SurpriseMe: an integrated tool for network community structure characterization using Surprise maximization

Rodrigo Aldecoa and Ignacio Marín *

Instituto de Biomedicina de Valencia, CSIC

Abstract

Detecting communities, densely connected groups may contribute to unravel the underlying relationships among the units present in diverse biological networks (e.g., interactome, coexpression networks, ecological networks, etc.). We recently showed that communities can be very precisely characterized by maximizing Surprise, a global network parameter. Here we present SurpriseMe, a tool that integrates the outputs of seven of the best algorithms available to estimate the maximum Surprise value. SurpriseMe also generates distance matrices that allow to visualize the relationships among the solutions generated by the algorithms. We show that the communities present in small and medium-sized networks, with up to 10.000 nodes, can be easily characterized: on standard PC computers, these analyses take less than an hour. Also, four of the algorithms may quite rapidly analyze networks with up to 100.000 nodes, given enough memory resources. Because of its performance and simplicity, SurpriseMe is a reference tool for community structure characterization.

Availability and implementation

The source code is freely available under the GPL 3.0 license at http://github.com/raldecoa/SurpriseMe/releases. SurpriseMe compiles and run on any UNIX-based operating system, including Linux and Mac OS/X, using standard libraries.

Contact

imarin@ibv.csic.es

---

*to whom correspondence should be addressed

1 Introduction

Complex networks are extensively used for representing interactions among elements of a system. This approach is particularly useful in biology: analyzing networks provides relevant information in fields such as genetics [1], neuroscience [2], ecology [3], systems biology [4] or proteomics [5]. An interesting property of these networks is the fact that related nodes tend to create tightly knit groups, usually known as *communities*. By unraveling the close relationships among certain units, community structure characterization improves our understanding of the system as a whole.

In the last years, many strategies have been devised to detect the optimal division into communities of a network. However, none of them alone is able to achieve high quality solutions in all kind of networks [6–8]. In recent works, we demonstrated that Surprise (S) [9–11] is an effective measure to evaluate the quality of any partition of a network [7, 8, 11]. In several complex benchmarks, composed of networks with very different structures, it has been shown that the partition of maximum S corresponds to the real community structure, with a minimal/null degree of error [7, 8, 11]. Although a simple algorithm to maximize S has not been yet devised, it was shown that combining the output of seven high-quality algorithms, always choosing the one that provided the maximum value of S, was sufficient to solve the structure of the networks tested. These algorithms were CPM [12], Infomap [13], RB [14], RN [15], RNSC [16], SCluster [10] and UVCluster [9, 10].

In this article we present SurpriseMe, a tool integrating those seven algorithms. SurpriseMe accelerates the research process by simply accepting a network as input, running internally all those algorithms and outputting their solutions together with their Surprise values. SurpriseMe also calculates distances among the solutions provided by the algorithms, an information that allows to understand how congruent they are [8].

Methods

SurpriseMe: S maximization and distances among solutions

SurpriseMe requires as an input a text file indicating the list of links that characterize the network. Each line of the file contains a link, represented as a pair of nodes separated by a tab or space character. From this text file, the software provides the different programs with the appropriate input files.

As indicated above, SurpriseMe analyses are focused on maximizing the Surprise (S) parameter. Given a partition of a network into communities, S calculates the unlikeliness of finding the observed number of intra-community links in a random network. It is based on a cumulative hypergeometric

distribution [9, 11]:

$$S = -\log \sum_{j=p}^{\min(M,n)} \frac{\dbinom{M}{j}\dbinom{F-M}{n-j}}{\dbinom{F}{n}} \tag{1}$$

where F is the maximum possible number of links of the network, n is the actual number of links, M is the maximum possible number of intra-community links and p is the actual number of links within communities. SurpriseMe calculates either the S values for the seven algorithms or of a subset of them chosen by the user, establishing which one is the best, maximum one. The program also compares all the solutions using either the Variation of Information (VI) [17] or the value that corresponds to (1 - NMI), where NMI means Normalized Mutual Information [18]. In both cases, distance = 0 means that two solutions are identical, and the greater the value, the more different are two partitions. Details of the differences of using VI versus NMI can be found in [7, 8, 19]. The program also estimates the distances to two artificial solutions called "One" (all units of the network are in one community) and "Singles" (each node belongs to a different community). The distances to these two solutions provide additional clues about how each algorithm is behaving [8], All these distances are saved into two distance matrix files (one for VI, another one for 1-NMI) that can be directly imported into MEGA [20], a popular free software which allows an easy visualization of the hierarchical relationships among the different solutions, as shown in [8].

Performance

Given the substantial complexity of the algorithms involved, the current version of SurpriseMe is most useful for networks of small to medium size, typically up to 10.000 nodes. We established the performance of the software by analyzing two types of standard benchmarks. One consisted of networks based on a Relaxed Caveman (RC) configuration [21] with 10% rewiring, which means that well-defined communities are present [7, 8, 11]. The second was a set of Erds-Rnyi (ER) random graphs [22], essentially without community structure. This last benchmark provides an estimate of the maximum time and resources required.

Both with RC and ER structures, networks with up to 10.000 nodes are analyzed by the 7 algorithms in less than an hour using a conventional desktop PC, consuming less than 1 GB of memory. However, larger networks require more powerful hardware and it may be then advisable to switch off the most time- and resource-consuming programs, which are RN, SCluster and UVCluster. Although this obviously may limit S maximization in some cases, close to optimal solutions are generally provided by the four remnant programs in ways that are moreover complementary (i.e., they work optimally in different network structures; see [7, 8]), so their combination will still

generate either very high or maximum S values. With all the programs, we have estimated that a RC network of 50.000 nodes requires 140 hours of analysis and around 60 GB of memory. This is reduced to 40 minutes and 14 GB of memory (RC structure) or 8 hours and 39 GB of memory (ER configuration) if only the four fastest programs are used. For a RC network of 100000 nodes, we have determined that the four fastest algorithms take 3 hours and 30 GB of memory in RC benchmarks, which goes up to 21 hours and 66 GB of memory for ER networks.

Summary

Only few researchers have the time and skills to select, download, compile and run multiple community detection algorithms. SurpriseMe allows to very simply run a set of state-of-the-art algorithms and determine which one generates the best Surprise value, i.e., the best partition of the network. It also provides the user with distance matrices (with VI, 1-NMI values) that may help to understand how the solutions of the different algorithms compare. Very simple to use, it only needs as input a file containing the network to analyze. The well-established power of this type of analysis together with the simplicity of its use, make SurpriseMe an excellent tool for characterizing the community structure of complex networks.

Acknowledgements

References

[1] Michael Costanzo, Anastasia Baryshnikova, Jeremy Bellay, Yungil Kim, Eric D Spear, Carolyn S Sevier, Huiming Ding, Judice LY Koh, Kiana Toufighi, Sara Mostafavi, et al. The genetic landscape of a cell. *Science*, 327:425–431, 2010.

[2] Ed Bullmore and Olaf Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat. Rev. Neurosci.*, 10: 186–198, 2009.

[3] Jordi Bascompte, Pedro Jordano, and Jens M Olesen. Asymmetric coevolutionary networks facilitate biodiversity maintenance. *Science*, 312: 431–433, 2006.

[4] Albert-László Barabási and Zoltan N Oltvai. Network biology: understanding the cell's functional organization. *Nat. Rev. Genet.*, 5:101–113, 2004.

[5] Benno Schwikowski, Peter Uetz, and Stanley Fields. A network of protein–protein interactions in yeast. *Nat. Biotechnol.*, 18:1257–1261, 2000.

[6] Michael T Schaub, Jean-Charles Delvenne, Sophia N Yaliraki, and Mauricio Barahona. Markov dynamics as a zooming lens for multiscale community detection: non clique-like communities and the field-of-view limit. *PloS one*, 7:e32210, 2012.

[7] Rodrigo Aldecoa and Ignacio Marín. Surprise maximization reveals the community structure of complex networks. *Sci. Rep.*, 3:1060, 2013.

[8] Rodrigo Aldecoa and Ignacio Marín. Exploring the limits of community detection strategies in complex networks. *Sci. Rep.*, 3:2216, 2013.

[9] Vicente Arnau, Sergio Mars, and Ignacio Marín. Iterative cluster analysis of protein interaction data. *Bioinformatics*, 21:364–378, 2005.

[10] Rodrigo Aldecoa and Ignacio Marín. Jerarca: Efficient analysis of complex networks using hierarchical clustering. *PloS one*, 5:e11585, 2010.

[11] Rodrigo Aldecoa and Ignacio Marín. Deciphering network community structure by surprise. *PloS one*, 6:e24195, 2011.

[12] Vincent A Traag, Paul Van Dooren, and Y Nesterov. Narrow scope for resolution-limit-free community detection. *Phys. Rev. E*, 84:016114, 2011.

[13] Martin Rosvall and Carl T Bergstrom. Maps of random walks on complex networks reveal community structure. *Proc. Natl. Acad. Sci. USA*, 105:1118–1123, 2008.

[14] Jörg Reichardt and Stefan Bornholdt. Statistical mechanics of community detection. *Phys. Rev. E*, 74:016110, 2006.

[15] Peter Ronhovde and Zohar Nussinov. Local resolution-limit-free potts model for community detection. *Phys. Rev. E*, 81:046114, 2010.

[16] Andrew D King, N Pržulj, and Igor Jurisica. Protein complex prediction via cost-based clustering. *Bioinformatics*, 20:3013–3020, 2004.

[17] Marina Meilă. Comparing clusterings – an information based distance. *J. Multivar. Anal.*, 98:873–895, 2007.

[18] Leon Danon, Albert Diaz-Guilera, Jordi Duch, and Alex Arenas. Comparing community structure identification. *J. Stat. Mech.*, page P09008, 2005.

[19] Rodrigo Aldecoa and Ignacio Marín. Closed benchmarks for network community structure characterization. *Phys. Rev. E*, 85:026109, 2012.

[20] Koichiro Tamura, Daniel Peterson, Nicholas Peterson, Glen Stecher, Masatoshi Nei, and Sudhir Kumar. Mega5: molecular evolutionary genetics analysis using maximum likelihood, evolutionary distance, and maximum parsimony methods. *Mol. Biol. Evol.*, 28:2731–2739, 2011.

[21] Duncan J Watts. *Small worlds: the dynamics of networks between order and randomness.* Princeton university press, 1999.

[22] Paul Erdős and Alfréd Rényi. On random graphs. *Public. Mathem. Debrecen*, 6:290–297, 1959.