

2010

The Structure of Political Discussion Networks: A Model for the Analysis of Online Deliberation

Sandra Gonzalez-Bailon

Oxford Internet Institute, University of Oxford, Oxford, UK, sandra.gonzalezbailon@oii.ox.ac.uk

Andreas Kaltenbrunner

Barcelona Media – Innovation Centre, Barcelona, Spain

Rafael E Banchs

Barcelona Media – Innovation Centre, Barcelona, Spain

Follow this and additional works at: <https://aisel.aisnet.org/jit>

Recommended Citation

Gonzalez-Bailon, Sandra; Kaltenbrunner, Andreas; and Banchs, Rafael E (2010) "The Structure of Political Discussion Networks: A Model for the Analysis of Online Deliberation," *Journal of Information Technology*. Vol. 25: Iss. 2, Article 11.

DOI: 10.1057/jit.2010.2

Available at: <https://aisel.aisnet.org/jit/vol25/iss2/11>

This material is brought to you by the AIS Affiliated Journals at AIS Electronic Library (AISeL). It has been accepted for inclusion in Journal of Information Technology by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.



Research article

The structure of political discussion networks: a model for the analysis of online deliberation

Sandra Gonzalez-Bailon¹, Andreas Kaltenbrunner², Rafael E Banchs²

¹Oxford Internet Institute, University of Oxford, Oxford, UK;

²Barcelona Media – Innovation Centre, Barcelona, Spain

Correspondence:

S Gonzalez-Bailon, Oxford Internet Institute, University of Oxford, 1 St Giles, Oxford OX1 3JS, UK.

Tel: +44 (0) 1865 287 233;

Fax: +44 (0) 1865 287 211;

E-mail: sandra.gonzalezbailon@oii.ox.ac.uk

Abstract

This paper shows that online political discussion networks are, on average, wider and deeper than the networks generated by other types of discussions: they engage a larger number of participants and cascade through more levels of nested comments. Using data collected from the Slashdot forum, this paper reconstructs the discussion threads as hierarchical networks and proposes a model for their comparison and classification. In addition to the substantive topic of discussion, which corresponds to the different sections of the forum (such as Developers, Games, or Politics), we classify the threads according to structural features like the maximum number of comments at any level of the network (i.e. the width) and the number of nested layers in the network (i.e. the depth). We find that political discussion networks display a tendency to cluster around the area that corresponds to wider and deeper structures, showing a significant departure from the structure exhibited by other types of discussions. We propose using this model to create a framework that allows the analysis and comparison of different internet technologies for the promotion of political deliberation.

Journal of Information Technology (2010) **25**, 230–243. doi:10.1057/jit.2010.2

Published online 23 March 2010

Keywords: e-democracy; e-deliberation; online forums; Slashdot; radial trees; social networks; political discussions

Online networks and the political process

The internet and related technologies allow novel forms of political participation. Recent events in Iran following the alleged fraud in the presidential elections of June of 2009 have made the impact of new technologies particularly visible to the general public. Social media like Twitter or YouTube were considered instrumental in the coordination and diffusion of the activities surrounding those protests and mobilisations. Cases like this have shown that social-networking sites may allow citizens to often (not always) overcome censorship and spread information beyond authoritarian control; but the actual role that these new technologies play in strengthening civic networks and enhancing their ability to organise is still a disputed matter. This is due, in part, to the journalistic and anecdotal evidence on which the account of contentious events is usually based. In more democratic societies, however, there is growing evidence that an increasing share

of the population go online to engage in the political process (Bimber, 2003; Chadwick, 2006). This trend challenges claims suggesting that civil society networks are shrinking (Paxton, 1999; Putnam, 2000; McPherson *et al.*, 2006). While it is obvious that internet technologies are facilitating information exchange, it is less clear how the resulting networks form and evolve, and to what extent their structural properties are responsible for a more plural flow of information.

Discussion networks play a crucial role in the democratic process because they give citizens the opportunity to engage in political talk and assess conflicting ideas (Lazarsfeld *et al.*, 1968; Zuckerman, 2005; Mutz, 2006). By discussing politics, people become more acquainted with their own opinions, which can result in a stronger political engagement; and they become more aware of oppositional arguments, which can lead to higher tolerance and even trust in those who hold different views. Empirical research

assessing the consequences of networks for political participation is far from conclusive (Mutz, 2002; Klofstad, 2007). But there is emerging consensus that discussion networks unfold mechanisms of social influence that cannot be grasped just by focusing on isolated individuals. Differences arise when estimating the effects of such influence because the evidence suggests that it can actually work in two directions: discussion networks can amplify preferences if individuals interact with like-minded people or they can bring positions closer and build consensus if they span different pools of opinion (Sunstein, 2007). Because of the different consequences that political discussions can generate, empirical research often leads to normative conclusions about the desirability of political discussions. This opens a point of connection with the theory of deliberation.

Deliberative theory is based on the claim that not all discussions count as deliberation because for it to take place a set of ideal conditions are needed first. Some of those conditions, like equality of all participants or the representativeness of the arguments exchanged, are broadly shared by all proponents of deliberation. Empirical research has sought to test how close political discussions are to this normative ideal. Much of this research has focused on internet technologies (Browning, 1996; Hacker and Van Dijk, 2000; Agren, 2001; Becker, 2001; Grönlund, 2001; Mahrer and Krimmer, 2005). Yet there is still a wide chasm separating normative, often ambiguous, definitions on one side and limited empirical data on the other. The former makes the identification of deliberation (and therefore its empirical evaluation) more difficult, allowing different researchers to focus on different aspects of the deliberative process in a non-cumulative manner; the latter often results in forced operationalisations, or in the analysis of one-time experiences, which also undermines the generality of the theory. This paper attempts to bridge the divide between the theory of deliberation and the practice of political discussions. It builds, on the one hand, on a stylised model that characterises discussion networks with features that reproduce prerequisites for deliberation; and it builds, on the other, on large-scale analyses of the dynamics in which people engage when discussing about politics. This paper seeks to empirically anchor some of the assumptions made about online interactions, in particular the belief that social networking sites fulfil a function regardless of the type of interactions in which users actually engage. The most fundamental claim that this paper makes is that not all networks allow the same flows of information, and that users might form different networks even when using the same internet technology.

This paper is based on observational data that captures what people actually do when they engage voluntarily in public discussions; and it focuses on features of the discussions that are not contingent to the forum itself, but identifiable in other settings. These two aspects allow us to construct a general framework that can help us compare the performance of different online platforms when enabling political talk. This comparative dimension is particularly relevant if we are to evaluate how different internet technologies approximate normative conceptions of political deliberation. By understanding how different internet technologies contribute to shape discussion

networks, we can ultimately design better tools for the promotion of political participation or find ways to effectively incorporate existing tools into decision-making processes. The main contribution of our model is, in that sense, to show that the structure of the discussions can give us relevant information about the underlying dynamics driving public exchange of information. Our focus on the structure of discussion networks, as opposed to the contents, aims to facilitate future work comparing political discussions in different online settings.

The paper proceeds as follows. The following section identifies the main elements behind the ideal of deliberation, and it specifies a model that can test this ideal empirically. The next section introduces the data, which include all the discussions that were held in the online forum Slashdot (<http://slashdot.org>) during the years 2005 and 2006. The subsequent section presents the rationale of our model, which reconstructs discussion threads as hierarchical networks, and uses two of their structural properties, width and depth, to distinguish four ideal types of social interactions. The main theoretical claim of our model is that one of those types provides better conditions for political deliberation than the other three types. We propose using this classification as a benchmark to assess how well online discussions approximate the deliberative ideal. The penultimate section analyses the distribution of Slashdot discussions in the analytical map provided by the model, showing that discussions classified as political exhibit a structure that significantly departs from other types of discussion and, most importantly, fall in the area classified as the deliberative type. The paper concludes with an evaluation of our findings and future lines of research.

The theory and practice of deliberation

Deliberative democracy is based on the normative assumption that public, plural discussions offer a superior form of collective decision making. In contrast with other forms of political participation like voting, which consists on the aggregation of choices that individuals make privately, deliberation is based on social interactions between heterogeneous individuals that are able to revise their preferences in the light of the arguments defended by others. According to the literature, by revealing private information deliberation is able to overcome the impact of bounded rationality, and to build consensus and improve the intellectual qualities of the discussants (Elster, 1998: 11). The normative principles of deliberation stem from the nature of communication, which is seen as an educative process where preferences are transformed rather than aggregated (Habermas, 1984, 1987). Deliberative theory also underlies the notion of 'strong democracy' whereby representative institutions should be supplanted by more participatory ones in order to realise the principle of self-government (Pateman, 1970; Cohen, 1989; Fishkin, 1991; Barber, 1998). This normative ideal resonates with the sociological literature that explores the impact of discussion networks on political participation. According to this literature, discussions also help to educate citizens and to make them more autonomous (Verba *et al.*, 1995; La Due Lake and Huckfeldt, 1998; Paxton, 1999; Putnam, 2000; Zuckerman, 2005). Unfortunately the theory of deliberation

has so far defied a strong connection with empirical research. There are two main reasons for this: the lack of conceptual clarity specifying which types of discussions classify as the deliberative type, and the confusion between the causes and the consequences of deliberation.

Much of the literature on deliberation derives from disagreements over the necessary and sufficient conditions that are required for deliberation to take place (Thompson, 2008). Without these conditions, deliberation is a moving target: it is difficult to match with any particular instance of public discussion, and it can always be argued that some crucial element is missing that disqualifies the entire empirical approach. The problem with this lack of conceptual clarity is not only that it goes against the basic principle of scientific refutability, hampering the development of the theory, but also that it blurs the boundaries between the definition of deliberation and its evaluation (Mutz, 2008). Empirical approaches to political deliberation can help develop the theory by, first, turning the normative assumptions into testable hypotheses and, second, progressively identifying a set of necessary conditions required to distinguish deliberation from other types of discussions.

There are two types of axioms in deliberative theory. The first, procedural, refer to the conditions that define the process of the discussion (such as representative participation). The second, consequentialist, refer to the effects of that discussion, such as being able to filter and choose the most legitimate option (Landa and Meirowitz, 2009). What the theory of deliberation usually does not acknowledge is that these axioms refer to different, and logically independent, realities. Having the right deliberative conditions does not necessarily lead to the best decision: experiments show that, even when the conditions for deliberation are carefully designed, the effects of the discussion might not work in the expected direction – discussions can actually make people adopt more extreme positions than those they originally had (Schkade et al., 2007). And likewise, deliberative theory cannot exclude as a matter of principle the empirical possibility that the best decisions could also be reached by means of non-deliberative forms of participation. Empirical research can help differentiate these two areas of inquiry by clearly specifying whether deliberation acts as the dependent or the independent variable, that is, as a set of conditions to be met or as the conditions that contribute to generate certain outcomes. For this, we need a starting point to measure deliberation that can ultimately lead to the entire set of necessary and sufficient conditions. This paper proposes one such starting point, in line with the arguments that follow.

On a normative level, deliberative democracy is mostly concerned with the issue of legitimacy: one of its core assumptions is that legitimate public decisions do not derive from the predetermined will of individuals but from the process of its formation, that is, from deliberation itself (Manin, 1987: 351–352). The practical implication of this normative requirement is that individuals need to have access to a pool of multiple points of view against which they can contrast their own values and beliefs; and they need to engage in a process of persuasion and argumentation that will help them shape their eventual opinion. From an empirical point of view, this implies an institutional framework that maximises the representativeness of the




Figure 1 Prerequisites of deliberation. Source: Adapted from Ackerman and Fishkin (2002).

debate by including as many different voices as possible; and that also intensifies the amount of public argumentation by allowing participants to engage in an exchange of competing arguments. These two features – the extent of representation and the intensity of argumentation – set deliberation apart from other forms of decision making, as illustrated in Figure 1.

This figure gives a simple map of democratic possibilities that takes into account who participates in the discussion and what kinds of opinions are expressed (Ackerman and Fishkin, 2002: 149). Of these four possibilities, only that represented by quadrant I falls in line with the requirements of mass deliberation: it is the only option that maximises the number of people involved in the discussion while also maximising the extent of persuasion and argumentation leading to the formation of preferences. Quadrant II corresponds to the deliberation of a select group of experts or elite, hence diminishing representativeness; quadrant III corresponds to the type of poll-directed mass democracy promoted by the media, which insert in the public dialogue the unfiltered preferences of a random sample of citizens, using them as an approximation to the private (and therefore non-deliberative) opinions of the general population; and finally quadrant IV corresponds to plebiscitary democracy, where the private preferences of the mass public are aggregated again without any discussion (*ibid*: 150–152). This is a simplistic map of democratic possibilities, but it provides a useful criterion to start differentiating, on the empirical level, types of public communication using two of their features: how representative they are and how much persuasive effort they contain.

A question related to the identification of the prerequisites of deliberation is what kinds of scenarios are more likely to engender those conditions. The same contextual features, like the size and publicity of the deliberation, can play contradictory roles when enhancing the quality of the discussion: on the one hand, the larger the deliberation is, the greater is also the risk that the debate will be dominated by a small number of charismatic speakers; but, on the other hand, the greater the publicity of the discussion, the harder it will be for individuals to be motivated by self-interest alone and the more incentives they will have to genuinely engage in argumentation (Elster, 1998: 111). New technologies have lengthened the list of conflicting features: the internet is seen by some researchers as an

unprecedented opportunity to promote democratic participation (Shane, 2004) but others attach to it new reasons for concern, mostly related to the ability it grants users to filter out contents based on similarity of opinions (Sunstein, 2007). The main limitation of these studies, however, is that they are either based on the technical possibilities of the internet rather than on actual usage or that they are too specific to the idiosyncrasies of one particular deliberative forum to allow drawing general conclusions.

The increasing use of internet technologies to take part in the political process will inevitably stir even more the debate about the best way to promote deliberation. But to advance in that debate we need to devise tools that can help us assess in a unified and systematic manner the impact that different settings have on the discussions and, in particular, on the degree of representativeness and persuasion involved. The approach we propose here is complementary to other approaches that try to determine whether online forums meet the conditions for rational deliberation (Dahlberg, 2001). Rather than looking into the content of the discussions, or assess the nature of the arguments exchanged, we propose focusing on the structure of the interactions in which discussants participate. Our aim is to identify the network features that set the necessary (albeit not sufficient) conditions to reach the ideal of deliberation, and ultimately test how close to that ideal discussion networks are when formed in different online settings. This opens a framework for comparative analysis that is largely missing in the literature.

Discussion networks in slashdot: empirical data

Our empirical strategy consists of analysing observational data of thousands of discussions as they take place in a natural environment. The aim is to differentiate discussions in terms of their representativeness and the amount of argumentation they contain, and use this variation as the criterion to set apart the discussions that reproduce the most prosperous conditions for deliberation. We chose to analyse data collected from the technology news site Slashdot because, unlike younger platforms such as Twitter or Digg, Slashdot (founded in 1997) has had enough time to evolve and consolidate, overcoming the problems associated to spam or misbehaviour and proving its robustness as a discussion forum. Although much buzz has surrounded the use of social media in recent protests, the exceptionality of these events makes it difficult to assess how representative they are of more sustainable forms of political participation; it is also difficult to obtain reliable data that can shed additional light to their patchy and anecdotal observation. Slashdot combines elements from the earlier discussion groups of USENET with features of the more modern Web 2.0 technologies, of which it is considered to be one of the earliest precursors. At the moment we write these lines, Slashdot is still an active forum, refusing to be outdated by more recent Web technologies, and still influencing public perceptions and awareness of the topics discussed. Because of all these reasons, Slashdot offered a perfect subject for our study, providing us with neutral data of participation under conditions of political normality.

Discussions in Slashdot start with short-story posts that often carry fresh news and link to sources of information where readers can find additional details. These posts incite many readers to contribute comments, generating discussions that may trail for hours or even days. In that sense, Slashdot is a hub in a large and intricate information network: it acts as a source from where users get the news and the opinions that they will then post on their blogs or discuss with friends. Most of the commentators in this forum register and comment under their nicknames, although a considerable amount participates anonymously. Slashdot allows users to express their opinions freely, but moderation and meta-moderation mechanisms are employed to judge comments and enable readers to filter contributions by quality. Each comment receives a score from -1 to $+5$, initially starting at a value (in the range from -1 to $+2$) which depends on the reputation of the author and with $+1$ being the default value. Anonymous comments start with a score of 0. When users gain a good reputation (i.e. have a positive karma in Slashdot jargon), the moderation system occasionally grants them points that they can use to modify (by $+1$ or -1) the ratings given to other comments. A meta-moderation system is used to rate the moderators themselves and either remove them from the pool of eligible moderators or reward them with more points.

This moderating system, analysed in detail by Lampe and Resnick (2004), is the main mechanism of the site to sort out high and low quality comments. Moderation in Slashdot fulfils the purpose of organising contributions according to their intrinsic value, sorting nonsense, spam, repetitive or otherwise potentially offensive messages and the most valuable contributions 'from the steady stream of information' (Malda, 1999) and thus allowing users to uphold the quality of the discussions. This feature was particularly important for the purpose of our analyses because it ensures that the information exchanged in the discussions is substantive and relevant for the topic at hand; this provides us with quality data to carry out an inductive classification of discussions.

Several studies have focused on Slashdot precisely because of the high quality interactions it enables. Poor (2005), Halavais (2001) and Baoill (2000) have conducted independent inquiries on the extent to which the site represents a public sphere or a 'virtual public'. While Baoill concluded that Slashdot has as many features of a public sphere as deviations from the ideal model, Poor suggested that Slashdot does broadly fulfil Habermasian's normative requirements, particularly given the effects of the moderation system and the strong relation of the forum with the open source community. Halavais reaches a similar conclusion, warning about the threats that a commercial web could pose to the representativeness of this public space. Poor and Baoill coincide in that one requirement, the universality of participation, is not met due to language restrictions and to unequal access to the technology. Our analyses respond to the same motivation as these studies but differ in two fundamental aspects: first, we use a bottom-up approach to deliberation based on the collective dynamics in which people engage rather than on aprioristic expectations of how those dynamics should look like; and second, we analyse the properties of those interactions

using features that are not contingent to Slashdot. We do not attempt to map onto Slashdot the entire set of conditions required by deliberation because, as the previous section claimed, the theory of deliberation is still engaged in a debate to agree on a unique set of necessary and sufficient conditions. Instead, we use the two broadly accepted features illustrated in Figure 1, the degree of representativeness and of argumentation, to differentiate types of discussion. This gives us a relative rather than an absolute approximation to deliberation: some discussions will be closer to the deliberative ideal depending on their score in those two features; our model aims to identify what these discussions look like.

We analyse all posts and comments published on Slashdot between 26 August 2005 and 31 August 2006. This period does not contain particularly salient events (like, for instance, a presidential election) that could cause higher levels of activity in political discussions; in that sense, our sample corresponds to the lowering tide of the political cycle, which minimises the risk of overestimating the effects of exogenous events on the discussions. The data were obtained in form of raw HTML-pages by a web-crawling process that started in September 2006 and took 4.5 days to complete. These pages were transformed into XML files containing the information summarised in Table 1.¹ All the XML files were imported into Matlab where the statistical analysis was performed. The data set was first presented by Kaltenbrunner *et al.* (2008) and contains roughly about two

million comments written by approximately 10,000 different users, generating approximately 10,000 different discussions. The exact numbers can be found in Table 1.

The distributions for both the number of comments per post and the number of unique users per post are right skewed: the median number of comments (174) and users (104) per discussion is lower than the corresponding averages (207 and 122, respectively). About 18.6% of all comments are posted anonymously. We consider these comments as written by a single user in the forthcoming analysis. Although this assumption is arguably unrealistic, it does not substantively affect the findings we present here. In additional analyses, not reported in this paper but available upon request, we used the alternative (and likewise unrealistic) assumption of treating every anonymous comment as written by a different unique user, but the results did not vary significantly. We also considered the possibility of omitting anonymous comments. However, this would have implied excluding also the responses to those comments, cropping artificially the structure of the discussion trees, and this would have introduced more artificial changes in the data than those implied by the original assumption of a unique anonymous commentator. In any case, our choice only affects marginally some of our results, reported in section ‘Refined model: principal component analysis’.

The posts in our data set, around which discussions originate, are organised into the categories (or sub-domains) listed in Table 2. This list does not necessarily coincide with the list of categories that appear on Slashdot today. Some of these categories have been removed since the data were collected and new categories have been added. The category ‘main’ is the only one in our list that has no clear relation with the context of the posts: it merely contains all those posts which do not fit into any of the other sub-domains. This seems to be an artefact of Slashdot at the time when the data were retrieved. Nowadays, all posted stories do appear in one of the sub-domains with a content descriptor; we therefore decided not to consider the category ‘main’ in the analyses that follow.

The categories displayed in Table 2 are not exclusive. For example, the post entitled ‘Lawmakers Try to Protect Kids from Spam’ is hosted under the sub-domain ‘IT’, but has the primary topic descriptor ‘Spam’ and two secondary descriptors: ‘Communications’ and ‘Politics’. In our analyses we classify as political all posts that have the category ‘Politics’ as one of the descriptors. It is because of the non-exclusive nature of the categories that the second and third columns of Table 2, which show the share of these categories on the total amount of posts and comments in the data set, do not add up to 100%. About 12.4% of all posts belong to more than one category, and most of them to two; only 38 posts appear in three categories.

Users are very heterogeneous in the intensity of their participation. The distribution of the number of comments per user, analysed by Kaltenbrunner *et al.* (2008), has a heavy tail and can be approximated with a truncated log-normal distribution. Most users write only a very few comments during the time-span that our data set covers: 30% write only one and 75% write less than 11; the remaining 25% are responsible for more than 87% of the comments and they are active in most of the categories

Table 1 Information crawled from the Slashdot site

Variables	Description
<i>Posts (N = 10,016)</i>	
Id	Unique identifier of the post
Title	Headline of the post
Sub-domain	Category under which each post is classified
Main topic	Descriptor of the main topic of the post
Secondary topics	One or several descriptors of topics related to the post
Time stamp	Date and hour of publication
Author Identifier	Slashdot editor publishing the post
Body	Text of the post
<i>Comments (N = 2,075,085)</i>	
Id	Unique identifier of the comment
Parent id	Post id or comment id to which the comment replies
Author id	Identifier of the comment’s author
Author name	Name of the comment’s author
Time stamp	Date and hour of publication
Score	Integer value between -1 and 5 indicating the comment’s score obtained from Slashdot’s moderation system
Body	Text of the comment
<i>Commentators (N = 93,636)</i>	
Id	Unique identifier of the commentator

Table 2 Distribution of domain categories classifying posts

Sub-domain	Posts		Comments		Users	
	#	%	#	%	#	%
main	1464	14.6	357,361	17.2	46,747	49.9
apache	10	0.1	1844	0.1	993	1.1
apple	418	4.2	120,444	5.8	22,365	23.9
ask	750	7.5	137,856	6.6	30,864	33.0
backslash	25	0.2	5820	0.3	2694	2.9
books	170	1.7	30,002	1.4	10,690	11.4
bsd	50	0.5	8263	0.4	2991	3.2
developers	377	3.8	78,856	3.8	18,545	19.8
features	5	0.0	1145	0.1	716	0.8
games	2139	21.4	257,334	12.4	33,110	35.4
hardware	1182	11.8	239,820	11.6	37,222	39.8
interviews	29	0.3	7515	0.4	3439	3.7
it	1132	11.3	250,674	12.1	37,360	39.9
linux	605	6.0	127,532	6.1	23,549	25.1
politics	549	5.5	155,830	7.5	25,713	27.5
science	1409	14.1	317,856	15.3	40,667	43.4
yro	982	9.8	269,076	13.0	35,661	38.1

listed in Table 2. Users that only comment sporadically also contribute to discussions classified under different sub-domains; this explains the high percentage of unique users in the different categories. The most interesting categories are ‘games’, which contains 21.4% of all posts but only 12.4% of all comments, and, on the other end of the distribution, topics such as ‘yro’ (short for ‘your rights online’) and ‘politics’, which receive many more comments than the average post. Other topics, such as ‘hardware’ or ‘developers’ seem to be quite average, having nearly the same share on posts and comments.

The structure of the discussion networks: a theoretical model

The model we present here responds to two main questions: Do all discussions share similar features? And if not, which ones approximate better the deliberative ideal? To answer these questions we reconstructed the discussion threads as radial trees, a form of hierarchical networks, following the procedure used by Gómez *et al.* (2008). This network representation places comments to the original post in the first layer, and the comments to these comments in additional nested layers that unfold adding new paths or branches to the tree. This gave us the basic structure of the discussions held in the Slashdot forum. This strategy is based on the implicit assumption that users follow a sequential posting behaviour, that is, that their contributions are submitted as a direct reply to the comment they refer to. This might not be the case for all the contributions, and indeed some comments also refer to previous messages to which they do not reply directly. However, our analyses focus on the general trends: if this unstructured, random posting behaviour were significant, it would not allow us to identify differences in the structure of the discussions; as the following sections show, this is not the case, making sequential posting a reasonable assumption. Once the

discussion trees were reconstructed, we characterised their structures using two basic features: their width, which measures the maximum number of comments at any layer of the network, and their depth, which counts the number of layers through which the discussion unfolds. We chose to focus on these two features because they are a good approximation to the number of different people involved in the discussion and also to the intensity of the argumentation: the deeper a discussion tree, the longer the chains of exchange between participants.

If we assemble these two attributes as a double entry matrix where the horizontal axis measures the width of the discussion networks, and the vertical axis, the depth, we can hypothesise the existence of four ideal types of discussions, as illustrated in Figure 2.

According to the conceptual matrix depicted in the figure, there are four types of networks that can emerge in discussion forums. Networks of Type I are those that attract the attention of a larger number of users and exhibit a higher intensity in the interactions: they maximise both the width and the depth of the networks that are theoretically possible. Networks of Type II are those that capture high-intensity interactions in which only a few participants engage: they exhibit long chains of exchange but only between a few users. Networks of Type III and IV, on the other hand, map discussions where participants are not very much engaged in a dialogue with other users (hence their short branches or paths) but they differ in the amount of users that still want to contribute with their comments: discussions of Type IV are more successful than discussions of Type III in attracting that sort of general attention. If we follow the numbers summarised in Table 2, posts like those classified under the ‘Games’ category tend to generate discussion networks of Type II or III because they do not seem to attract that much attention from commentators when compared to the average post, whereas posts like

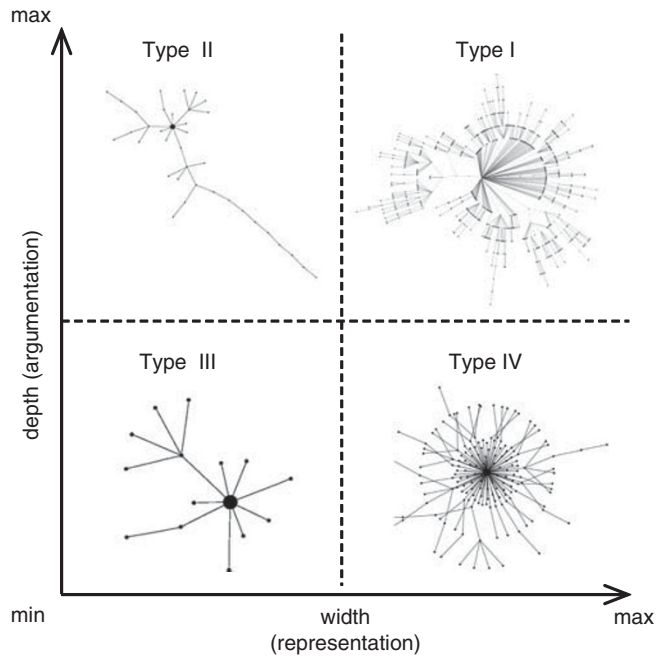


Figure 2 Types of discussions according to the width and depth of the interactions.

those classified under ‘Politics’ tend to generate networks of Type I or IV precisely because of the high number of comments they attract.

Our main theoretical claim with this model is that these two structural features, width and depth, act as good proxies to the two deliberative conditions identified in Figure 1, representativeness and argumentation. The model suggests that quadrant I in Figure 1, which corresponded to the idea of mass deliberative democracy, finds a correlate in discussion networks classified here as Type I: relative to the other networks, networks of Type I maximise the amount of people engaged in the discussion and the amount of persuasive effort they make, much in the same way as quadrant I represented a type of decision making that maximised the number of people involved and the amount of argumentation given. If Figure 1 was a simplistic map of democratic possibilities, Figure 2 is a simplistic map of discussion types, among other things because it just focuses on the structure, not on the content of the information being exchanged. But – our claim is – the structure contains enough information to allow us to differentiate mere discussion from deliberation; or, at least, it allows us to identify discussions that set the most prosperous conditions for deliberation to take place.

The model illustrated by Figure 2 sets the ground to start sorting empirical instances of public talk according to the types of dynamics they generate. One of the basic flaws hindering communication between deliberative theory and empirical research was the lack of conceptual resources to establish when instances of public talk are deliberation or only discussion (Thompson, 2008: 501–502). Our model tackles this issue explicitly by providing a conceptual map that covers all possibilities in a continuum that gradually

approximates the preconditions for deliberation. It also helps direct empirical research by suggesting the following two questions: How many discussions fall in the area identified as the deliberative type? And are these discussions politically relevant? Ideally, political discussions should exhibit the Type I features more often than discussions around other topics, and therefore they should show a tendency to cluster in the upper-right cell of Figure 2: that would mean that political talk has some intrinsic qualities that make it a valuable asset for the democratic process, as it is so often assumed in the literature. We use the Slashdot data to provide an empirical screenshot of how different kinds of discussions scatter in this possibility space, and to illustrate how the model can be used to assess the performance of this or other discussion forums in promoting deliberation.

Data analysis and results

In this section we investigate the spatial distribution of Slashdot discussions in the plane depicted by the model introduced above. We follow a two-step strategy: we first analyse the distribution of discussions according to how wide and deep they are, and we then apply a more sophisticated technique to provide a better approximation to the number of unique people involved in the discussion and to the degree of persuasive effort contained. For this, we consider a narrow and a broad definition of the width and depth of the discussion trees: in the simple model, we only take into account the maximum number of comments at any layer of the network, and the number of nested layers, to define the properties of the discussions; in the refined model, we include additional variables to control for the presence of prolific authors and for the effects of particularly conflictive comments in the overall structure. By including these variables into the analyses we get a richer picture of the actual degree of representativeness and argumentation present in the discussions.

Original model: width-depth analysis

In this section we distribute the discussions from Slashdot in the possibility space opened by the model of Figure 2, using the narrow definition of width and depth. The aim is twofold: to differentiate types of discussions according to the collective dynamics they generate, and to identify where political discussions fall within that space. If we differentiate political from non-political discussions we obtain two groups: in the non-political category we have a total of 9464 discussions, which constitute 94.52% of all available posts; and the remaining 549 discussions fall under the political category, representing 5.5% of all posts. Figure 3 presents the width and depth frequencies for these two categories, as well as their spatial distributions in the width-depth model plane.

The upper half of the figure presents the depth (left) and width (right) relative frequencies for both categories. The overall depth of discussions ranges from 1 to 17, while the maximum width ranges from 2 to 706. As seen from both relative frequency plots, political discussions exhibit a slight tendency to have larger depth and width values than non-political discussions. The lower half of the figure plots all discussions into the width-depth plane. Mimicking

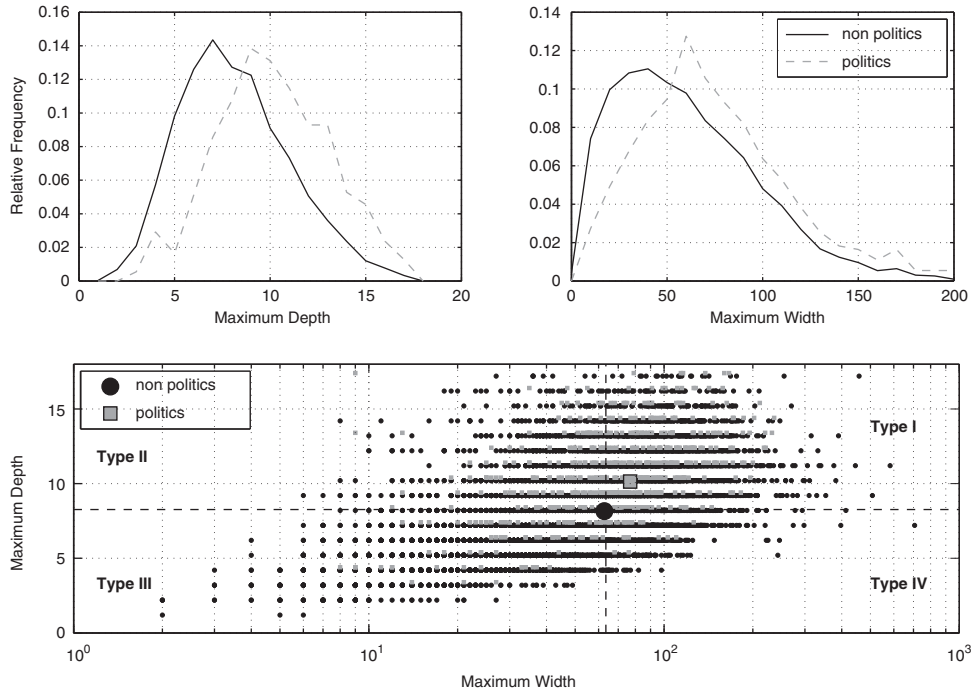


Figure 3 Width and depth frequencies and spatial distributions for *Politics* and *Non Politics* categories.

Table 3 Mean values of width and depth and estimated confidence intervals for the categories *Non-Politics* and *Politics*

Category	Variable	Average	95% confidence interval
Non-politics	Width	62.78	[61.91, 63.65]
	Depth	8.16	[8.10, 8.21]
Politics	Width	76.78	[73.31, 80.24]
	Depth	10.14	[9.89, 10.37]

Figure 2, we split the plane into four parts: every quadrant represents one discussion type, and the intersection corresponds to the mean width (63.6 comments) and the mean depth (8.3 layers) of all discussions. Discussions in the ‘non-politics’ category are depicted as black dots, and discussions in the ‘politics’ category are depicted as grey squares.² In addition to individual posts, a centroid (or average value) for each category is also plotted. These centroids were computed by considering mean values of depth and width for all the posts belonging to the corresponding category. Interestingly, the ‘politics’ mean value appears slightly deviated to the upper right, into the quadrant denominated Type I. The ‘non-politics’ value, on the other hand, falls just over the intersection of the boundaries of the four zones; this is not surprising since the means of 94.5% of all the available data should coincide quite well with those of the entire data set.

In order to find out whether the observed differences between the centroids of both categories are statistically significant, and determine if the tendency of political posts to fall within the Type I region is not attributable to mere chance, we performed a bootstrap test (Efron and Tibshirani, 1986) with $n = 10,000$ of the means for the two

considered categories. This allows us to determine confidence intervals for each estimator and gives us an idea of how significant the observed difference is. The results of these bootstrap tests are summarised in Table 3, which shows that the confidence intervals are quite narrow and do not overlap for the two considered categories.

Figure 4 gives a more intuitive map of how discussions differ by aggregating the posts by categories and plotting their average width and depth values. The centroids of the different categories are distributed quite widely over the plane. The categories ‘politics’, ‘apple’, ‘interviews’, ‘yro [your rights online]’ and ‘backlash’ are the most representative categories for the region mapping networks of Type I. Similarly, the category ‘games’ falls quite clearly in the region of Type III. The case of the remaining categories is less evident: they either cluster together around the four-zone intersection or are too close to a two-zone boundary line.

What these findings suggest is that discussions vary significantly in the type of dynamics they generate: some are more likely to attract a higher number of participants, and some are more likely to incite longer chains of exchange between the discussants. The results also suggest




Figure 4 Width-depth distributions of computed centroids for all post categories.

that these differences are not independent of the topics being discussed: discussions about politically relevant issues involve a wider pool of participants and make them engage in more intense interactions. However, this model is limited for a number of reasons: first, it shows that there is a clear correlation between the width and the depth of the discussions ($r=0.498$, significant at the 1% level) but it does not assess how much of the variance in the discussions results from this correlation; second, the measurement of width disregards the fact that the same users could be contributing the majority of the messages, which would undermine the value of this variable as a proxy for the representativeness of the discussion; and third, the measurement of depth does not take into account the influence that repeated mutual replies between only two users or particularly controversial messages might have in the structure of the discussions. The refined model presented in the next section aims to overcome these limitations.

Refined model: principal component analysis

The refined model adds four new variables to the analyses that provide alternative approximations to the width and depth of the discussions. In addition to the maximum number of comments present at any layer of the network, we use the total number of comments and the total number of unique users participating in the discussion as measures of width. And in addition to the number of nested layers, we use two versions of the h-index as a measure of depth. The h-index has been initially proposed to rank researchers by their scientific outputs, and considers the number of papers published by researchers and the number of times that these papers are cited: if a scientist has an h-index of

11, it means that he has written 11 papers with at least 11 citations each (Hirsch, 2005). For the analyses presented in this section, we used two adapted versions of the index: we calculated an index both for the comments (considering the number of layers and the number of comments per layer) as proposed by Gómez *et al.* (2008), and for the commentators (ordering users by the number of comments they contribute in each post). We use both indexes as alternative measures of depth intended to weight in the controversy of certain comments and the engagement of the most active users.

We carried a correlation analysis of the six variables, resulting in the coefficients reported in Table 4 (all are statistically significant). For the width variables (total number of comments, total number of users and maximum number of comments at any layer), we used the logarithmic transformations since they exhibit log-normal distributions (such as the one depicted in the upper-right plot of Figure 3) – hence the slight difference with the correlation coefficient reported in the previous section. The highest association takes place among the width variables, which means that discussions that contain a high number of comments also tend to contain a high number of unique participants. In addition, the coefficients confirm the significant association between the width and the depth of discussion networks. In the light of these coefficients, we decided to apply principal component analysis (PCA, see Jolliffe, 2002) to reduce the six variables to the two dimensions that maximise the amount of variance explained.

Our analyses showed that the first two components explain 92.76% of the variance of the data. All six variables under consideration contribute positively to component 1, which reflects their high degree of correlation and suggests that a discussion with a wider structure than average will

Table 4 Cross-correlation coefficients for the six variables considered in the PCA

	<i>Comments</i>	<i>Width</i>	<i>Users</i>	<i>h-comm.</i>	<i>Depth</i>	<i>h-users</i>
Comments	1	0.963	0.987	0.848	0.639	0.793
Width	0.963	1	0.971	0.730	0.571	0.697
Users	0.987	0.971	1	0.809	0.640	0.733
h-comm.	0.848	0.730	0.809	1	0.780	0.816
Depth	0.639	0.571	0.640	0.780	1	0.743
h-users	0.793	0.697	0.733	0.816	0.743	1

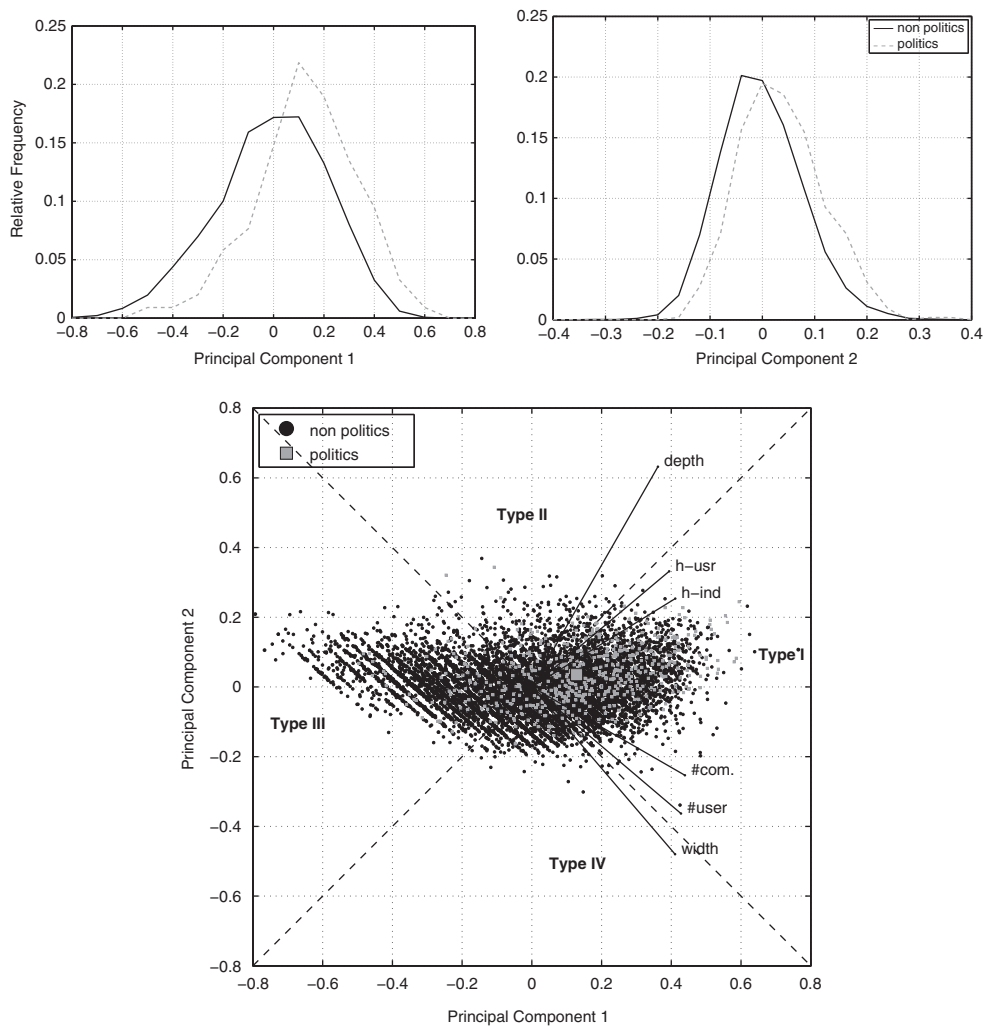


Figure 5 Frequencies and spatial distributions of the two first PCA components for the *Politics* and *Non Politics* categories.

also tend to be deeper than average. Component 2, in turn, discriminates between the variables measuring depth and width. Discussions that are deeper than their width would predict, show a positive value for component 2, whereas discussions that are wider than their depth would suggest, show a negative value for this component.

Figure 5 presents the frequencies of these two principal components for the ‘non-politics’ and ‘politics’ categories,

as well as their spatial distributions in the new model plane. The relative frequencies of the two main principal components are presented in the upper half of the figure, where the categories ‘politics’ and ‘non-politics’ are again represented. What the figure shows is that the distributions of political discussions exhibit a clear deviation to the right with respect to non-political discussions, confirming the trends identified in the previous section. The lower half of

Table 5 Mean values of the two principal data components and estimated confidence intervals for the categories *Non-Politics* and *Politics*

Category	Variable	Mean	95% confidence interval
Non-politics	Comp. 1	-0.0074	[-0.0120, -0.0030]
	Comp. 2	-0.0020	[-0.0036, -0.0005,]
Politics	Comp. 1	0.128	[0.111, 0.145]
	Comp. 2	0.035	[0.029, 0.042]




Figure 6 Spatial distributions of centroids for all post categories within the PCA plane.




Figure 7 Spatial distributions of centroids for low-variability categories within the PCA plane (Confidence Intervals as Gray Shadows).

the figure shows the spatial distribution of the discussions. The boundaries of the four regions are different because we are now using the two principal components to define them: the lines are rotated 45 degrees to account for the rotation of the axes that results from the PCA. However, the relative distribution of individual posts and centroids is still very similar. In particular, the centroid for political discussions falls again in the Type I region, whereas the average value for non-political threads falls over the intersection of the two boundary lines.

Table 5 summarises the results of the bootstrap test ($n = 10,000$) conducted for these two new centroids. They again show that the average values of the ‘politics’ and ‘non-politics’ categories do not overlap each other and that this difference can be considered as statistically significant.

The distribution of average values for each category is shown in Figure 6, which again reproduces the PCA plane. What the figure shows is that, similarly to the original width-depth model, the categories ‘politics’, ‘apple’, ‘interviews’ and ‘yro’ continue to be the most representative of networks of Type I; and the category ‘games’ continues to be the most representative of the networks of Type III. The significance of these observations was assessed running another bootstrap test. We included only the categories with at least 350 posts because otherwise the samples are too small to allow any significant conclusions.

The confidence intervals of these categories are presented as the grey elliptical shadows in Figure 7. With the

exception of the ‘linux’ and ‘developers’ categories, which are positioned on the border of two different regions, all other categories clearly fit into one particular type. More specifically, the ‘hardware’ category belongs to Type IV, categories ‘ask’ and ‘games’, to Type III, and the remaining five categories, including ‘politics’, are all enclosed in the area identified as Type I. This figure also allows us to understand the relation between the different variables used to describe the discussions. As mentioned above, all six variables contribute in a similar manner to the first principal component; the projections of the grey lines on the x-axis show the proportion of this contribution to the first principal component.

This refined model allows us to observe more detailed similarities among discussion networks. We can clearly see that the categories ‘apple’, ‘politics’ and ‘yro [your rights online]’ contain discussions with very similar structures: they have a high, above-average number of users and comments, and raise more controversy than the average discussion, as measured by the h-indexes. Other categories such as ‘hardware’ (of Type IV) attract more users than the average post, but these users do not interact very much with one another and the discussions on these topics are reduced to mere comments on the original story. The contrary seems to be the case for the category ‘developers’. Although it is just on the edge between networks of Type II and I, these discussions attract less users than average; however, these users seem to entangle quite frequently in deep



discussion. Finally, the posts in the 'games' category (one of the most prominent in terms of number of posts contributed) generate the simplest type of discussions: those that include a low number of users and comments, and exhibit a very flat structure.

These results are relevant for the empirical analysis of deliberation for two reasons: first, because they draw a clear picture of why not all types of public discussion entail the same collective dynamics; and second, because they show that political talk has some intrinsic quality that sets it apart from other type of debates. The results reported here allow us to draw conceptual distinctions (in this case, between types of discussion) that are unambiguously rooted in empirical features: the width of discussion networks as an approximation to the representativeness of the discussion, and the depth of the networks as an approximation to the degree of argumentation and persuasion involved. Our model uses simplistic measures of heavily loaded theoretical concepts, but this is where their most valuable contribution lies: they provide a criterion to operationalise and measure the preconditions for deliberation, which is a necessary step if we are to progressively devise better measurement devices. The following section considers some further implications.

Discussion and further research

The analyses presented in this paper show that the structure of discussion networks exhibit significant differences when the topic being discussed changes. Our aim was to identify if discussions in online forums resemble the deliberative ideal. We focused on two preconditions for deliberation, the representativeness of the discussion and the presence of persuasive effort, and we modelled them using two features: the width and the depth of discussion networks. We found that some discussions approximate the deliberative ideal better than others, and that the contents of the discussion matter to determine that approximation. One interesting result was the similarity between discussions of the categories 'apple', 'politics' and 'yro [your rights online]'. While the similarity between the last two is arguably obvious, given the relatedness between political topics and issues concerning privacy, censorship or open access to information (some of the topics discussed under 'yro'), it is surprising to find similar discussion dynamics under the category 'apple'. However, this seems less surprising if we take into account the ideological and missionary zeal with which many users preach the supremacy of either Linux, Apple or Windows (the type of discussions frequently found in the 'apple' category), which is comparable to the type of ideological discussions that develop around most political issues. Slashdot is, in the end, a technology news site oriented to a savvy IT audience.

The implications of these findings are threefold: on the theoretical level, the model provides an unambiguous empirical definition to start differentiating deliberation from mere discussion; on the empirical level, the model proposes a framework to undertake comparative analysis and assess different online settings in the light of the deliberative dynamics they promote; and on the policy level, our approach offers tools to start thinking about the best way to incorporate already successful platforms of

public discussion into decision-making processes. The last few years have witnessed a considerable amount of government-sponsored initiatives to use new technologies for the promotion of political engagement; however, these initiatives have usually not attracted large amounts of participation or they have not been successful at becoming permanent features of the participatory landscape (Coleman and Norris, 2005). This contrasts with the proliferation on the internet of self-organised communities that give expression to the concerns of citizens and have become permanent settings for political participation. Analysing the mechanisms that underlie the functioning (and, therefore, success) of these online communities will provide valuable information to assess and redesign government-initiated projects. Likewise, it will help us give better accounts of the role that social media play in articulating contentious politics (of which the protests in Iran are an example), acknowledging the fact that the same technology can enable very different types of collective dynamics. This paper proposes a strategy to uncover those mechanisms and the consequences they have for the flow of information.

An important element of our approach is that it allows us to classify a given discussion by quantifying the extent to which it belongs to a particular type, but there are still elements related to the growth of those networks that need further attention. Although the results presented here are based on the final structure of the discussions, our model can also be used to track the discussions as they evolve over time. This would give a much better insight into the mechanisms that underlie the formation of online discussion networks, and test some of the theoretical claims about the effects that discussions have on individual engagement. Discussions should all start as networks of Type III (with a low number of participants) and then systematically grow towards networks of Type I, monotonically increasing the value of the first component as time goes by. The interesting aspect of this progression lies in the second component, which we expect to oscillate to form networks of different structures. *A priori*, we would expect to find very different types of trajectories in the plane depending on the type of discussions. We could also improve our approach by adding more variables to the calculation of the components, including more information about the total number of users that read the forum but do not participate: taking into account the amount of time these passive users devote to follow discussions might increase the discriminative power of our model. However, the availability of this kind of data is much more restricted. Finally, we also plan to use text mining techniques to complement our structural approach to the discussions with a more qualitative account of the actual emotional content (and potential for disagreement) of the discussions.

Our findings apply to one particular online forum, Slashdot, which, as has been already mentioned, is known by the high-quality discussion it generates thanks to its moderation and filtering system. Still we have found a number of interesting things that we could try to replicate in other online discussion forums. We were interested in identifying deliberation by analysing the structure of the discussions, and we found that not only political but also other types of discussions meet the preconditions for

deliberation. It remains to be seen if other forums, like newsgroups, are so successful in attaining these outcomes. The ultimate aim of our model is to contribute to this comparative analysis: we can use the same PCA-plane we have defined here to calculate the location of discussions in other forums using their own descriptive variables. It would even be possible to compare the entire set of discussions from different forums by calculating their distribution in the same plane.

All in all, the model we presented in this paper provides researchers with a framework that can help them undertake comparative analysis of how different internet technologies promote political deliberation. The model uses simple metrics that give us relevant information about the structure of interaction in which users engage when discussing in public forums. Advancing in this type of analysis will give us a better understanding of how individuals engage in the political process using online technologies, and it will also allow us to devise better tools to enhance that participation. A non-intuitive fact suggested by our model is that the same setting can lead to different forms of discussion. This adds an important dimension to the institutional design approach to deliberation: there might be a number of contextual variables that, as explained in the literature review, affect the outcomes of public discussions, for instance the size of the community and the publicity of the discussion; but our model suggests that, even when these variables are held constant, there might still be differences in the structure exhibited by the discussions. This invites us to look at other variables beyond the scenario in which discussions take place, like the content of the discussions themselves.

The internet has increased the range of spontaneous, self-organised, bottom-up forms of political participation. The influence of new communication technologies has become particularly salient in countries where there is no open access to information and new media are the only resort citizens have to spread their discontent. But before we can consider how to best link these initiatives to formal politics, or how to empower citizens in their fight against authoritarian regimes, we need a better understanding of how these forms of participation and communication work. This paper is an attempt to move in that direction, bridging in the way normative and empirical theories of political deliberation. Ultimately, we expect this approach to be able to provide the knowledge to inform both policy-makers and practitioners about how to exploit internet technologies to make governments more receptive to the concerns, opinions and demands of citizens.

Acknowledgements

This work has been partially funded by the Càtedra Telefónica de Producció Multimèdia and by the Ramon y Cajal program funded by the Spanish Ministry of Education and Science; it has also benefited from the R + D project SEJ2006-00959/SOCI. We are grateful to the attendants of the Nuffield-OII Networks seminar and the OII-Berkman Center workshop on Internet and Democracy for their comments to previous versions of this paper, and to Bernie Hogan for his helpful suggestions. We are also indebted to three anonymous reviewers for their advice.

Notes

- 1 Post-processing caused by the presence of duplicated comments was necessary due to an error of representation on the website. This explains discrepancies in the total number of comments according to our study and to the Slashdot figures for certain posts.
- 2 The maximum depth for political discussions has been offset by a small amount in order to avoid overlap with the non-political posts.

References

- Ackerman, B. and Fishkin, J.S. (2002). Deliberation Day, *The Journal of Political Philosophy* 10(2): 129–152.
- Agren, P.O. (2001). Is Online Democracy in the EU for Professionals Only? *Communications of the ACM* 44: 36–38.
- Baoill, A.O. (2000). Slashdot and the Public Sphere, *First Monday* 5(9), http://www.firstmonday.org/issues/issue5_9/baoill/index.html.
- Barber, B.R. (1998). Three Scenarios for the Future of Technology and Strong Democracy, *Political Science Quarterly* 113(4): 573–590.
- Becker, T. (2001). Rating the Impact of New Technologies on Democracy, *Communications of the ACM* 44: 39–43.
- Bimber, B. (2003). *Information and American Democracy. Technology in the Evolution of Political Power*, Cambridge: Cambridge University Press.
- Browning, G. (1996). *Electronic Democracy. Using the Internet to Influence American Politics*, Pemberton: Wilton CT.
- Chadwick, A. (2006). *Internet Politics. States, Citizens, and New Communication Technologies*, New York, NY: Oxford University Press.
- Cohen, J. (1989). Deliberative Democracy and Democratic Legitimacy, in A. Hamlin and P. Pettit (eds.) *The Good Polity*, Oxford: Blackwell, pp. 17–34.
- Coleman, S. and Norris, D.F. (2005). A New Agenda for E-Democracy, Oxford Internet Institute Discussion Papers, n. 4.
- Dahlberg, L. (2001). The Internet and Democratic Discourse: Exploring the prospects of online deliberation forums extending the public sphere, *Information, Communication & Society* 4(1): 615–633.
- Efron, B. and Tibshirani, R. (1986). Bootstrap Methods for Standard Errors, Confidence Intervals, and Other Measures of Statistical Accuracy, *Statistical Science* 1(1): 54–75.
- Elster, J. (1998). *Deliberative Democracy*, Cambridge: Cambridge University Press.
- Fishkin, J.S. (1991). *Democracy and Deliberation*, Yale: Yale University Press.
- Gómez, V., Kaltenbrunner, A. and Lopez, V. (2008). Statistical Analysis of the Social Network and Discussion Threads in Slashdot, in WWW 2008: Proceedings of the 17th International Conference on World Wide Web (Beijing, China 21–25 April).
- Grönlund, A. (2001). Democracy in an IT-Framed Society, *Communications of the ACM* 44: 22–26.
- Habermas, J. (1984). *The Theory of Communicative Action, I*, Cambridge: Polity.
- Habermas, J. (1987). *The Theory of Communicative Action, II*, Cambridge: Polity.
- Hacker, K. and Van Dijk, J. (2000). *Digital Democracy: Issues of theory and practice*, London: Sage.
- Halavais, A.C. (2001). The Slashdot Effect: Analysis of a large-scale public conversation on the world wide web, Ph.D. Thesis submitted to the Department of Communication, University of Washington, Seattle, WA.
- Hirsch, J.E. (2005). An Index to Quantify an Individual's Scientific Research Output, *Proceedings of the National Academy of Sciences* 102(46): 16569–16572.
- Jolliffe, I.T. (2002). *Principal Component Analysis*, New York, NY: Springer.
- Kaltenbrunner, A., Gomez, V., Moghnieh, A., Meza, R., Blat, J. and Lopez, V. (2008). Homogeneous Temporal Activity Patterns in a Large Online Communication Space, *IADIS International Journal on WWW/INTERNET* 6(1): 61–76.

- Klofstad, C.A. (2007). Talk Leads to Recruitment. How Discussions about Politics and Current Events Increase Civic Participation, *Political Research Quarterly* 60(2): 180–191.
- La Due Lake, R. and Huckfeldt, R. (1998). Social Capital, Social Networks, and Political Participation, *Political Psychology* 19(3): 567–584.
- Lampe, C. and Resnick, P. (2004). Slash(dot) and Burn: Distributed moderation in a large online conversation space, in *CHI'04: Proceedings of the SIGCHI conference on human factors in computing systems*, New York, NY: USA, ACM Press.
- Landa, D. and Meirowitz, A. (2009). Game Theory, Information, and Deliberative Democracy, *American Journal of Political Science* 53(2): 427–444.
- Lazarsfeld, P., Berelson, B. and Gaudet, H. (1968). *The People's Choice: How the voter makes up his mind in a presidential campaign*, New York, NY: Columbia University Press.
- Mahrer, H. and Krimmer, R. (2005). Towards the Enhancement of E-Democracy: Identifying the notion of the 'Middleman Paradox', *Information Systems Journal* 15(1): 27–42.
- Malda, R. (1999). Slashdot moderation, <http://slashdot.org/moderation.shtml>.
- Manin, B. (1987). On Legitimacy and Political Deliberation, *Political Theory* 15(3): 338–368.
- McPherson, M., Smith-Lovin, L. and Brashears, M.E. (2006). Social Isolation in America: Changes in core discussion networks over two decades, *American Sociological Review* 71: 353–375.
- Mutz, D.C. (2002). The Consequences of Cross-Cutting Networks for Political Participation, *American Journal of Political Science* 46(4): 838–855.
- Mutz, D.C. (2006). *Hearing the Other Side: Deliberative versus participatory democracy*, Cambridge, MA: Cambridge University Press.
- Mutz, D.C. (2008). Is Deliberative Democracy a Falsifiable Theory? *Annual Review of Political Science* 11: 521–538.
- Pateman, C. (1970). *Participation and Democratic Theory*, Cambridge: Cambridge University Press.
- Paxton, P. (1999). Is Social Capital Declining in the United States? A Multiple Indicator Assessment, *American Journal of Sociology* 105(1): 88–127.
- Poor, N. (2005). Mechanisms of an Online Public Sphere: The website slashdot, *Journal of Computer Mediated Communication* 10(2), <http://jcmc.indiana.edu/vol10/issue2/poor.html>.
- Putnam, R.D. (2000). *Bowling Alone. The Collapse and Revival of American Community*, New York, NY: Simon and Schuster.
- Schkade, D., Sunstein, C. and Reid, H. (2007). What Happened on Deliberation Day? *California Law Review* 95: 915–940.
- Shane, P.M. (2004). *Democracy Online: The prospects for political renewal through the internet*, New York, NY: Routledge.
- Sunstein, C. (2007). *Republic.com 2.0*, Princeton, NJ: Princeton University Press.
- Thompson, D.F. (2008). Deliberative Democratic Theory and Empirical Political Science, *Annual Review of Political Science* 11: 497–520.
- Verba, S., Schlozman, K.L. and Brady, H.E. (1995). *Voice and Equality: Civic voluntarism in American politics*, Cambridge, MA: Harvard University Press.
- Zuckerman, A.S. (ed.) (2005). *The Social Logic of Politics: Personal networks as contexts for political behavior*, Philadelphia, PA: Temple University Press.

About the authors

Sandra Gonzalez-Bailon received a DPhil in Sociology from the University of Oxford in 2007. She is currently a Research Fellow at the Oxford Internet Institute and a member of Nuffield College. Her research explores the formation of politically relevant networks on the Internet and the mechanisms that explain individual contributions to the formation of online public goods. She is a co-convenor of the Nuffield-OII Networks Seminar Series and member of the Nuffield Network of Network Researchers (more information on her projects and publications can be found at <http://users.ox.ac.uk/~lady2042/>).

Andreas Kaltenbrunner received his Ph.D. from the University Pompeu Fabra (Barcelona, Spain) in Computer Science and Digital Communication in 2008, with a research topic on Social Media, and obtained a master degree in Applied Mathematics from the Vienna University of Technology (Austria) in 2000. Currently he works as senior researcher in the Information, Technology and Society Group of the Barcelona Media Innovation Centre where he performs empirical research on the characterisation and modelling of social networks and social media. His research interests include: neural networks, synchronisation, human communication, temporal and structural patterns in social and discussion networks, computational sociology (more information can be found at <http://www.dtic.upf.edu/~akalten/>).

Rafael E Banchs received his Ph.D. in Electrical Engineering from The University of Texas at Austin in 1998, and was awarded a 'Ramón y Cajal' fellowship from the Spanish Ministry of Education and Science in 2004. Currently, he works as a research scientist at the Speech and Language Research Group of Barcelona Media Innovation Centre, where his research activity is mainly focused on information retrieval and text mining technologies and their application to specific problems in the media industry and the Web. He has been author and co-author of more than 50 publications in international conferences and journals. He has also taught several undergraduate and graduate courses in different universities around the world (more information can be found at <http://varoitus.barcelona.media.org/rafael/index.html>).