On Random Graphs. I. *Publicationes Mathematicae*. 6, (1959), 290–297.

[15] Frank, O. and Strauss, D. 1986. Markov Graphs. *Journal of the American Statistical Association*. 81, 395 (1986), 832–842.

[16] Gillani, N. and Eynon, R. 2014. Communication patterns in massively open online courses. *The Internet and Higher Education*. 23, (Oct. 2014), 18–26.

[17] Handcock, M.S. et al. 2008. statnet: Software Tools for the Representation, Visualization, Analysis and Simulation of Network Data. *Journal of Statistical Software*. 24, 1 (2008), 1–11.

[18] Haythornthwaite, C. 2001. Exploring Multiplexity: Social Network Structures in a Computer Supported Distance Learning Class. *The Information Society*. 17, 3 (2001), 211–226.

[19] Hoff, P. et al. 2002. Latent space approaches to social network analysis. *Journal of the American Statistical Association*. 97, 460 (2002), 1090–1098.

[20] Holland, P.W. and Leinhardt, S. 1981. An exponential family of probability distributions for directed graphs (with discussion). *Journal of the American Statistical Association*. 76, (1981), 33–65.

[21] Huang, J. et al. 2014. Superposter behavior in MOOC forums. *Proceedings of the first ACM conference on Learning @ scale conference - L@S '14*. (2014), 117–126.

[22] Huang, Y. et al. 2009. The Formation of Combat Groups in Online Games. *the Fourth Annual INGRoup Conference*.

[23] Hunter, D.R. 2007. Curved exponential family models for social networks. *Social Networks*. 29, 2 (2007), 216–230.

[24] Kay, R.H. 2006. Developing a comprehensive metric for assessing discussion board effectiveness. *British Journal of Educational Technology*. 37, 5 (Sep. 2006), 761–783.

[25] Kellogg, S. et al. 2014. A social network perspective on peer supported learning in MOOCs for educators. *The International Review of Research in Open and Distributed Learning*.

[26] Kortemeyer, G. 2007. Correlations between student discussion behavior, attitudes, and learning. *Physical Review Special Topics - Physics Education Research*. 3, 1 (Jan. 2007), 010101.

[27] Laghos, A. and Zaphiris, P. 2006. Sociology of Student-Centred E-Learning Communities : a Network Analysis. *Learning*. (2006), 98–106.

[28] Lazega, E. and van Duijn, M. 1997. Position in formal

structure, personal characteristics and choices of advisors in a law firm: A logistic regression model for dyadic network data. *Social Networks*. 19, (1997), 375–397.

[29] Pattison, P.E. and Wasserman, S. 1999. Logit models and logistic regressions for social networks: II. Multivariate relations. *British Journal of Mathematical and Statistical Psychology*. 52, (1999), 169–193.

[30] Rabbany, R. et al. 2013. Collaborative Learning of Students in Online Discussion Forums : A Social Network Analysis Perspective. *Educational Data Mining: Applications and Trends*. (2013), 1–30.

[31] Ramos, C. and Yudko, E. 2008. "Hits" (not "Discussion Posts") predict student success in online courses: A double cross-validation study. *Computers & Education*. 50, 4 (May 2008), 1174–1182.

[32] Robins, G.L. et al. 2007. An introduction to exponential random graph (p*) models for social networks. *Social Networks*. 29, 2 (2007), 173–191.

[33] Robins, G.L. et al. 1999. Logit models and logistic regressions for social networks, III. Valued relations. *Psychometrika*. 64, (1999), 371–394.

[34] Robins, G.L. et al. 2007. Recent developments in exponential random graph (p*) models for social networks. *Social Networks*. 29, 2 (2007), 192–215.

[35] Siemens, G. 2005. Connectivism: A Learning Theory for the Digital Age. *International Journal of Instructional Technology and Distance Learning*. 2, 1 (2005).

[36] Snijders, T. a. B. and Nowicki, K. 1997. Estimation and Prediction for Stochastic Blockmodels for Graphs with Latent Block Structure. *Journal of Classification*. 14, 1 (1997), 75–100.

[37] Snijders, T.A.B. and Baerveldt, C. 2003. A Multilevel Network Study of the Effects of Delinquent Behavior on Friendship Evolution. *The Journal of Mathematical Sociology*. 27, 2 (2003), 123–151.

[38] Soller, A. 2001. Supporting social interaction in an intelligent collaborative learning system. *International Journal of Artificial Intelligence in …*. 12, 1 (2001).

[39] Song, L. and McNary, S. 2011. Understanding students' online interaction: Analysis of discussion board postings. *Journal of Interactive Online Learning*. 10, 1 (2011), 1–14.

[40] Sweet, T.M. et al. 2012. Hierarchical Network Models for Education Research: Hierarchical Latent Space Models. *Journal of Educational and Behavioral Statistics*. 38, 3 (Sep. 2012), 295–318.

[41] Valente, T.W. et al. 2013. Variations in network boundary and type: A study of adolescent peer influences. *Social Networks*. 35, 3 (Jul. 2013), 309–316.

[42] Vygotsky, L. 1978. Interaction between learning and development. *Mind and Society*. Harvard University Press. 79–91.

[43] Wang, P. et al. 2013. Exponential random graph model specifications for bipartite networks—A dependence hierarchy. *Social Networks*. 35, 2 (May 2013), 211–222.

[44] Wang, Y. et al. 2014. A Longitudinal Study on Learner Career Advancement in MOOCs. *Journal of Learning Analytics*. 1, 3 (2014), 203–206.

[45] Wasserman, S. and Faust, K. 1994. *Social network analysis: Methods and applications*. Cambridge University Press, Cambridge.

[46] Wasserman, S. and Pattison, P.E. 1996. Logit models and logistic regressions for social networks: I. An introduction to Markov graphs and p*. *Psychometrika*. 61, 3 (1996), 401–425.

[47] Xiao, W.-K. et al. 2007. Empirical study on clique-degree distribution of networks. *Physical Review E*. 76, 3 (2007), 37102.

[48] Zhu, M. and Feng, G. 2015. An Exploratory Study Using Social Network Analysis to Model Eye Movements in Mathematics Problem Solving. *Proceeding of the 5th International Learning Analytics and Knowledge Conference (LAK '15)* (Poughkeepsie, NY, 2015).

9. APPENDIX

Tables for ERGM results for weekly networks.

<u>**Model 1**</u>

	Week 1	**Week 2**	**Week 3**	**Week 4**	**Week 5**	**Week 6**	**Week 7**	**Week 8**
Effect	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***
Density (edge)	-4.36* (0.05)	-4.09* (0.09)	-3.58* (0.09)	-3.63* (0.12)	-3.01* (0.14)	-3.05* (0.10)	-3.29* (0.15)	-2.59* (0.15)
Tendency for Hubs (alt. k-stars)	-1.71* (0.12)	-0.99* (0.19)	-1.21* (0.22)	-0.79* (0.26)	-0.94* (0.33)	-1.46* (0.26)	-0.72* (0.32)	-1.64* (0.37)

* indicates $p<0.05$; † indicates $p < 0.1$

<u>**Model 2**</u>

Effect	**Week 1**	**Week 2**	**Week 3**	**Week 4**	**Week 5**	**Week 6**	**Week 7**	**Week 8**
	Estimate (S.E.)	**Estimate (S.E.)**	**Estimate (S.E.)**	**Estimate (S.E.)**	**Estimate (S.E.)**	**Estimate (S.E.)**	**Estimate (S.E.)**	**Estimate (S.E.)**
Density (edge)	-4.49* (0.15)	-4.06*(0.15)	-3.51* (0.15)	-3.43* (0.15)	-2.74* (0.17)	-2.98* (0.14)	-3.48* (0.23)	-2.68* (0.16)
Tendency for Hubs	-1.63* (0.13)	-0.93* 0.19)	-1.15* (0.23)	-0.65* (0.27)	-0.53 (0.38)	-1.41*(0.27)	-0.54 (0.34)	-1.58* (0.31)

(alt. k-stars)								
score (current week)	-0.14* (0.05)	0.19* (0.10)	0.20† (0.11)	0.16 (0.14)	-0.60* (0.22)	0.31† (0.17)	0.06 (0.21)	-0.59† (0.31)
score difference (current week)	0.02 (0.08)	0.18 (0.13)	0.07 (0.14)	0.09 (0.15)	-0.08 (0.21)	-0.20 (0.17)	0.03 (0.20)	0.64* (0.32)
post views (current week)	-0.71* (0.29)	-0.14 (0.34)	-0.42 (0.34)	-0.24 (0.33)	0.33 (0.83)	-0.59 (0.77)	-2.99* (1.20)	0.34 (0.43)
lectures (current week)	0.83* (0.25)	-0.49* (0.20)	-0.30* (0.15)	-0.55* (0.20)	0.71 (2.00)	-1.21 (0.81)	0.37 (0.27)	0.11 (0.19)

* indicates p<0.05; † indicates p < 0.1

Model 3

	Week 2	**Week 3**	**Week 4**	**Week 5**	**Week 6**	**Week 7**	**Week 8**
Effect	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***
Density (edge)	-4.12* (0.30)	-3.51* (0.16)	-3.35* (0.16)	-2.76* (0.18)	-2.97* (0.15)	-3.48* (0.23)	-2.85* (0.18)
Tendency for Hubs (alt. k-stars)	-0.76* (0.20)	-1.12* (0.23)	-0.51 (0.28)	-0.47 (0.39)	-1.26* (0.29)	-0.11 (0.37)	-1.23* (0.33)
score (current week)	0.05 (0.11)	0.38* (0.17)	0.43† (0.26)	-0.94* (0.40)	-0.27 (0.30)	1.03† (0.58)	-0.46 (0.33)
score difference (current week)	0.18 (0.13)	0.07 (0.14)	0.09 (0.15)	-0.08 (0.22)	-0.20 (0.17)	0.03 (0.19)	0.61† (0.31)
post views (current week)	-1.36* (0.47)	-0.67 (0.43)	0.08 (0.43)	0.58 (1.31)	-0.73 (0.89)	-4.07* (1.43)	0.45 (0.56)
lectures (current week)	-0.62* (0.28)	-0.32 (0.22)	0.34 (0.39)	-3.40 (4.76)	1.38 (1.41)	2.38* (0.60)	0.05 (0.66)
Score (previous week)	0.28* (0.13)	-0.25 (0.17)	-0.26 (0.27)	0.39 (0.38)	0.73* (0.29)	-1.13† (0.63)	-0.30 (0.26)
posts reads (previous week)	2.01* (0.50)	0.81 (0.53)	-0.84 (0.61)	-0.16 (1.07)	0.24 (0.61)	2.72† (1.55)	-0.75 (1.07)
post (previous week)	-1.20* (0.42)	-0.40 (0.42)	0.71 (1.36)	-0.65 (0.69)	-0.46 (0.64)	-2.20 (1.47)	2.22 (1.38)
lectures (previous week)	-0.38 (0.64)	0.06 (0.31)	-0.91* (0.39)	0.54 (0.56)	-7.38* (2.87)	-9.34* (2.54)	0.30 (0.70)

* indicates p<0.05; † indicates p < 0.1

Model 4

	Week 2	**Week 3**	**Week 4**	**Week 5**	**Week 6**	**Week 7**
Effect	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***	***Estimate (S.E.)***
density (edge)	-3.95* (0.30)	-3.60* (0.17)	-3.35* (0.16)	-2.95* (0.19)	-2.97* (0.15)	-3.56* (0.24)
tendency for hubs (alt. k-stars)	-0.63* (0.21)	-0.89* (0.25)	-0.49† (0.29)	-0.03 (0.42)	-1.26* (0.29)	0.0005 (0.39)
score (current week)	-0.43* (0.18)	-0.29 (0.25)	0.26 (0.31)	-0.39 (0.59)	-0.46 (0.45)	0.65 (0.55)
score difference (current week)	0.17 (0.13)	0.07 (0.14)	0.09 (0.16)	-0.09 (0.21)	-0.19 (0.17)	0.02 (0.20)
post views (current week)	-1.61* (0.49)	-0.96* (0.47)	0.02 (0.44)	-0.18 (1.48)	-0.87 (0.95)	-4.90* (1.48)
lectures (current week)	-1.26* (0.37)	-0.49 (0.36)	0.54 (0.48)	10.88† (5.99)	0.67 (1.80)	0.40 (0.96)
Score (previous week)	0.32* (0.13)	-0.56* (0.20)	-0.40 (0.28)	-0.30 (0.44)	0.73* (0.31)	-0.64 (0.66)
posts reads (previous week)	2.39* (0.53)	1.12* (0.57)	-0.84 (0.61)	-0.12 (1.20)	0.33 (0.65)	3.85* (1.70)
post (previous week)	-1.21* (0.42)	-0.55 (0.45)	0.71 (1.37)	-0.19 (0.73)	-0.58 (0.67)	-2.70† (1.46)
lectures (previous week)	-1.05 (0.67)	0.15 (0.32)	-0.82* (0.40)	0.30 (0.61)	-7.64* (2.94)	-9.31* (2.58)
score (next week)	0.38* (0.15)	1.02* (0.25)	0.38 (0.30)	0.56 (0.49)	0.17 (0.42)	-0.13 (0.33)
lectures (next week)	0.62* (0.24)	0.13 (0.33)	-3.15 (3.73)	-9.14* (2.07)	0.26 (0.42)	1.98* (0.79)

* indicates p<0.05; † indicates p < 0.1