2009 International Conference on Intelligent Networking and Collaborative Systems

Exploring structural prestige in learning object repositories: some insights from examining references in MERLOT

Miguel-Angel Sicilia, Salvador Sánchez-Alonso, Elena García-Barriocanal and Daniel Rodríguez-García
Computer Science Department
University of Alcalá
Polytechnic Building, Ctra. Barcelona, km. 33.6, 28871 – Alcalá de Henares, Madrid (Spain)
{msicilia, salvador.sanchez, elena.garciab, daniel.rodriguezg}@uah.es

Abstract— **Several existing learning object repositories provide mechanisms for users to arrange personal collections with their selection of resources or to provide reviews and ratings for other's resources, creating a kind of community dynamics. The resulting information can be used to build structural prestige models for the creators of the resources. This paper reports preliminary explorations on relational models that could be used to develop metrics of quality and prestige for learning object authors. Concretely, Social Network Analysis tools are used to analyze the overall community structure of a dataset obtained from the MERLOT repository. Networks extracted from the indirect reference between users through references in personal collections and reviews are examined with regards to the position of relevant community members.**

Keywords- learning object; repositories; social network analysis, prestige, MERLOT

I. INTRODUCTION

Learning object repositories provide a platform for the sharing of Web educational resources, and most of them provide some mechanisms for building community dynamics around their resource base. The community dimension and its social dynamics have been found to be an important aspect for the success of these repositories. For example, Brosnan (2005) provided a conceptualization for that importance based on social capital theory, and Monge, Ovelar and Azpeitia (2008) analyzed the potential impact of Web 2.0 strategies to foster social dynamics and participation in repositories. In a similar direction, Han et al (2008) reported an empirical study on the LON-CAPA repository in which a non-explicit community model was identified on the basis of co-contribution of resources to the same courses by popular authors. In spite of the scattered reports and studies available, our understanding of community dynamics in learning object repositories is still in an inception phase. Evidence has been found that repositories grow linearly with varied patterns of contributor productivity and popularity (Ochoa and Duval, 2008), but we are still far from fully understanding the social aspects of these systems and the motivations and patterns of interaction of their users. Further, measurements of the relevance of contributors have still not been fully developed.

A source of empirical evidence about structural prestige (Wasserman and Faust, 1994) that can be found in some repositories is the availability of references in different forms. These references are often indirect, e.g. some community members store personal "favorite link" collections or provide comments and ratings about some particular resources. Connecting these resources with their authors provides a way to explore indicators of prestige or to seek for potential sub-networks or cluster models. That referencing bears some similarity with citation measures as have been developed extensively in bibliometrics, and they share also the intuition behind models for the ranking of Web pages as the popular PageRank (Page et al., 1999), that are based on outgoing and incoming links between Web pages. However, the assumptions under which citation measures used in the context of scholarly literature have been developed need not be directly transferable to the domain of learning resources, as scientific findings and educational material are created, maintained and evaluated very differently. The same occurs with models as the PageRank, which were developed for the open, non-specific context of the Web, which is fairly different from the smaller and more focused context of a learning resource repository, especially when they are organized around disciplinary domains as is the case of MERLOT[1]. Further, more elaborated and rich forms of referencing can be found in reviews and ratings provided by community members in learning object repositories, beyond those used in these other models.

The consideration of referencing leads to models of *structural prestige*, i.e. a form of relevance determined by the implicit network structure of references or assessments created by community members when forming personal resource collections or when evaluating other's resources. This paper presents an initial exploration on quantitative aspects for structural prestige in the context of learning object repositories. The ultimate objective of that exploration is twofold. On the one hand, providing reliable ranking metrics for learning object authors may eventually play a similar role as citation-based measures does for scientific output. And on the other hand, they would serve as a quality assessment mechanism for resources, which is a need of special relevance in open repositories, in which contribution is not restricted a priori. Indeed, some repositories as Connexions [2] have followed a completely open and

[1] http://www.merlot.org/
[2] http://cnx.org/

978-0-7695-3858-7/09 $26.00 © 2009 IEEE
DOI 10.1109/INCOS.2009.12

unrestricted approach to contribution, providing referential post-publication approaches to quality accreditation, e.g. the concept of Lens in Connexions (Kelty, Burrus and Baraniuk, 2008).

The rest of this paper is structured as follows. Section 2 provides the basic definitions used to build the structural models that are later applied to a concrete case study. Then, Section 3 reports on the analysis done on the social structure of MERLOT (Cafolla, 2002) based on its mechanisms for personal collections and reviews. Discussion following the analysis is provided in Section 4. Finally, conclusions and outlook are provided in Section 5.

II. Social Net Analysis and Repository Communities

Some repositories provide interaction mechanisms for their members, as for example, discussion forums or ways of communicating that are similar to the functionalities provided by social network sites as `facebook`[3]. However, that kind of interaction is usually informal and difficult to inspect, as it is in some cases not available for external data gathering. However, several repositories have started to provide services by which members can share their personal collections of favorite resources and comment or review other's resources. This information is in those sites openly available in some cases, and represents an objective account of usage and expression of preference, at least as much as links in the Web do. Even though it does not directly reflect personal ties between community members, it is a reliable empirical material to analyze the patterns of shared interest between peers, which represents a starting point to gain insights on the communities behind repositories.

The social network model for the repository studied here was based on two different relations that account for the two abovementioned kinds of data: personal collections and reviews. As the base data for the two models is of a different nature, we take the assumption that it is representing two different kinds of referential ties. In both cases, the set of actors $M = \{m_1, ..., m_n\}$ is defined as the user community of the repository considered. Also, the set of learning objects in the repository is denoted as *LO*.

A. Model based on personal collections

The first model is based in a directed, non-valued network built from gathering the personal collections compiled by repository users. A personal collection is simply a set of bookmarks to resources available in the repository that is openly exposed by a given repository user. The relation *P* of personal collection references represents ties in the form of ordered pairs $P<m_i, m_j>$, with m_i and m_j being users, and coming from the following transformation:

$$P < m_i, m_j > \rightarrow \exists PC < m_i, lo > \wedge author(lo, m_j)$$

Where *author(lo, x)* is true if *x* is (one of) the authors of the resource, and *PC<m, lo>* is a bimodal, directed relation between *M* and *LO* representing the learning objects bookmarked by users inside the repository. This kind of indirect referencing has been proposed for modeling prestige in the Web, e.g. in the case of PeopleRank (García-Barriocanal and Sicilia, 2005).

Authors are the original creators of the resources. However, as MERLOT does not stores the resources themselves but only the metadata describing them, it is frequent that the person that initially described the resource is not (one of) its original authors. Then we make a difference from authors and contributors of the resources, being both sets overlapping only to some extent. All the contributors are users in *M*, however, some authors are not MERLOT users, as their resources have been described in MERLOT by others. If instead of *author(lo, x)*, we use the alternative *contributor(lo, x)* a variant *P'* network is obtained. This alternative deserves exploration, as ratings and prestige could be hypothesized to be attached to contributors rather than to authors, since all contributors are MERLOT members and they take over the work of selecting useful resources from the Web.

B. Model based on reviews

The model based on learning object reviews can be represented as a valued graph in which values represent the ratings given to resources by users evaluating them. The relation *R* of review-based references represents ties in the form of valued ordered pairs $R<m_i, m_j>: v$, coming from the following transformation:

$$R < m_i, m_j >: v \rightarrow \exists RR < m_i, lo >: r \wedge author(lo, m_j)$$
$$\wedge v = \Phi(m_i, m_j)$$

Where relation *RR* represents the relation between users and the learning objects they have reviewed. The aggregation operator Φ(x, y) has the purpose of combining the opinion of user *x* to resources authored by user *y*. In a simple formulation, it can be implemented as an average of the ratings.

The domain of *v* is determined by the rating scale given by the repository, and it might represent a simple rating or a vector of ratings. This later case occurs in eLERA, in which the multi-item rating instrument LORI (Vargo et al., 2003) is used for the evaluation of the objects. As in the previous case, changing *author* to *contributor* results in a related but different network *R'*. That additional network should be analyzed separately.

III. Data Analysis

A database from the MERLOT repository was gathered May 2009 by using a crawler that systematically traversed the Web pages of the repository, similar in functionality to the one reported by Biletskiy, Wojcenovic and Baghi (2009). Information of a total of 69,248 users was extracted, of which 1,393 were also recognized as resource authors. 434 of these authors had no declared organization, and the rest of the frequencies of occurrence of individuals from the same organization were below 10, which allows us to discard a

[3] http://www.facebook.com/

possible bias coming from a dominant institution behind the community of users. The fact that there is a significantly higher number of contributors than authors is significant as it points out that MERLOT is more a community of contributors than of authors. This can be a result of MERLOT being a repository only storing metadata and not the contents themselves, as occurs in other systems as Connexions.

A. *Overall structure of the networks*

Table 1 provides some basic measures on the size and structure of the nets defined by the relations considered (for weak components a minimum component size of 10 is required).

The data included for P and R is restricted to those ties in which the authors are also community members. This is why the networks are significantly smaller, however they potentially provide a completely different kind of information. The four networks are relatively sparse, especially in the case of personal collection based networks. In the case of P, no strong components can be found in the network, but there is a large weak component covering the majority of the nodes. The other weak component is a centralized network around a couple of resources, one on “tooth decay” and a second on “oral radiography” which represents a very concrete topical area of interest. This is a case of “highly specific” interest that deserve further analysis, as learning object repository communities can be hypothesized to be highly focused in topical areas. In the case of P’, a strong component with a significantly higher density (.095) can be found, which is showed in Figure 1.

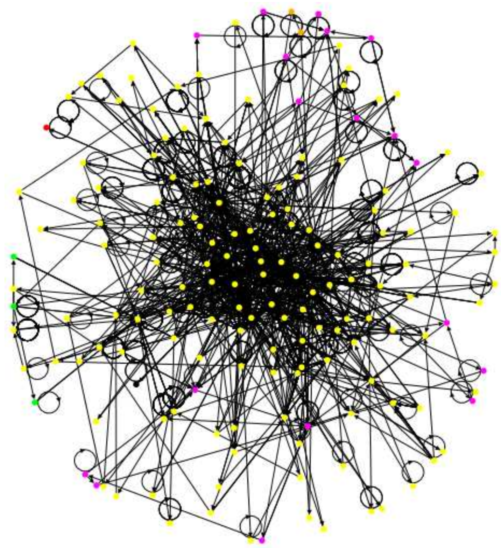

Figure 1. Strong component of size 171 in P’

Parameter	P	P’	R	R’
Number of ties	6998	31935	756	5170
Density	.00058	.00060	.00161	.00117
Diameter	5	14	3	21
Strong components actor size	no strong components of size > 3	{171}	no strong components of size > 3	(4, 3, 27, 21, 4, 5, 3, 10, 9, 3, 3, 3)
Weak component actor size	{3286, 10}	{7143}	{230, 10, 11, 11, 13}	(1786, 10, 13)

TABLE I. DENSITY, CENTRALITY AND COHESION IN THE NETWORKS STUDIED

It can be appreciated in Figure 1 that there is a small but significant amount of self-referencing in personal collections (depicted as loops in the network), concretely P’ has 1978 self loops. These loops have not been removed as they can’t be assumed to be a kind of purposeful bias introduced by users, but a reflection o their interests in exposing their own resources as relevant.

P’ features a distribution of indegrees (one of the major tools to analyze structural prestige) that appear to follow a kind of power law. Concretely, Figure 2 shows the log-log plot of the indegree in the y-axis and the rank of each vertex in the x-axis. The points appear to fall along a single line segment, which is typical of Zipf law distributions, except for the higher rank values.

There are no strong components in *R*, but there are some weak components of a small size. The case of R’ is different, as there are several strong components of a very small size, while there is a larger weak component.

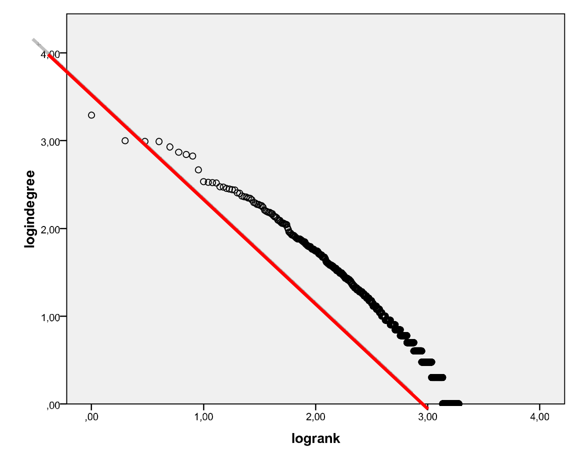

Figure 2. Log-log plot of indegree and vertex rank in P’

The distribution of ratings in R and R' is showed in Figure 3. Most reviews are positive (i.e. ratings above three). In consequence, comments can be considered in general as supporting statements, with only a few acting as negative filters. This needs to be taken into account when considering rankings based on rating aggregation.

Rating	R freq.	R % freq.	R' freq.	R'% freq.
1	18	2.4%	90	1.7%
2	25	3.3%	172	3.3%
3	114	15.1%	733	14.2%
4	308	40.7%	2204	42.6%
5	292	38.5%	1971	38.1%

TABLE II. RATING FREQUENCIES IN R AND R'

B. *Examining central actors*

Central actors in P' were extracted by first computing the indegree partition of the net, and then taking the subgraph with indegree equal or above 10, resulting in a net of 418 members. That network featured a single strong component of 85 actors. Spearman's correlation rank for this subgraph showed a positive correlation of 0.7. A positive correlation of outdegree and betweenness of 0.55 was also found.

MERLOT provides a number of elements describing different aspect of recognition of member behavior, as showed in Figure 3. These include the colored ribbons quantifying different kind of contributions to MERLOT in the categories: gold, silver, bronze and regular.

Table 3 summarizes the network centrality scores and MERLOT recognition level for the collections of users that are in the top-ten according to their indegrees. As can be seen from the table, all the central users have been recognized with either gold or silver ribbons and some of most of them collaborate in peer reviewing inside the repository. User 32414 received the 2008 MERLOT House Cup as a Discipline Editorial Board Member. Even though a complete quantitative analysis of the relation between internal recognitions and network structure is still missing, data points to the fact that there is some form of correlation. This suggests that ranking measures based on incoming arcs may provide good indicators for community service. Table 3 also shows the position in the indegree ranking of each user for R'. Even though some of the higher-ranked users also rank high in R', there are exceptions. Concretely, non-peer reviewers rank lower or do not appear at all. This suggests that comments are providing information substantially different from P'.

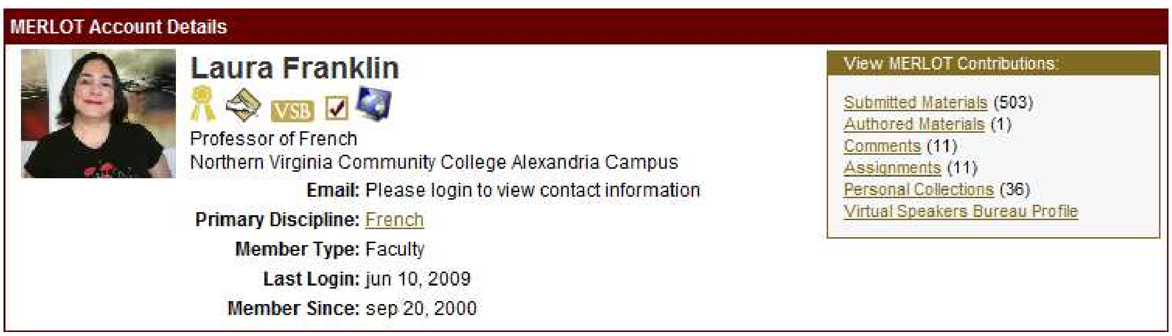

Figure 3. A fragment of a member profile in MERLOT

member	indegree	Outdegree (maximum value 214)	Betweenness (maximum value 1.0)	Closeness (maximum value 3.3)	contributor ribbon	peer reviewer	R'?
23735	198	214	1	1.571	gold	yes	10
17644	192	188	0.014	2.238	gold	yes	18
28799	147	103	0.666	1.583	gold	yes	12
36101	94	82	0.009	2.143	gold	yes	6
18858	65	17	0.434	1.833	gold	yes	17
273543	65	29	0.003	2.357	silver	no	487
31187	55	35	0.054	1.905	silver	yes	102
226210	54	49	0.021	2.119	gold	no	-
4393	46	52	0.212	1.750	silver	yes	88
32414	37	52	0.403	1.667	gold	yes	33

TABLE III. TABLE 3. INDICATORS OF TOP MEMBERS IN P'

Member 23735 in P' has the highest value of betweenness, covering a clearly defined structural hole. That role is covered by 36101 in R'.

The central users in R' were computed as for P'. The resulting structure is shown in Figure 4. The more cohesive subnetwork shaded in the upper left part of Figure 4 is formed by Faculty and students related to computing disciplines (Information Technology, Computer Science, Information Systems, etc.). Other disciplinary subnetworks can be identified in the network, but they are much less sharply defined. Also, some small cohesive subnetworks can be found that represent members from the same institution. An example is a three-node subnetwork shaded in the right bottom part of Figure 3. This indicates that reviews have a potential in identifying communities of interest, which could be used to develop contextualized prestige indicators, resembling the topical separation of bibliometric measures as citation indexes.

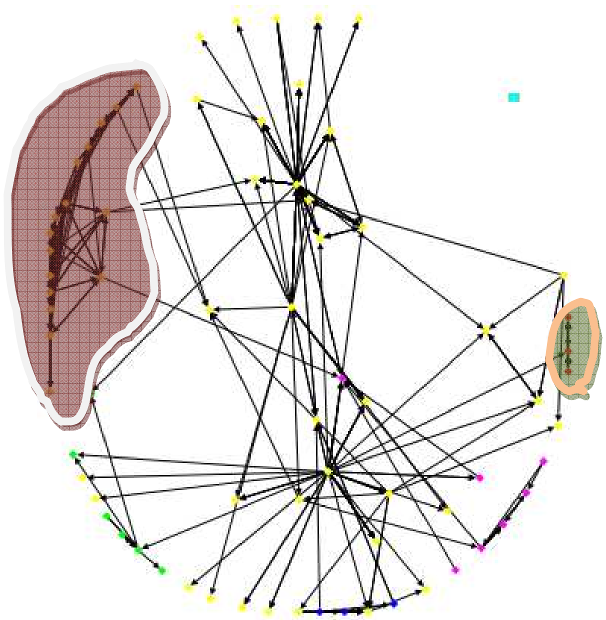

Figure 4. Central users in R'

IV. DISCUSSION

In general terms, P and R represent smaller networks and are less cohesive than P' and R', offering less opportunities for analysis.

P' has single strong and weak components of relatively large size, while in R' strong components are of a very small size. An interesting case is the strong component of size 27 in R'. This appears to be highly centralized, as it has a node with indegree 110, and the rest have indegrees below 6. However, the node has 103 self-references. This represents an example of the case of contributors that rate and comment their own contributions. Such behavior is legitimate, as the inspection of the self-references clearly evidences that the user is trying to provide some form of report of the rationale of his selection of resources included in the repository, which is a clear added value for readers. However, this needs to be taken into account and eventually removing loops from the analysis to avoid a bias caused by this type of use of the platform.

The distribution of indegrees both for *P'* and *R'* appear to follow a Zipf law, even though more evidence is needed to generalize that hypothesis. This has been observed as a common pattern in the Web (Adamic & Huberman, 2002). An important aspect of considering indegrees as a measure of prestige is the existence of correlation between activity in the community and prestige. Ochoa and Duval (2008) found that the popularity of learning objects is independent of the number of objects contributed (productivity). In this study, we have tested the correlation of indegrees and outdegrees both for P' and R' using Spearman's rank correlation coefficient. In P', a significant but negative correlation of -0.448 exists. In the case of R', the correlation coefficient is -0.212. In both cases, a moderate correlation exists but it is negative, which indicates that the hypotheses that more active users (doing more reviews and having larger personal collections) are not necessarily being considered more important (having larger indegrees). However, taking "central users" in P' (see above the procedure followed) results in a set of users in which indegree and outdegree are highly correlated. This can be interpreted in the sense that restricting the analysis to central users, those receiving more votes tend also to be also active voters.

In the case of R', the analysis has found clear discipline-related structural relationships.

The findings in Figure 4 suggest that Social Network Analysis tools might be useful to detect groups sharing common interests which could be contrasted with the pre-established disciplinary structure provided by MERLOT. Figure 5 shows a blockmodel of central actors in R'. The model in Figure 5 shows a highly cohesive cluster in the upper left corner, being a subset of the computing-related subnetwork in Figure 4. The second cluster in the same row contains again a number of actors from computing in the columns. In the rest of the disciplines, the structural arrangement of comments is not clearly discipline-oriented according to the cluster model. The main benefit of blockmodeling is that produces automatic output that is quantitative and can be contrasted with disciplines to detect sub-communities. Contrasting the accuracy of blockmodeling as a tool detecting groups of related interest requires further elaboration, but the analysis described here points out to a good potential of the tool for that task.

The initial exploration reported here has certain limitations that need to be addressed in future studies. An important aspect is that the display order of resources in MERLOT may be producing an effect of increasing the references to resources that are displayed first, even when accessing discipline-related collections[4].

[4] MERLOT provides subcollections for major disciplines, e.g.art or business.

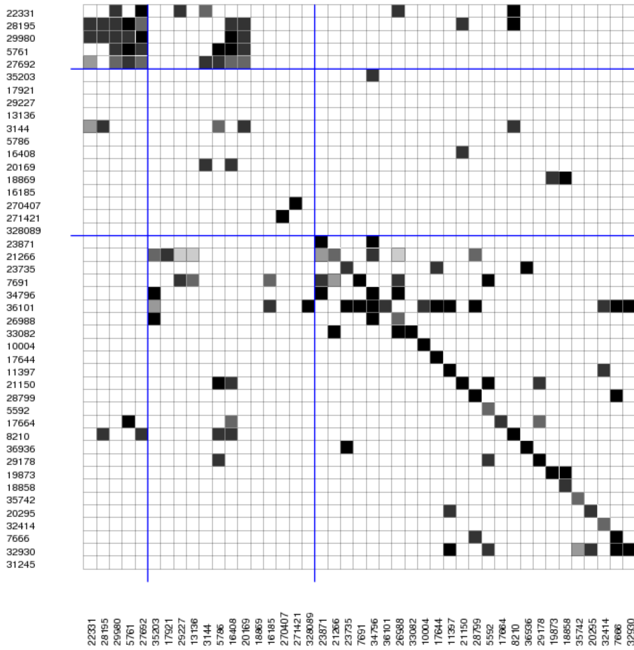

Figure 5. Blockmodel for some[5] central users in R'

V. Conclusions and Outlook

The examination the networks underlying personal collections and reviews between authors and contributors in MERLOT have resulted in some initial insights that deserve more attention in future studies. In general, considering *authors* instead of *contributors* results in less data, as the proportion of authors registered in the repository is relatively small. In addition, as we are considering the community aspect, the main discussion has been around the latter instead of the former, and the conclusions provided here are referred to the contributor networks also. This can be hypothesized to be a consequence of MERLOT being a repository of metadata not storing the resources themselves, and could be contrasted with other systems.

Regarding the main parameter to measure prestige, namely the *indegree*, evidence has been found that it follows a kind of power law, using both kinds of references (personal collections and reviews). This initial result can be taken as a point of departure to the validity of quality ranking systems that take into account this particularity.

Overall, the networks are not cohesive, but in P' a single strong and a single weak component of moderate sizes have been found, which may be hypothesized to conform a "core" community. In the case of R', some cohesive networks among central users can be found and associated clearly with discipline-related communities, as in the case of the computing disciplines. This suggests that approaching quality from a structural perspective might consider the different disciplines separately. Also, clustering based on structural relations can be used to extract those associations from the network in the cases in which the disciplinary frontiers are not clearly observable from the network structure.

The examination of indegrees of central users suggests that the internal recognition system inside the repository (e.g. ribbons or awards) do actually reflect to some extent structural importance, even though P' and R' substantially differ in their rankings.

Future work should account for the disciplinary aspects in the repository. Social network analysis using discipline data can be done by considering the disciplines of the resources in each user's personal collection, rather than using the field "Primary Discipline" appears in member profiles which in some cases seems to be not enough informative. As these repositories are organized around formal or informal communities interested in concrete educational topics, considering discipline may be useful to gain further understanding of community structure.

Also, further work should study the applicability of ranking measures used in Web search as the PageRank (Page et al., 1998) for these kinds of repositories, together with person-centered measures of referencing as those used in scientometrics, e.g. the h-index (Hirsch, 2005).

The relationships between R' and P' as indicators of prestige require also further analysis, as their rankings differ significantly. However, the distribution of ratings in Table 2 suggests they might be related to some extent, as comments are in most cases exposing positive statements, which is also the interpretation of personal collections.

[5] Nodes with less than 10 indegree have been excluded from the analysis tomake the diagram smaller.

Acknowledgments

The results presented in this project have been partially funded by the University of Alcalá and the CAM (Comunidad de Madrid), as part of project MARIA, code CCG08-UAH/TIC-4178, and by the Spanish Ministry of Education through project MAPSEL, code TIN2009-14164-C04-01.

References

[1] K. Brosnan, "Developing and sustaining a national learning-object sharing network: A social capital theory perspective," In: J.B. Williams, & M.A. Goldberg (Eds.), Proceedings of The ASCILITE 2005 Conference, pp. 105-114, Brisbane: Australia.

[2] S. Monge, R. Ovelar and I. Azpeitia, "Repository 2.0: Social Dynamics to Support Community Building in Learning Object Repositories". Interdisciplinary Journal of E-Learning and Learning Objects, 4, 2008, http://ijklo.org/Volume4/IJELLOv4p191-204Monge.pdf

[3] P. Han, G. Kortemeyer, B.J. Kramer and C. von Prummer, "Exposure and Support of Latent Social Networks Among Learning Object Repository Users". Journal of Universal Computer Science (J.UCS), 14, 2008.

[4] X. Ochoa and E. Duval, "Quantitative Analysis of Learning Object Repositories". Proceedings of the World Conference on Educational Multimedia, Hypermedia and Telecommunications, 2008, pp. 6031-6048, Chesapeake, VA: AACE.

[5] S. Wasserman and K. Faust. Social network analysis: methods and applications. Cambridge University Press, Cambridge, New York, Melbourne, 1994.

[6] L. Page, S. Brin, R. Motwani and T. Winograd, “The PageRank citation ranking: Bringing order to the web”. Technical report, Stanford Digital Library Technologies Project, paper SIDL-WP-1999-0120 (version of 11/11/1999).

[7] C.M. Kelty, C.S. Burrus and R.G. Baraniuk, R.G. “Peer Review Anew: Three Principles and a Case Study in Postpublication Quality Assurance”. Proceedings of the IEEE vol 96, issue 6, 2008, pp.1000-1011.

[8] R. Cafolla, “Project Merlot: Bringing Peer Review to Web-based Educational Resources”. Proceedings of the USA Society for Information Technology and Teacher Education International Conference, 2002, pp. 614– 618.

[9] E. García-Barriocanal and M.A. Sicilia, “Filtering Information with Imprecise Social Criteria: A FOAF-based backlink model”. Proceedings of the Fourth Conference of the European Society for Fuzzy Logic and Technology, EUSFLAT 2005, pp. 1094-1098.

[10] J. Vargo, J.C. Nesbit, K. Belfer and A. Archambault, “Learning object evaluation: Computer mediated collaboration and inter-rater reliability”. International Journal of Computers and Applications, vol. 25 issue 3, 2003, pp. 198-205.

[11] Y. Biletskiy, M. Wojcenovic, and H. Baghi. “Focused Crawling for Downloading Learning Objects. – An Architectural Perspective”. Interdisciplinary Journal of E-Learning and Learning Objects vol 5, 2009, pp. 169-180, http://ijklo.org/Volume5/IJELLOv5p169-180Biletskiy416.pdf

[12] L.A. Adamic and B.A. Huberman, “Zipf's law and the internet.” Glottometrics vol. 3, 2002, pp. 143- 150.

[13] J.E. Hirsch, “An index to quantify an individual's scientific research output”. PNAS 102 (46), 2005, pp. 16569–16572.