

Community Structure for Efficient Information Flow in ‘ToS;DR’, a Social Machine for Parsing Legalese

Reuben Binns, David Matthews
Web Science Doctoral Training Centre
University of Southampton
Southampton, UK
{rb5g11,dm1x07}@soton.ac.uk

ABSTRACT

This paper presents a case study of ‘Terms-of-Service; Didn’t Read’, a social machine to curate, parse, and rate website terms and privacy policies. We examine the relationships between its human contributors and machine counterparts to determine community structure and information flow.

Categories and Subject Descriptors

H.1.2 [User/Machine Systems]: Human information processing; G.2.3 [Discrete Mathematics]: Applications

General Terms

Human Factors, Measurement, Theory

Keywords

Social Machines; Terms of Service; Privacy; Legal Informatics; Network Science, Community structure

1. INTRODUCTION

Social machines are new combinations of human and machine activity that engage in complex activities and solve problems which were previously difficult or costly for humans or machines to do alone [6]. An interesting application of social machines is to the curation, analysis and rating of legal documents and advice. Contracts, terms-of-service, privacy policies, and licenses all serve important functions in a range of online and offline interactions. However, they incur significant costs on interacting parties as they are often written in complex legalistic language (sometimes referred to as ‘legalese’) which is time-consuming and difficult for individuals to understand without formal training. In addition, despite advances in automatic parsing techniques made in semantic processing and legal informatics [2, 7], computational processes alone may be insufficient. Instead, social machines aimed at handling these tasks and functions are emerging. These systems aim to capture the activity of human actors who write, develop, read, modify, sign, or abide

by such texts, and aggregate that activity to produce new content and services.

We used an existing classificatory framework for social machines to search for and identify systems in this area [6]. For the purposes of this paper, our broader analysis of the domain of social machines for parsing legalese has been discarded. We focus here on just one case study, the ‘Terms of Service; Didn’t Read’ project, a platform for analysing and rating website terms and privacy policies.

2. CASE STUDY: TOS;DR

The ToS;DR platform has around 500 users, who communicate primarily through an open mailing list. Its stated aim is to ‘fix the biggest lie on the internet’ - namely, the statement that ‘I have read and agreed to the terms’. Participants identify, discuss, and annotate clauses in terms-of-service and privacy policies, rating them as either ‘good’, ‘bad’, or ‘neutral’. This activity generates the raw data which drives the service - as they explain it, ‘every thread on the mailing list is a data point’. The number of good, bad and neutral points are tallied to produce an overall rating of a policy. The ratings serve data via an API to other services including a browser plugin.

In order to keep track of changing policies, an automated web crawler called ToSBACK regularly crawls an index of websites and notifies the human participants of any changes to the policies which may need to be reviewed. Any participant can add a website to the index, which is a list of XPath addresses.

The ToS;DR platform provides a case study of how legalese text can be parsed through a social machine. The system receives an input of ‘raw legalese’, analyses its parts in terms of ‘good’ and ‘bad’, and outputs an overall score. An interesting aspect of this process is the community structure and the way information flows within it. In order to explore this we used the Girvan and Newman (GN) algorithm for identifying communities in complex networks [3]. We ran the same analysis on the data before and after the introduction of the ToSBACK crawler, in order to test how the introduction of a computational actor affected community structure.

2.1 Method

We built two graphs G_1 and G_2 corresponding to the mailing list archive before and after the ToSBACK bot began contributing. The vertex sets of both graphs are made up of all the regularly interacting contributors to the mailing list. In total, 2,345 interactions were analysed, from between 28-35 active users. An edge was placed between two vertices if

the relevant contributors responded, or were responded to by another contributor. These edges were then weighted by the number of interactions between the relevant contributors. Finally we applied the GN algorithm to both graphs. GN detects bottlenecks in the graph of social interactions and splits the graph along these bottlenecks into a nested hierarchy of communities [See appendix 4].

2.2 Results

Tree T_1 (Figure 1) was built via the GN algorithm from the mailing list *prior* to the introduction of the ToSBack bot, and corresponds to a nested hierarchy of communities. The vertex labelled 1 represents the entire mailing list and every interaction between contributors, whereas the vertices 2 and 17 represent sub-communities that we will call C_2 and C_{17} respectively. Since 2 is adjacent to 3 and 4 there also exist sub-communities of C_2 that we denote C_3 and C_4 . One can deduce from the GN algorithm that leaf vertices of T_1 correspond to a single contributor, thus C_3 is a community of precisely 1 person. On the other hand C_4 indicates a much larger community – (C_2 without one contributor).

In general a subtree of the tree associated to a network using the GN algorithm that is “close” to a line indicates a hub and spoke community. We have highlighted four line-like subtrees in T_1 ; these subtrees can be recognised as the areas of discussion initiated by or frequently involving particular members.

We can formalise the notion of “line-like” subtrees by analysing the *symmetry* of a tree. One can measure how symmetric a tree is by calculating the number of permutations, $|\text{Aut}(T)|$, of the vertices (of that graph) that preserve adjacent vertices [1]. We call the set of such permutations the automorphism group of T . In general a low $|\text{Aut}(T)|$ coincides with the presence of one or several long line-like subtrees. We calculated that $|\text{Aut}(T_1)| = 16$ and $|\text{Aut}(T_2)| = 24$. This shows a marked increase in the cardinality of the automorphism group and therefore increased symmetry which is further entrenched by normalisation of $\text{Aut}(T)$.





Figure 1: The tree, T_1 , built via the GN algorithm that corresponds to the mailing list *prior* to the ToSBack bot contributions.



5. REFERENCES

- [1] BOLLOBÁS, B. *Modern graph theory*, vol. 184. Springer, 1998.
- [2] FRANCESCONI, ENRICO, E. A. *Semantic processing of legal texts: Where the language of law meets the law of language*. Springer, 2010.
- [3] GIRVAN, M., AND NEWMAN, M. E. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences* 99, 12 (2002), 7821–7826.
- [4] GUIMERA, ROGER, E. A. Self-similar community structure in a network of human interactions. *Physical review E* 68, 6 (2003).
- [5] MURRAY, C. D. The physiological principle of minimum work: I. the vascular system and the cost of blood volume. In *Proceedings of the National Academy of Sciences of the United States of America* (1926).
- [6] SHADBOLT, NIGEL R., S. D. A. S. E. V. K. M. Y. Y., AND HALL, W. Towards a classification framework for social machines. In *SOCM2013: The Theory and Practice of Social Machines, Rio de Janeiro, Brazil* (2013).
- [7] SPINOSA, PIERLUIGI, E. A. Nlp-based metadata extraction for legal text consolidation. In *Proceedings of the 12th International Conference on Artificial Intelligence and Law*. (2009), ACM.

APPENDIX

Girvan and Newman Algorithm:

Let G be a simple weighted graph with edge and vertex sets $E(G)$ and $V(G)$ respectively. The importance of an individual edge $e_{ij} \in E(G)$ is commonly calculated in terms of *edge betweenness centrality* which is defined as follows:

$$\beta(e_{ij}) = \sum_{u \neq v \in V(G)} \frac{\sigma_{uv}(e_{ij})}{\sigma_{uv}}$$

where σ_{uv} is the number of shortest paths from vertex u to v and $\sigma_{uv}(e_{ij})$ is the number of those shortest paths that pass through edge e_{ij} .

In order to determine community structure we will apply the Girvan and Newman (GN) algorithm [3] which associates a binary, rooted tree, T , with a simple weighted graph G as follows:

- (i) The root of T is assigned to be the whole graph G .
- (ii) The edge, e_{ij} , with the highest betweenness centrality in G is determined.
- (iii) Edge e_{ij} is removed from G .
- (iv) If step (iii) disconnects G then connect two vertices to the root of T (these vertices correspond to the connected components of G).
- (v) Iterate until there are no remaining edges in G .

It has been shown [3] that the degree of cohesion in a network can be detected via the GN algorithm and it has been used to identify communities in structures as diverse as scientific collaboration networks, food webs and e-mail networks [3, 4].