

Taxonomy and clustering in collaborative systems: The case of the on-line encyclopedia Wikipedia

A. CAPOCCI¹, F. RAO² and G. CALDARELLI^{3,2,4}

¹ *Dipartimento di Informatica e Sistemistica, Università “La Sapienza” - via Ariosto, 25 00185 Rome, Italy*

² *Centro Studi e Ricerche e Museo della Fisica “E. Fermi” - Compendio Viminale, 00185 Rome, Italy*

³ *SMC Centre, INFN-CNR, Dipartimento di Fisica, Università “La Sapienza” - P.le A. Moro 2, 00185 Rome, Italy*

⁴ *Linkalab, Center for the Study of Complex Networks - 09100 Cagliari, Italy*

received 14 June 2007; accepted in final form 13 November 2007
published online 17 December 2007

PACS 89.75.Fb – Structures and organization in complex systems
PACS 89.75.Hc – Networks and genealogical trees
PACS 89.75.-k – Complex systems

Abstract – In this paper we investigate the nature and structure of the relation between imposed classifications and real clustering in a particular case of a scale-free network given by the on-line encyclopedia Wikipedia. We find a statistical similarity in the distributions of community sizes both by using the top-down approach of the categories division present in the archive and in the bottom-up procedure of community detection given by an algorithm based on the spectral properties of the graph. Regardless of the statistically similar behaviour, the two methods provide a rather different division of the articles, thereby signaling that the nature and presence of power laws is a general feature for these systems and cannot be used as a benchmark to evaluate the suitability of a clustering method.

Copyright © EPLA, 2008

Many real systems can be modeled by means of a scale-free network [1–5]. By such mathematical representation it is often possible to better understand the development of these systems and possibly to discover some unexpected behaviour. Much scientific attention has recently focussed on their community structure, often revealed by highly clustered regions of a network. Dividing a network into communities of nodes sharing some given property gives a coarse-grained representation of the whole system. The paramount example of this is given by information networks such as the World Wide Web (WWW). The WWW is a network composed by html documents connected by hyperlinks and for its giant structure only partial studies on its community structure have been produced [6,7]. Once a large information network such as the WWW is decomposed into communities, data mining can be performed in a more efficient way by restricting the data search to smaller regions of the WWW where the desired information is more probable to be found.

Another well-known example of information network is the on-line, user-generated encyclopedias Wikipedia available at <http://www.wikipedia.org> in several languages. Articles of each encyclopedia can be represented as nodes, and the hyperlinks from an article to another within the

Wikipedia form directed networks shaped by the article creations and edits of thousands of individual users around the world. The versions of Wikipedia we explored display statistical properties [8–10] typical of complex networks such as the WWW, whose Wikipedia is a subset, even though their microscopic growth processes differ noticeably: while in the first case users need “administrator” rights to edit webpages, Wikipedia articles can be edited by any user.

Wikipedia networks have varying sizes depending on the language and activity of the underlying users’ community and ranging from a few hundred to more than one million articles. Here we present an analysis of a sample of this set of graphs (hereafter Wikigraphs) collected in September of 2007 from the web site <http://download.wikimedia.org/>. In particular, we study similarities and differences between two possible classifications of Wikipedia articles: their internal categorization and the partition, by a suitable algorithm, of the network formed by articles and hyperlinks between them.

Wikipedia articles are gathered into categories according to their topics. The classification of articles, the creation or the deletion of categories are decided upon the agreement of the whole Wikipedia community. In turn,

categories are organised hierarchically according to their generality. However, articles and categories do not strictly form a perfect tree, since an article or a category may happen to be the child of more than one parent category. Therefore, the taxonomy of articles can be represented a direct acyclic graph [11].

Much work has been devoted to the study of the statistical features of taxonomies, in order to understand whether their overall properties could reveal any general pattern of organization. In most of the cases, one observes power law distributions in the number of offspring that can be explained by means of Yule processes or by the inherent properties of supercritical trees. The first explanation has been proposed for taxonomies of natural species [12] showing a power law decay in the frequency of the number of species for a given genus. Based on such data, Yule [13] introduced a model to explain how mutations in a population of individuals may eventually form a series of different species in the same genus. The results of this process have a rather good agreement with the observed data. Yule processes represent a fundamental mechanism in the production of power laws, though they do not reproduce completely the richness of the scale invariance presented in natural taxonomies. Indeed, when looking at the statistical distribution of the sizes of trees (which corresponds to the distribution of genera in the same family and different families in the same order) a similar power law relation has also been found [14,15]. In this case, the value of the power law exponent may also depend upon the observed ecosystem type [16]. Following this experimental evidence, one may decide to model the development of the whole hierarchical tree of the taxonomy by using a random branching process. It has been analytically shown that the subtree size distribution of a random tree displays a power law decay, $P(s) \propto s^{-\tau}$. The exponent τ is $3/2$ for critical random trees, where the branching number is 1 [17], and equal to 2 [18] if the branching number is larger than 1 as this is often the case in several real instances of growing networks. Therefore, the presence of power laws with an exponent nearby 2 can be considered just a consequence of the parent-child structure of a taxonomy [19].

Beside their classification in categories, Wikipedia articles may also be clustered by the analysis of the network that, through hyperlinks, connect them. Such task is nowadays performed by a number of algorithms [20]. Methods based on edge betweenness and clustering coefficient assume that edges lying on most of the shortest paths in the graph or with low clustering coefficient are likely to connect separate communities. Alternatively, the betweenness can be computed by taking into account all paths between two nodes. Such measure is based on the statistics of random walks taking place on the network [21]. By recursively deleting the edges with larger betweenness or low clustering, the graph splits into its communities [22,23]. Methods that optimise the network modularity, instead, form a cluster of nodes so that the

density of links within the communities is maximised against the number of links among communities [24,25]. Finally, spectral methods are based on the analysis of the eigenvalues and eigenvectors of suitably chosen functions of the adjacency matrix A , whose size is given by the number of vertices n in a graph and whose elements a_{ij} are equal to 1 if an edge exists between nodes i and j and zero otherwise [24,26,27].

In principle one should choose among these algorithms those yielding the highest modularity in the partition obtained. While any of such methods can be applied in small graphs, unfortunately they turn to be unusable in larger networks since they require exceeding computational resources or time. Following the size of the systems analysed, we used only algorithms whose computational time can be kept under control. In particular, the detection of strongly interconnected communities of nodes in a network can still be achieved by finding the attraction basins of random walks on the graph. This is obtained through the method we adopted in our investigation, the MCL algorithm, which provides a fast response in a reasonable time even for networks including thousands of nodes, and can be tuned opportunely in order to maintain its efficiency for even larger systems. However, it has to be noted that the MCL algorithm too is unable to cluster the larger available Wikigraphs (up to millions of articles).

The MCL algorithm [28,29] finds the partition of a network at the desired resolution as follows:

- 1) start with the transition matrix A of the network and normalise each column of the matrix to obtain a stochastic matrix S ;
- 2) compute S^2 ;
- 3) take the p -th power ($p > 1$) of every element of S^2 and normalise each column to unity;
- 4) go back to step 2).

The physical meaning of this procedure is the following: through step 2) we compute the probability that a random walk visits edges two steps apart the starting position. If a walk starts within a community, with greater probability it will remain inside it. By raising these probabilities to a power (step 3)) and then normalising them, we enhance these paths with respect to the others. The effect is to create a statistical matrix S' corresponding to an adjacency matrix (and hence a graph) in which edges between communities are removed.

After some iterations, MCL converges to a matrix $S_{\text{MCL}(p)}$ which is invariant under transformations 2) and 3). Only a few lines of $S_{\text{MCL}(p)}$ have non-zero entries, yielding the nodes' clusters as separated basins (there is in general exactly one non-zero entry per column). As noted above, step 3) reinforces the high-probability walks at short-time scale at the expense of the low-probability ones. The whole process of iteration, on physical grounds, of the

MCL algorithm corresponds to simulating many random walks on the networks and strengthening their flow where it is already strong and weakening it where it is weak. The parameter p tunes the granularity of the clustering. If p is large, the effect of step 3) becomes stronger and the random walks are likely to end up in a greater number of smaller basins of attraction, or communities. On the other hand, a small p produces larger communities. In the limit of $p = 1$, only one cluster is found. The MCL method, thus, has a parameter to be tuned, determining the resolution of the resulting division of the network. In order to compare the communities emerging from the MCL analysis and the taxonomy established by Wikipedia contributors, we set such parameter p to produce approximately the same number of categories observed in the data.

In our work, we have investigated the relation between the category structure of Wikipedia and the clustering properties of the underlying graph representing articles and hyperlinks between them. The category system is a tool to let users browse the content of Wikipedia with the greater ease. Due to the large amount of information to be handled by users, a self-organised categorization system would help in classifying pages without human intervention and discussion. Clustering methods, often based on the topology of a graph, are used for this aim, especially to deal with the user-generated content on the WWW [30–32]. However, the application of automatic clustering methods in each specific context has to be validated by comparing their yielding to the results of manual indexing. The aim of this paper, thus, is to compare and discuss the partition of the graph based on the built-in taxonomy, *i.e.* the categories, and the one obtained by means of the MCL algorithm, applied to the network of the Wikipedia pages connected by internal links. The various data sets analyzed can be downloaded from the archive of the encyclopedia (<http://download.mediawiki.org>). We chose some data dumps selected according to the number of articles S . In particular, the largest set we considered is the English Wikipedia (En, 2042361 articles), while the smallest one is the Norman one (Nrm, 2750 articles), as reported in table 1. They span several orders of magnitude in the number of nodes, ranging at that time from the few thousand articles of the *Norman* archive to the millions of articles of the *English* one. This way, we are able to check whether finite-size effects affect our observations. For each Wikipedia, we analyzed the data sets reporting the category structure and the internal link structure.

We start by considering the category and cluster size distributions, *i.e.* the distribution of the number of Wikipages contained in a category or in a cluster. The statistical properties of Wikipedia taxonomies appear to display a remarkable regularity in the size distribution of categories. As shown in fig. 1 the category size distribution $P(s)$ is heavy tailed, following approximately a power law $P(s) \propto s^{-\gamma}$ with $\gamma \simeq 2.2$ for very different sizes of the system.

Table 1: The Wikipedia versions sampled in the analysis and their size.

Language	Wikipedia	Articles
English	En	2042361
German	De	650241
Italian	It	357538
Norwegian	No	134943
Catalan	Ca	81660
Danish	Da	70757
Croatian	Hr	35932
Galician	Gl	28113
Simple English	Simple	19921
Latin	La	15602
Neapolitan	Nap	12603
Occitan	Oc	10359
Afrikaan	Af	8443
Aragonese	An	7144
Venetian	Vec	5974
Corsican	Co	5324
Interlingua	Ia	3652
Alemannic	Als	3141
Norman	Nrm	2750




Fig. 1: The frequency of category sizes in a sample of the Wikigraphs analyzed here. The legend refers to table 1. The solid line represents $s^{-2.2}$.

We then applied the MCL algorithm to measure the size distribution of topology-based communities. Unfortunately, our survey has to limit itself to Wikipedia smaller than a given size, above which the problem of clustering the network becomes computationally intractable. Interestingly, the clustering-based partition we obtain follows a very similar cluster size distribution with a power law decay for large values of s , as can be observed in fig. 2. In the experiment, we have tuned the granularity parameter of the MCL algorithm in order to obtain approximately the same number of communities and categories.

Nonetheless, fig. 2 shows only a similar partition structure, which does not necessarily imply that the partitions themselves are similar. To compare the two partitions, we




Fig. 2: The category size distribution (triangles) compared to the cluster size distribution (circles) obtained by the MCL algorithm for the Danish (a), Croatian (b), Galician (c), Simple English (d), Latin (e), Neapolitan (f) Wikipedia. The solid lines represent $s^{-2.2}$ in all graphs.

adopt as a measure the *adjusted Rand index* as it has been recently generalised to soft partitions [33]. Standard Rand index [34] results from a pairwise comparison of the elements in two different partitions P, Q . If we denote as

- a) number of pairs of data objects belonging to the same class in P and to the same class in Q ,
- b) number of pairs of data objects belonging to the same class in P and to different classes in Q ,
- c) number of pairs of data objects belonging to different classes in P and to the same class in Q ,
- d) number of pairs of data objects belonging to different classes in P and to different classes in Q ,

we can compute the Rand index R as

$$R = \frac{a + d}{a + b + c + d} = \frac{2(a + d)}{n(n - 1)}, \quad (1)$$

since $a + b + c + d$ is given by the total number $n(n - 1)/2$ of pairs in the system. This is a measure of the agreement between partitions since terms a and d contribute to consistent classifications (agreements), whereas terms b and c are measures of inconsistent classifications (disagreements). Unfortunately, in the case of a partition composed by many clusters, the d element dominates such that the quantity R can be close to 1 even if the partitions substantially differ. To overcome this, the adjusted Rand index R_a has been introduced,

$$R_a = \frac{a - \frac{(a+c)(a+b)}{a+b+c+d}}{\frac{2a+b+c}{a+b+c+d} - \frac{(a+c)(a+b)}{a+b+c+d}}, \quad (2)$$

which is equal to 0 if the two partitions P and Q are randomly drawn [35].

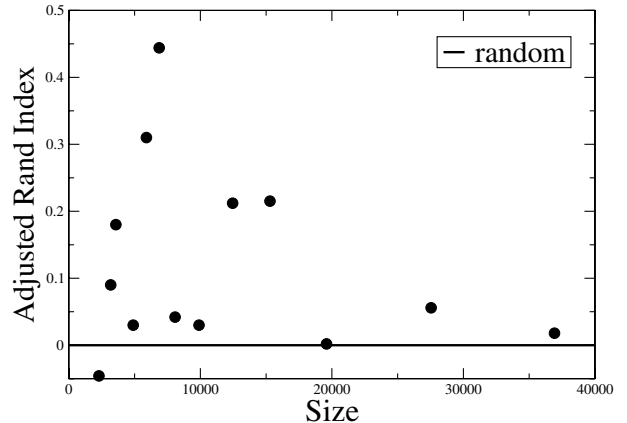


Fig. 3: The overlap between category-based and clustering-based partitions measured by the adjusted Rand index for a sample of Wikipedia networks. The solid line represents the expected value for two random partitions.

It has to be noticed that articles in Wikipedia can lie in more than one category. Therefore, the taxonomy has to be treated as a soft partition, *i.e.* a partition where classes intersection is not null and elements can belong to more than one class with varying intensity. Accordingly, we adopted the generalization of the adjusted Rand index for fuzzy partitions recently introduced [33].

The adjusted Rand index takes very different values when measured in different systems. Moreover, its value seems to be uncorrelated with the network size, as reported in fig. 3 where only articles assigned to at least one category are taken into account. This shows that the categorization of Wikipedia articles does not necessarily correspond to the clustering patterns emerging from the MCL algorithm. The latter, in fact, could result in some case in a quite different organization of knowledge.

From all the above analysis we can conclude that the two divisions of the graphs represent truly different local and global processes on the network, depending upon the decentralised users' action and the consensual collective choice, respectively. This behaviour does not reflect into a different frequency distribution $P(s)$ of the category and cluster sizes. Rather, this quantity is distributed with the same scale-invariant distribution given by $P(s) \propto s^{-2.2}$. This suggests that the presence of power laws in these quantities is more related to the fractal nature of the branching in the category structure [18] when approaching the problem from top to bottom, or conversely to the Zipf's law [36] when considering the inverse bottom-up process of cluster formation. The varying agreement between clustering and categorization across the studied versions of Wikipedia suggests that links in Wikipedia do not necessarily imply similarity or relatedness relations. From a technological point of view, this observation implies that, before switching to automatic categorization of items in Wikipedia and in other information networks, it should be tested how the selected clustering algorithm performs with respect to manual indexing.

AC and GC acknowledge the European project DELIS for support. Authors acknowledge enlightening discussions with S. LEONARDI.

REFERENCES

- [1] ALBERT R. and BARABÁSI A.-L., *Rev. Mod. Phys.*, **74** (2002) 47.
- [2] DOROGOVTSSEV S. N. and MENDES J. F. F., *Adv. Phys.*, **51** (2002) 1079.
- [3] BOCCALETTI S., LATORA V., MORENO Y., CHAVEZ M. and HUANG D.-U., *Phys. Rep.*, **4** (2006) 175.
- [4] NEWMAN M. E. J., *SIAM Rev.*, **45** (2003) 167.
- [5] CALDARELLI G., *Scale-Free Networks* (Oxford University Press, Oxford) 2007.
- [6] ADAMIC L. A. and HUBERMAN B. A., *Science*, **287** (2000) 2115.
- [7] DONATO D., MILLOZZI S., LEONARDI S. and TSAPARAS P., in *Proceedings of the Eighth International Workshop on the Web and Databases (WebDB 2005), June 16-17, 2005, Baltimore, Maryland*.
- [8] BRODER A., KUMAR R., MAGHOUL F., RAGHAVAN P., RAJAGOPALAN S., STATA R., TOMKINS A. and WIENER J., *Comput. Netw.*, **33** (2000) 309.
- [9] ZLATIĆ V., BOŽIČEVIĆ M., ŠTEFANČIĆ H. and DOMAZET M., *Phys. Rev. E*, **74** (2006) 016115.
- [10] CAPOCCI A., SERVEDIO V. D. P., COLAIORI F., BURIOL L. S., DONATO D., LEONARDI S. and CALDARELLI G., *Phys. Rev. E*, **74** (2006) 036116.
- [11] For more detailed guidelines in the categorization of Wikipedia articles, see the web page <http://en.wikipedia.org/wiki/Wikipedia: Categorization>.
- [12] WILLIS J. C. and YULE G. U., *Nature*, **109** (1922) 177.
- [13] YULE G. U., *Philos. Trans. R. Soc. London, Ser. B*, **213** (1925) 21.
- [14] BURLANDO B., *J. Theor. Biol.*, **146** (1990) 99.
- [15] BURLANDO B., *J. Theor. Biol.*, **163** (1993) 161.
- [16] CARETTA CARTOZO C., GARLASCHELLI D., RICOTTA C., BARTHELEMY M. and CALDARELLI G., q-bio.PE/0612023, preprint (2006).
- [17] HARRIS T. E., *The Theory of Branching Processes* (Springer-Verlag, Berlin) 1963.
- [18] DE LOS RIOS P., *Europhys. Lett.*, **56** (2001) 898.
- [19] CALDARELLI G., CARETTA CARTOZO C., DE LOS RIOS P. and SERVEDIO V. D. P., *Phys. Rev. E*, **69** (2004) 035101(R).
- [20] DANON L., DÍAZ-GUILERA A., DUCH J. and ARENAS A., *J. Stat. Mech.* (2005) P09008.
- [21] NEWMAN M. E. J. and GIRVAN M., *Phys. Rev. E*, **69** (2004) 026113.
- [22] GIRVAN M. and NEWMAN M. E. J., *Proc. Natl. Acad. Sci. U.S.A.*, **99** (2002) 7821.
- [23] RADICCHI F., CASTELLANO C., CECCONI F., LORETO V. and DOMENICO PARISI D., *Proc. Natl. Acad. Sci. U.S.A.*, **101** (2004) 2658.
- [24] DONETTI L. and MUÑOZ M.-A., *J. Stat. Mech.* (2004) P10012.
- [25] CLAUSET A., NEWMAN M. E. J. and MOORE C., *Phys. Rev. Lett.*, **70** (2004) 066111.
- [26] KLEINBERG J. M., *J. ACM*, **46** (1999) 604.
- [27] CAPOCCI A., SERVEDIO V. D. P., CALDARELLI G. and COLAIORI F., *Physica A*, **352** (2005) 669.
- [28] ENRIGHT A. J., VAN DONGEN S. and OUZOUNIS C. A., *Nucl. Acid Res.*, **30** (2002) 1575.
- [29] VAN DONGEN S., PhD Thesis (University of Utrecht) (2000).
- [30] HOTHO A., JÄSCHKE R., SCHMITZ C. and STUMME G., *Proceedings of the Workshop on Applications of Semantic Technologies, Informatik 2006, October 2-6, 2006, Dresden, Germany*.
- [31] VOSS J., *Proceedings 10th International Symposium for Information Science, May 30-June 1, 2007, Cologne, Germany*. Available at the URL <http://arXiv.org/abs/cs/0701072>.
- [32] WATANABE Y., ASAHARA M. and MATSUMOTO Y., *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL), Prague, Czech Republic, June 2007*, pp. 649–657.
- [33] CAMPELLO R. J. C. B., *Pattern Recognition Lett.*, **28** (2007) 833.
- [34] RAND W. M., *J. Am. Stat. Assoc.*, **66** (1971) 846.
- [35] HUBERT L. and ARABIE P., *J. Classif.*, **2** (1985) 193.
- [36] ZIPF G. K., *Human Behavior and the Principle of Least Effort* (Addison-Wesley, Reading, Mass.) 1949.