Smartocracy:
Social Networks for Collective Decision Making

Marko A. Rodriguez[*], Daniel J. Steinbock[†], Jennifer H. Watkins[*], Carlos Gershenson[‡],
Johan Bollen[*], Victor Grey[§], and Brad deGraf[¶]

[*]Los Alamos National Laboratory, Los Alamos, New Mexico 87545
[†]Stanford University, Stanford, California 94305
[‡]Vrije Universiteit Brussel, Brussel, Belgium 1160
[§]2idi Corporation, Concord, California 94524
[¶]Media Venture Collective, Mill Valley, California 94941
Email: marko@lanl.gov

Abstract— Smartocracy is a social software system for collective decision making. The system is composed of a social network that links individuals to those they trust to make good decisions and a decision network that links individuals to their voted-on solutions. Such networks allow a variety of algorithms to convert the link choices made by individual participants into specific decision outcomes. Simply interpreting the linkages differently (e.g. ignoring trust links, or using them to weight an individual's vote) provides a variety of outcomes fit for different decision making scenarios. This paper will discuss the Smartocracy network data structures, the suite of collective decision making algorithms currently supported, and the results of two collective decisions regarding the design of the system.

I. INTRODUCTION

The recent explosion of so-called 'social software' on the Internet has been characterized by democratic approaches to content generation [1]. The Wiki is an exemplar of this approach, where all users have equal power to add or modify the content of any hypertext [2]. Wikis are social software that support democratic collaborative authorship. Systems exist to support generative collaboration in many fields including, but not limited to, journalism[1], scholarly citation[2], photography[3], and hypertext bookmarking[4].

While these systems vary in the forms of collaboration they support, they have in common an egalitarian social structure and they all aggregate user contributions into shared representations of collective belief. For instance, in addition to contributing news stories, users of the Digg web service vote for stories they think highly of and can view the most popular stories for different subjects. Del.icio.us users save and categorize bookmarks for websites they like into a common pool and can easily view the most popular sites for any category. Similar patterns of contribution and aggregation occur with varying prominence throughout the social software sphere. Despite the proliferation of such systems with traditional methods of aggregation, there has yet to emerge a generalized software model for the intelligent aggregation of individual contributions beyond mere vote-counting. If such a model did exist, it could systematically improve the state of the art in social software design and promote the innovation of systems geared more directly toward the aggregation of individual knowledge into collective knowledge, i.e. software supported collaborative problem solving and decision making [3], [4].

Smartocracy[5] is a web-based social software system for collective problem solving/decision making. Smartocracy uses a problem-solution model where individuals pose problems (i.e. issues, questions) to the Smartocracy community and propose potential solutions (i.e. options, answers). The proposed solutions are voted on and the aggregate preferences of all individual users yield the collective decision. To further facilitate vote-based decision making, a trust-based social network is used to represent the relations among users and, in some cases, to support the automated delegation of decision making power along paths of trust. This paper discusses trust-based decision making theory, the Smartocracy system implementation, and concludes with the results of the Smartocracy community's aggregate system development decisions as of March 2006.

II. COLLECTIVE DECISION MAKING

Collective decision making, in the context of Smartocracy, involves individuals generating problems, providing potential solutions, voting for solutions, and the software aggregating individual votes and ultimately deriving a final collective decision. Figure 1 provides an outline of Smartocracy's collective decision making process.

---

[1]Digg available at: http://digg.com
[2]CiteULike available at: http://citeulike.org
[3]Flickr available at: http://flickr.com
[4]Del.icio.us available at: http://del.icio.us
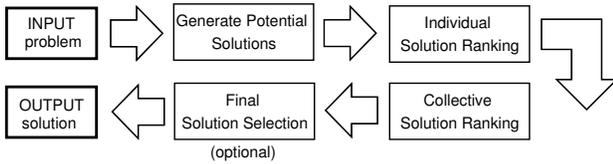
[5]Smartocracy available at: http://www.smartocracy.net

Fig. 1. Smatocracy's collective decision making stages

Given a particular problem, any individual in Smartocracy can propose a potential solution. Individuals can also vote for a proposed solution. By voting for particular solutions, individuals are explicitly providing their individual ranking of the problem's solution set. The suite of Smartocracy collective solution ranking algorithms aggregates all the individual votes into a collective solution ranking. Collective ranking algorithms currently supported by Smartocracy are direct democracy [5], dynamically distributed democracy [6], and proxy vote [7]. The result of the collective solution ranking algorithm serves as input to a solution selection function. For nominal-, or categorical-, based solution sets, the highest ranked solution is considered the collective decision. For numeric-, or gradient-, based solution sets, a weighted average selection function is calculated based on their respective rank. The result returned by the selection function is the collectively derived solution to the problem. It is important to note that the solution selection function is an optional step that is used only if a single definitive solution is required. In some cases, the solution to the problem may be the collective ranking itself.

III. THE NETWORKS

Social networks are used to represent the relationships between individuals of a population. Most commercial social network systems (e.g. MySpace[6] and LinkedIn[7]) denote static relational ties between users, such as 'is a friend' or 'is a colleague'. For collective decision making, friendship is not necessarily the desired semantic relationship. In collective decision making, stating that an individual is a friend does not identify them as a good decision maker [8]. Therefore, a collective decision making social network must make explicit the notion of decision making trust. For instance, if individual A provides individual B with a trust edge, then A is stating that he or she trusts B to make a good decision.

There exist other artifacts in the system besides human individuals. These artifacts include problems (i.e. current issues being addressed by the collective) and solutions (i.e. potential options for a particular problem). Therefore, the graph data structure underlying the Smartocracy system is a graph connecting individuals according to their relative trust of one another and connecting individuals to particular

solutions via edges that denote a vote. The network data structure describing this system can be formally denoted as $G = \{N, W\}$ where N is a collection of human (H), problem (P), and solution (S) nodes and W is the set of weighted semantic relationships between these entities. Note that $N = H \cup P \cup S$, $H \cap P \cap S = \varnothing$, and $w^{\lambda}_{n_i, n_j}$ states that there exists a directed, weighted relationship of semantic λ between n_i and n_j. The only semantic relationships currently supported by Smartocracy are $\lambda = \{\text{trusts}, \text{votedFor}\}$.

The connection between the trust-based social network and the vote-based decision network is represented in Figure 2 where the solid edges represent the weighted trust relations amongst individuals and the dashed edges represent the weighted votes that individuals provide for particular solutions to problems. The following two sections will further discuss each network's constructs.
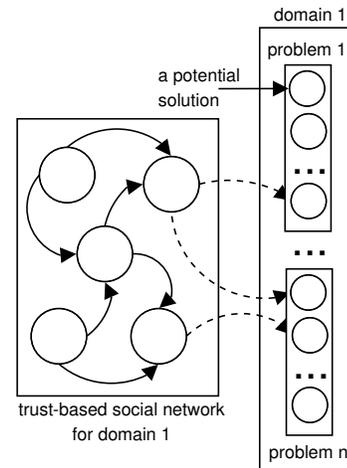

Fig. 2. The connection between the social and decision networks

A. Trust-Based Social Network

In the Smartocracy trust-based social network, an edge going from individual $h_i \in H$ to $h_j \in H$ states that h_i believes that individual h_j is 'good' at decision making where $i \neq j$ and $H \subseteq N$. In this model, h_j may not trust the decision making behavior of h_i. Therefore, the trust-based edges between individuals are directional. Furthermore, individual h_i may weight the relative trustworthiness of individual h_j and, in such cases, an edge can be represented by the conditional probability,

$$w^{\text{trusts}}_{h_i, h_j} = P(h_j \text{ is good} \mid h_i\text{'s knowledge of } h_j).$$

This relationship states that h_i trusts, according to some probability, that h_j will make a good decision given h_i's previous understanding of h_j. What this edge model lacks is a representation of the domain for which trust is given. Because humans are multi-dimensional entities composed of various

skills and beliefs, individual h_i may trust the decision making behavior of h_j in one domain but not another. Therefore, it may be more appropriate, given a decision making scenario, to state,

$$w_{h_i,h_j}^{\text{trusts}} =$$
$$P(h_j \text{ is good in domain } d_l$$
$$\mid h_i\text{'s knowledge of } h_j \text{ in domain } d_l).$$

There are multiple ways to represent domains in a social network context [7]. Currently, due to the focused use of Smartocracy, there exists only a single social network and therefore, no explicit representation of domains. As will be demonstrated in the results section, domain specific representation is a collectively desired feature for future implementation.

B. Vote-Based Decision Network

In Smartocracy, any individual that votes for a potential solution is creating a directed edge from themselves in H to a solution in S. The semantic of this relationship is $\lambda = \text{votedFor}$ and can be represented by the conditional probability,

$$w_{h_i,p_j(s_m)}^{\text{votedFor}} =$$
$$P(s_m \text{ is a good solution for problem } p_j$$
$$\mid h_i\text{'s knowledge of } p_j).$$

The above vote edge states that human h_i believes that, according to some probability, s_m is a good solution for problem p_j.

The weighted ranking that an individual provides for the solution set to a particular problem is their subjective evaluation of the relative optimality of the solutions for the problem. These are called the individual solution rankings. To move from an individual solution ranking to a collective solution ranking, an aggregation algorithm is required. The next section will discuss the algorithmic framework used to calculate collective decisions within the trust-based social network and the vote-based decision network.

IV. ALGORITHMS FOR COLLECTIVE SOLUTION RANKING

Given a trust-based social network and a vote-based decision network, Smartocracy supports a family of algorithms for aggregating individual votes into a collective solution ranking. The currently implemented algorithms are direct democracy, dynamically distributed democracy, and proxy vote. All of these algorithms are implemented under the parameterized particle swarm framework described in [7], [9] and are similar to the idea of constrained spreading activation [10].

In a particle swarm, a particle is considered an 'atom' of decision making influence. To calculate a collective solution ranking for a particular problem, particles begin their journey at human nodes and make their way, in a stochastic manner, to the solution nodes of that problem via the trust-based social network and vote-based decision network. Because particles are discrete, indivisible entities, the diffusion of particles through a network requires a sufficient initial distribution to expose the underlying network topology. The more particles initially supplied to the network, the more accurate the collective ranking [9]. At the end of the particle propagation algorithm, when all particles have either been destroyed or have reached a solution, the distribution of particles over the solution set of the problem determines the collective's solution ranking. This section will describe the specifics of each of the three collective solution ranking algorithms currently supported by Smartocracy.

A. Direct Democracy

Direct democracy embodies the idea of 'one-person/one-vote' [5]. In direct democracy, the trust-based social network is *not* used to calculate the collective decision. If the individual does not vote, then the individual does not participate. If an individual does not participate, then he or she does not influence the collective solution ranking. To implement this algorithm within the particle swarm framework, each individual is supplied with 100 particles. A particle can only traverse a votedFor edge, $w_{h_i,p_j(s_m)}^{\text{votedFor}}$. The probability of traversing one votedFor edge over another is dependent upon the edge weight assigned by the voter. Strongly weighted edges have a higher probability of being traversed by the particle. If the individual has not voted, then the particle destroys itself. After one step, all particles are either destroyed or are at a particular solution node to problem p_j. The distribution of particles over the solution nodes represents the direct democracy collective solution ranking.

Figure 3 provides an example of the possible paths of a direct democracy swarm. Humans h_1, h_2, and h_3 are each provided 100 particles. After the first step, because particle diffusion is a stochastic process biased by edge weights, s_1 will have accumulated approximately 50 particles, s_2 60 particles, and s_3 90 particles. Note that the particles given to h_2 are destroyed because h_2 has not voted on a solution. The normalized distribution over the solution set is $s_1 = 0.25$, $s_2 = 0.30$, and $s_3 = 0.45$. This normalized distribution is the direct democracy collective solution ranking. If the solution set is a nominal solution set, then s_3 is the collectively derived solution according to a highest rank solution function.

B. Dynamically Distributed Democracy

Dynamically distributed democracy, or DDD, was developed to handle fluctuating levels of participation. In
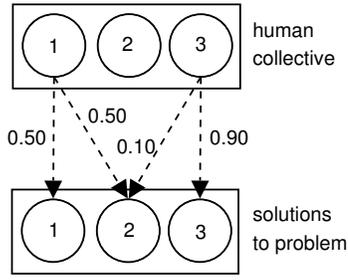
Fig. 3.   Direct democracy particle paths

such cases, ad hoc representative structures are created to ensure that every individual can influence the collective solution ranking even if only through a proxy representative. This algorithm has been shown to be an accurate way to model the collective's perspective as voter participation wanes [6].

In DDD, a particle, if it is unable to take a vote edge to a particular solution, uses the trust-based social network to move to a proxy representative. If that representative has voted, then the particle moves to one of the representative's chosen solutions. If the representative has not voted, the particle traverses a trust edge to move to yet another representative. This iterative process continues until a solution to the problem is found.

In Figure 4, human h_2 has not voted on a particular solution. In DDD, h_2's decision making influence (in the form of particles) is delegated to human h_1 because h_2 trusts h_1, $\exists w_{h_2,h_1}^{\text{trusts}}$. Therefore, h_1 has 200 particles to provide to his or her chosen solutions. After two steps, solution s_1 will have 100 particles, s_2 will have 110 particles, and s_3 will have 90 particles. The normalized distribution over the solution set is $s_1 = 0.33$, $s_2 = 0.36$, and $s_3 = 0.30$. If the selection function selects the highest ranked solution, then solution s_2 would be the collectively derived solution.
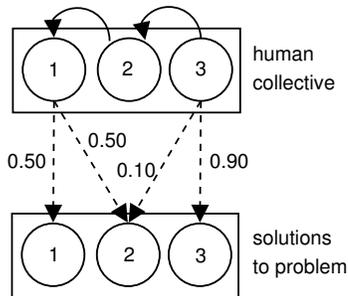

Fig. 4.   Dynamically distributed democracy particle paths

Notice that because h_3 provided a vote, the particles initially provided to h_3 do not traverse the social network. That is, $w_{h_3,h_2}^{\text{trusts}}$ is not used in the DDD computation.

Furthermore, in the case that every individual votes for a particular solution, the DDD algorithm will provide a collective solution ranking that is equivalent to the direct democracy algorithm's ranking. DDD is useful in problem domains where it is desirable for every individual to have an equal say (i.e. influence) in the collective's decision, but it is not feasible for every individual to actively participate via voting. In such cases of reduced participation, ad hoc representative structures emerge to simulate full participation.

C. Proxy Vote

Proxy vote is an extension of DDD where the initial distribution of particles is biased by the trust-based in-degree of an individual. That is, the more a particular individual is trusted, the more particles that individual initially receives. This algorithm is used for expert-based problem domains where an equal say for every individual is not desired. For proxy vote, it is assumed that the in-degree of the individual is a representation of his or her level of expertise.

In Figure 5, human h_1 is supplied with 200 particles, h_2 200 particles, and h_3 100 particles. The number of particles at h_1 and h_2 is double that of h_3 because both h_1 and h_2 have one incoming trust-based edge and each incoming edge provides an extra 100 particles to an individual. As in the DDD example, h_2 disseminates his or her received particles to h_1 because h_2 trusts only h_1 and h_2 has not voted on a solution to the problem. Therefore, h_1 has 400 particles to distribute to his or her chosen solutions. After the second step, solution s_1 will have 200 particles, s_2 210 particles, and s_3 will have 90 particles. The normalized distribution over the solution set is $s_1 = 0.40$, $s_2 = 0.42$, and $s_3 = 0.18$. Given a highest rank selection function, s_2 would be the collective decision.
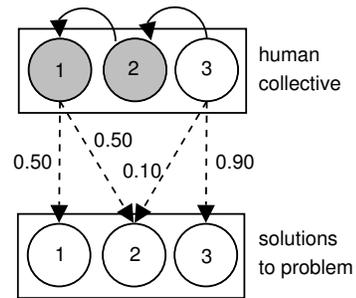

Fig. 5.   Proxy vote particle paths

V. SMARTOCRACY BETA RESULTS

Smartocracy was released for beta testing in February of 2006. New users join the system via an invitation from a currently existing user. An invitation of someone new automatically creates a directed edge from the inviter to the invitee. Throughout the life of the system, individuals can

reassign their outgoing trust-based edges. The interface to view one's incoming and edit one's outgoing trust edges is shown in Figure 6.
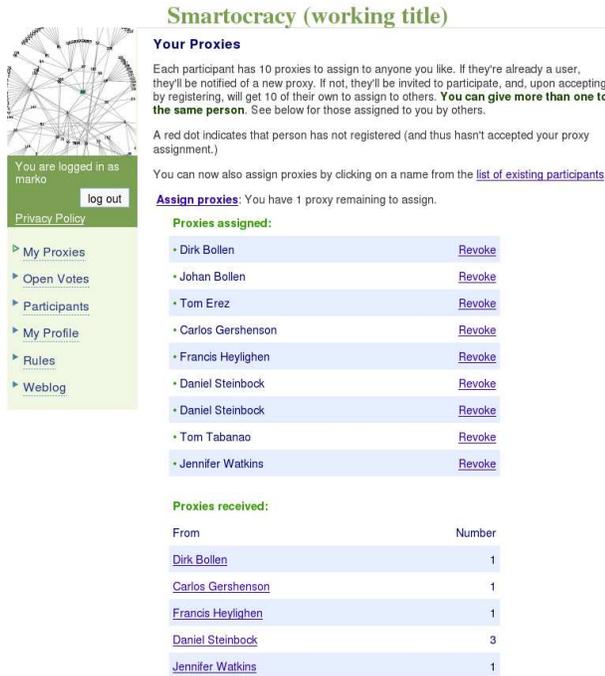

Fig. 6.   An individual's in and out trust-based edges

Figure 7 provides a visualization of the 276 participants in the Smartocracy trust-based social network as of March 7, 2006.
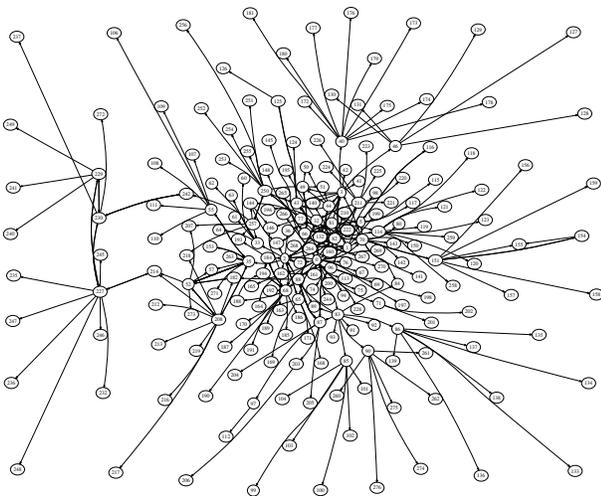

Fig. 7.   Visualization of the Smartocracy network as of March 7, 2006

The current in-degree edge distribution is presented on a log/log plot in Figure 8. Like most real world networks,

the Smartocracy social network's edge distribution follows a power-law. For the function $P(\text{in-degree}) = a \times \text{in-degree}^{-\gamma}$, $\gamma \approx 2.4$. This function computes the probability that any given node will have a particular in-degree. As can be seen in Figure 8, there exist many individuals with few incoming trust edges and few individuals with many incoming edges. As depicted in Figure 8, more than 75% of all individuals in the network have 1 incoming edge while only 0.4% of the individuals have 8 incoming edges.
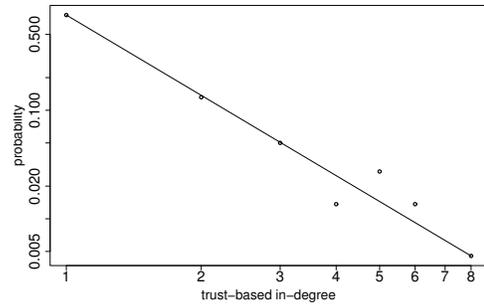

Fig. 8.   Smartocracy social network in-degree edge distribution

The remainder of this section will provide some preliminary results of two initial problems posed to the Smartocracy collective. The two problems, in the form of questions, are:

1) What should the name of the system be?
2) What features should we add to the system?

These problems are self-reflective, as they concern the system itself. The goal of the beta test period is to develop a system that is palatable to the collective. For this reason, it is thought best to allow the users to provide development solutions and thus regulate the system's evolution via the collective. It is important to note that the results of the various collective solution ranking algorithms differ. We are not attempting to state whether one algorithm is more optimal than another, only to demonstrate that different metrics for collective decision making exist and that it is important to study in which contexts which algorithms are most appropriate.

A. System Name

The system name problem was posed by the developers in order to determine the most appropriate name for the service. Any individual could provide a potential name solution as well as vote on which name they preferred. Table I provides the 13 potential names that were provided by the group.

A visualization of the vote interface of the system is provided in Figure 9. Between the time the initial decision making data was analyzed and the screenshot was taken, a

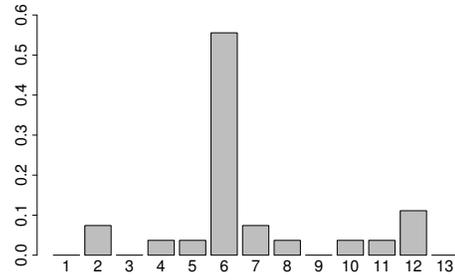|	potential solution
1	DecisionNet
2	DDD
3	Dis Dis Sys
4	Cell of the Elite
5	Smart Mob Rule
6	Smartocracy
7	Holocracy
8	netocracy
9	Wisism
10	Stupocracy
11	Cheerocracy
12	antpile
13	Decision Network

TABLE I

SYSTEM NAME SOLUTIONS

Fig. 10. Direct democracy system name ranking

new system name, Dynamocracy, was proposed. Note that the new solution expresses the dynamic nature captured in the difference between Table I and Figure 9.

Fig. 9. Creating and voting on solutions screenshot

Figure 10 provides the ranking of the 13 potential solutions according to the direct democracy algorithm. Of the 276 participants in the system, 95 voted. Therefore, only 34% of the group contributed to the direct democracy collective solution ranking. With a highest rank selection function, the name Smartocracy, is the final collective decision.

Figure 11 provides the ranking of the 13 potential solutions according to the dynamically distributed democracy algorithm. In DDD, the 181 non-participants delegated their decision making influence (particles) to their trusted proxy representatives. Given a highest rank selection function, Smartocracy is again the collective decision.
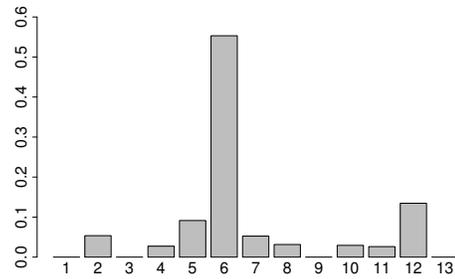

Fig. 11. Dynamically distributed democracy system name ranking

Finally, Figure 12 provides the ranking of the 13 potential solutions according to the proxy vote algorithm. Smartocracy is once again the highest ranked solution.
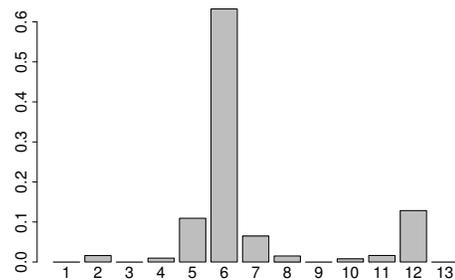

Fig. 12. Proxy vote system name ranking

Table II provides the Spearman ρ correlations between the various algorithms' collective solution rank distributions where dd is direct democracy, ddd is dynamically distributed democracy, and pv is proxy vote. The strong correlations, with p-value $< 2.2^{-16}$, between the algorithms demonstrates the definitiveness of the collective's desired name choice.

For this reason, Smartocracy is the current name of the system.

	dd	ddd	pv
dd	1.0	0.93	0.91
ddd	0.93	1.0	0.97
pv	0.91	0.97	1.0

TABLE II

SYSTEM NAME SPEARMAN ρ RANK CORRELATIONS

B. Future Features

In line with today's widely-used strategies for social software development, Smartocracy was originally implemented in a bare-bones fashion with the desire that the group steer its future development. With this idea in mind, the problem "What features should we add to the system?" was posed to the group. Table III provides the list of all proposed solutions generated by members of the group.

	potential solution
1	social network visualization
2	know if my proxies have taken action
3	domain specific proxies
4	Golightly Community
5	dialog with other participants
6	geo-based network viz

TABLE III

FUTURE FEATURES SOLUTIONS

Of the 276 participants in the system, 98 voted for a future feature. Figure 13 depicts the direct democracy collective solution ranking over the 6 proposed solutions. According to direct democracy with a highest rank selection function, the collective desires to have a feature that allows individuals to know if their representatives, or proxies, have voted for a solution to a problem.
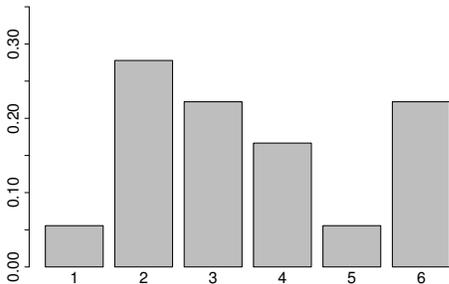
Fig. 13. Direct democracy collective solution ranking for features

In Figure 14, the DDD collective solution ranking provides a different perspective of the collective's desired future feature. Smartocracy was originally developed with a single trust-based social network. Therefore, decision making influence is delegated in the system irrespective of the domain of the problem. It became apparent to many users that they may trust someone in one domain, but not another. The DDD/highest rank decision shows that the collective is more interested in having domain specific proxies, or domain specific representatives. That is, decision making influence is delegated to representatives depending on the domain of the problem and the domain of the trust-based representation.
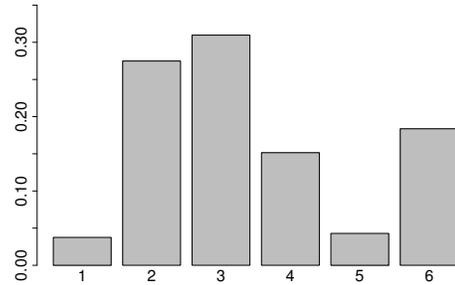
Fig. 14. Dynamically distributed democracy collective solution ranking for features

Finally, the proxy vote solution ranking is perfectly correlated with the DDD solution ranking and therefore, proxy vote shows that the collective is interested in having domain specific proxies. The proxy vote collective solution ranking is provided in Figure 15.
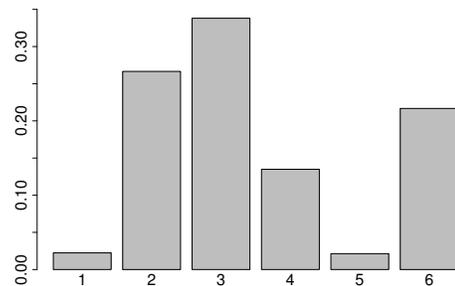
Fig. 15. Proxy vote collective solution ranking for features

The Spearman ρ correlations between the various algorithm's collective solution rankings are provided in Table IV where p-value $< 2.2^{-16}$. Given the perfect correlation between DDD and proxy vote, the developers of Smartocracy are currently working towards implementing domains into the system as articulated in [7].

	dd	ddd	pv
dd	1.0	0.88	0.88
ddd	0.88	1.0	1.0
pv	0.88	1.0	1.0

TABLE IV

FUTURE FEATURES ALGORITHM SPEARMAN ρ CORRELATIONS

VI. FUTURE DIRECTIONS

In spirit with the previously presented collective decision making processes, future development of the system will be guided by the requirements generated by the collective. As shown in the future features section, a major push in the development of this system will be the inclusion of domain specific representation. In [7], a model for adding domain specific representation to social networks is provided. Depending on the requirements of the Smartocracy group, one of the proposed models will be implemented. Furthermore, the current platform can be extended to support collaborative discourse as presented in [4], [11]. Collaborative discourse promotes interaction amongst the participants by means of argumentation.

Finally, in terms of future research in the area of collective decision making, we intend to explore various collective ranking algorithms for different classes of problems to understand which algorithms are best suited for particular decision making situations. Given a problem set with known optimal solutions, which algorithms best aggregate the perspectives of the individuals to yield the optimal solution? Such a study would expand the use of Smartocracy into the domain of collective intelligence research and would bring Smartocracy to a level that is a generalization of domain specific human collective intelligence systems like [4], [12]–[14].

VII. CONCLUSION

Social systems, in order to make the 'right' decisions, require sensors capable of delivering relevant information. Polls have been used to collect information relevant for societal-scale decision making. The drawbacks of polls are well known: either they require public participation from a large sector of the population, or they might be biased and therefore, not representative of the collective opinion. By exploiting social network structures, systems such as the one presented here can overcome some of these drawbacks. With many collective solution ranking algorithms, other than just direct democracy, it is possible to learn how these metrics effect the collective solution outcome. In turn, over time it may be possible to understand in which context which algorithm is best.

Smartocracy could be used to make social decisions within communities of any type, from frivolous to scientific, from political to educational: scientific peer review, project funding, school admissions, popularity contests, artistic awards, ethical decisions, judicial sentences. These are only some of the potential niches where the research presented in this paper could be exploited.

REFERENCES

[1] M. Tepper, "The rise of social software," *netWorker*, vol. 7, no. 3, pp. 18–23, 2003.

[2] B. Leuf and W. Cunningham, *The Wiki way: quick collaboration on the Web*. Addison-Wesley Longman Publishing Co., 2001.

[3] F. Heylighen, "Collective intelligence and its implementation on the web: Algorithms to develop a collective mental map," *Computational & Mathematical Organization Theory*, vol. 5, no. 3, pp. 253–280, 1999.

[4] M. Turoff, S. Hiltz, H. Cho, Z. Li, and Y. Wang, "Social decision support systems (SDSS)," in *35th Annual Hawaii International Conference on System Sciences (HICSS'02)*, vol. 1, Hawaii, USA, 2002, p. 11.

[5] A. Lijphart, *Democracies: Patterns of Majoritarian and Consensus Government in Twenty-One Countries*. Yale University Press, 1984.

[6] M. A. Rodriguez and D. Steinbock, "A social network for societal-scale decision-making systems," in *NAACSOS '04: Proceedingss of the North American Association for Computational Social and Organizational Science Conference*, Pittsburgh, PA, USA, 2004. [Online]. Available: http://arxiv.org/abs/cs.CY/0412047

[7] M. A. Rodriguez and D. J. Steinbock, "The anatomy of a large scale collective decision making system," Los Alamos National Laboratory Technical Report [LA-UR-06-2139], Tech. Rep., March 2006.

[8] C.-N. Ziegler and J. Golbeck, "Investigating correlations of trust and interest similarity - do birds of a feather really flock together?" *Decision Support Systems*, vol. [in press], 2006.

[9] M. A. Rodriguez and J. Bollen, "Simulating network influence algorithms using particle-swarms: Pagerank and pagerank-priors," *[submitted]*, 2005. [Online]. Available: http://arxiv.org/abs/cs.DS/0602002

[10] F. Crestani and P. L. Lee, "Searching the web by constrained spreading activation," *Information Processing and Management*, vol. 36, no. 4, pp. 585–605, 2000.

[11] M. Turoff, S. R. Hiltz, M. Bieber, J. Fejermestad, and A. Rana, "Collaborative discourse structures in computer mediated group communications," in *32nd Annual Hawaii International Conference on System Sciences (HICSS'02)*, vol. 1, Hawaii, USA, 1999.

[12] N. Johnson, S. Rasmussen, C. Joslyn, L. Rocha, S. Smith, and M. Kantor, "Symbiotic intelligence: self-organizing knowledge on distributed networks driven by human interaction," in *6th International Conference on Artificial Life*, 1998.

[13] C. A. Kaplan, "Collective intelligence: A new approach to stock price forecasting," in *Proceedings of the 2001 IEEE Systems, Man, and Cybernetics Conference*, New York, USA, 2001.

[14] D. Steinbock, C. Kaplan, M. A. Rodriguez, J. Diaz, N. Der, and S. Garcia, "Collective intelligence quanitifed for computer-mediated group problem solving," University of California at Santa Cruz, Tech. Rep., 2004. [Online]. Available: http://arxiv.org/abs/cs.CY/0412064