

Identifying Covert Sub-Networks Through Iterative Node Classification

Aram Galstyan and Paul R. Cohen
USC Information Sciences Institute
4676 Admiralty Way, Suite 1001
Marina del Rey, CA, 90292, USA
{galstyan,cohen}@isi.edu

Keywords: Link Analysis

Abstract

In this paper we study a problem of identifying a group of related individuals embedded in a larger population. We state the problem in terms of node classification in a social network, and present an iterative algorithm to classify individuals. We test it empirically on data generated by the Hats simulator. Despite its simplicity, the algorithm performs remarkably well. Like most iterative processes, iterative classification has characteristic dynamics. We demonstrate that the dynamics of classifying group members differs from the dynamics of classifying non-members. We call this phenomenon two-tiered dynamics. Our algorithm exploits this difference to identify group members with high accuracy.

1. Introduction

Traditionally most of the literature in social network analysis (SNA) has dealt with overt networks with transparent structures, such as scientific collaboration networks, online communities, and so on. After the September 11 attacks, SNA has increasingly been used to study adversary networks. While covert networks share some features with conventional networks, they are harder to identify because they mask their transactions. Another complicating factor is that adversaries are often embedded in a much larger population (i.e., adversaries have links with both covert and innocent individuals). Hence, it is very desirable to have tools to correctly classify individuals in covert network so that the resources for isolating them will be utilized more efficiently.

We present an iterative algorithm for classifying nodes in networks. Our algorithm utilizes an iterative label propagation scheme where nodes classified at each step help to classify neighboring nodes at the next step.

One can think of this scheme in terms of an epidemic spreading in a heterogeneous network that contains two subpopulations. A discrete epidemic model begins with some infected individuals, and then at each step new individuals become infected if they are linked to super-threshold numbers of infected individuals. Clearly, when two sub-networks are decoupled, the epidemic will be contained in one. When there are links between the subpopulations, epidemics might spread through the whole network, depending on the threshold value for propagating infection. However, when the patterns of links between subpopulations are sufficiently different, the dynamics of the epidemic will be separated in time. This is two-tiered dynamics. Where one finds two-tiered dynamics, one finds different subpopulations.

2. Algorithm and Results

We want to identify small sets of related individuals embedded in much larger populations. For the sake of concreteness, we concentrate on binary classification, although the problem and the algorithm we present below are rather general. That is, we assume that each individual in the network either belongs or does not belong to a certain class; for example, each is an “adversary” or “benign”. Initially, we know the correct class labels of a small subset of the individuals, called “known adversaries”, and the problem is to identify covert adversaries given a graph characterizing meetings between the individuals. Let C_A be the class of adversaries (initially C_A comprises of known adversaries only). At each iteration step, for each individual not in C_A we calculate the number of members of C_A . If for a certain individual, i , this number is greater or equal than some pre-established threshold H , individual i will be classified as an adversary itself, and added to the class C_A . Our algorithm iterates this procedure until a steady state is achieved, e.g., no new node is added to the adversary class. Note that the final state of the system will depend on the threshold value and the pattern of

links between the individuals. If the threshold value is set sufficiently low then the system will evolve to a state where every node has been classified as an adversary. On the other hand, if it is set too high, then no individual will be identified as covert at all. Hence, we expect the algorithm to be very sensitive to the threshold value.

We tested our algorithm on a data generated by the Hats simulator (Cohen and Morrison, 2004). The Hats simulator is a virtual society in which agents carry out individual and collective. Most agents are benign, some are covert terrorists, and a handful are known to be terrorists at the outset. The Hats society includes benign and covert terrorist organizations. Each covert terrorist belongs to at least one terrorist organization; benign agents belong only to benign organizations. Each meeting includes members of exactly one organization. Thus, a meeting planned by a terrorist organization will comprise only terrorists (known and covert), whereas meetings planned by benign organizations might include benign and terrorist individuals. The task for our algorithm is to identify the covert terrorists. A single run of the Hats simulator produces a graph (or adjacency matrix) A of hats' meeting activities. For the analysis presented here, we collected Hats data for $N=1200$ individuals: 50 known adversaries, 1000 benign, and 150 covert. The simulation ran for $T=2500$ time steps.

We ran our algorithm for small, large, and intermediate values of the threshold H . For small values of H , most individuals are classified as adversaries, after a short time. For large values of H , only a few individuals were classified as adversaries. The situation is dramatically different for intermediate values of H : The spread of the epidemic has a distinctive bimodal structure. After a sharp initial spread the epidemic seems to be saturated. However, upon further iterations, the number of infected nodes increases sharply, and in a few steps thereafter all the nodes in the network are infected. Clearly, this corresponds to some kind of threshold-phenomenon in the whole network, where infection of certain nodes causes an epidemic in the whole system. This is illustrated in Fig. 1(a-b) where we plot the number of infected individuals versus time, $N(t)$. Note that this behavior suggests a natural criterion for stopping the iteration. More precisely, in Fig. 1 (c-d) we plot the number of newly infected nodes at each times step versus time, i.e., $\Delta N(t) = N(t+1) - N(t)$. For the threshold value $H=4$ there is only a single peak in $\Delta N(t)$, whereas for $H=6$ one sees two well-separated peaks, denoted P_a and P_b , which are indicative of two-tier dynamics in the spread of the epidemic. If we assume that P_a corresponds to the epidemic spreading in the first (covert) sub-population, and P_b corresponds to the rest of the network, then the iteration should be stopped right before the infection starts to spread through the rest of the system (e.g., at $t = 5$). In this particular example, stopping the algorithm at this point correctly identifies 118 out of 150 covert terrorists and misidentifies only 8 of the 1000 benign nodes as terrorists. This is a very good level of accuracy for such a simple algorithm. More generally,

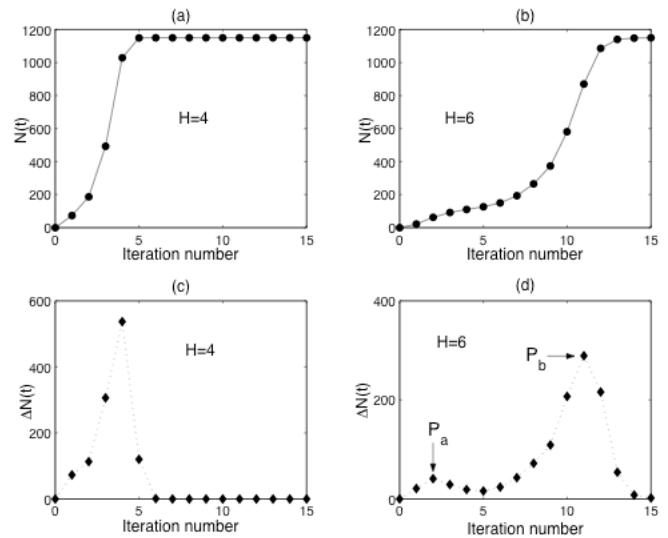


Figure 1: $N(t)$ (a-b) and $\Delta N(t)$ (c-d) for $H=4$ and $H=6$

experiments with Hats data indicate that although the detection error rate of the algorithm varies depending on the particular structure of the network, the amount of the available data, as well as the presence of noise, its performance is rather robust as long as the two-tier dynamics is observed.

In conclusion, we have presented a simple iterative scheme for identifying covert sub-networks embedded in a much larger benign population. Ours is not the first iterative classification algorithm (see, e.g., Macskassy and Provost 2003; Neville and Jensen 2000), however we believe it is the first to explicitly exploit the dynamics of newly classified class instances. Note that many classification and group-finding algorithms are sensitive to parameter settings (e.g., classification threshold). Iterative scheme presented in this paper also demonstrates sensitivity to the threshold parameter H . However, our algorithm uses this sensitivity in a self-consistent way. Namely, the threshold is chosen such as to produce the most pronounced two-tiered dynamics, when one achieves the largest time-separation in the classification process between two sub-populations.

References

- Cohen, P. and C. T. Morrison, C. T. 2004. The Hats Simulator. *Proc. Winter Simulation Conference, 2004*.
- Macskassy, S. A. and Provost, F. J. 2003. A Simple Relational Classifier, *Workshop on Multi-Relational Data Mining in conjunction with KDD-2003*.
- Neville, J. and Jensen, D. 2000. Iterative classification in relational data. *Proc. AAAI-2000 Workshop on Learning Statistical Models from Relational Data*, pages 13-20. AAAI Press, 2000.

