



HAL
open science

Don't Vote, Evolve!

Pietro Sperini Di Fenizio, Derek Paterson

► **To cite this version:**

Pietro Sperini Di Fenizio, Derek Paterson. Don't Vote, Evolve!. Second IFIP WG 8.5 International Conference on Electronic Participation (EPART), Aug 2010, Lausanne, Italy. pp.13-25, 10.1007/978-3-642-15158-3_2 . hal-01059026

HAL Id: hal-01059026

<https://inria.hal.science/hal-01059026v1>

Submitted on 29 Aug 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Don't Vote, Evolve!

Pietro Speroni di Fenizio and Derek Paterson

CISUC, Department of Informatics Engineering,
University of Coimbra, Coimbra, Portugal

speroni@dei.uc.pt

cellers@gmx.com

<http://cisucpt.dei.uc.pt/ecos/>

Abstract. We present an alternative form of decision making designed using a Human Based Genetic Algorithms. The algorithm permits the participants to tackle open questions, by letting all of them propose answers and evaluate each others answers. A successful example is described and some theoretical results are presented showing how the system scales up.

Keywords: e-Democracy, Voting Theory, Pareto Front, Genetic Algorithms, Multi-Criterion Decision Making

1 Introduction

The word “democracy” comes from the Greek “dēmokratia”, “dēmos” meaning “the people” and “-kratia” meaning “power, rule”. It literally means “rule of the people” [1]. The word has come to mean a government where either directly or indirectly the citizens decide, through voting, what should be done. Is this the only option? As technology advances the possibilities with which the core idea of democracy can be implemented, expand. It becomes then necessary to investigate, also formally, these new possibilities. Voting Theory is the branch of Mathematics which studies the different ways in which people can express their preferences, and how these preferences can be integrated to reach a final result [2]. But Voting Theory is based upon the assumption that we know what the options are, and inevitably the result of a vote is always selecting one alternative, among the many. And leave those who did not like the winning alternative unhappy with the result.

The amount of information that the citizens send to the government through voting each election is incredibly low. If we are able to chose between m candidates, each citizen is sending $\log_2(m)$ bits of information every few years. How could a government be able to represent correctly the desire of the citizens with such a tiny amount of information? But this represented correctly the amount of information that could be exchanged in the 18th century, when travel was done on horse, and modern democracies were designed. Now we are able to transmit a much broader amount of information. Apart from voting for our representatives, we have many other tools: from polls, to debates on television, to blogs, discussion boards, and wikis. Recently governments around the world have started to

open up their data and become transparent to the scrutiny of the public [3, 4]. Also some government have started posting online the laws they are about to implement, letting people comment on them [5]. All this is positive and permits a much broader integration between the representatives and the represented.

But the politicians can ignore all this, and still the only moment when people are directly affecting their government is during an election. And, although the tool might have in the meantime progressed from a paper ballot into an electronic form of voting, the essence of it hasn't changed. We are still only sending $\log_2(m)$ bits of information every election.

How could the citizen more directly participate in the governance process? We are not suggesting here a form of direct democracy, in fact quite the opposite; what we are trying to suggest in this paper is that our ability to design the democracy of the 21st century suffers an endemic lack of imagination. Now that we have a tool, the internet, that permits an instantaneous, and cheap, transfer of information between the citizens and the governing body, we don't know how to use it. We should go back to the basic idea behind *dēmos-kratia*, and rethink the whole process through. What this paper wants to suggest is a fundamentally new system. This by introducing two ideas which are not new in science, but are new in political science. . .

2 Evolving and Selecting Ideas.

2.1 Evolving

Our basic critique of Voting Theory is that by assuming that the alternatives are easy to find out sets up the stage for its too limited results. Not only are the people who decide these alternatives given an unbalanced amount of power with respect to the voters, but also the best alternative that can be reached is one which has been decided beforehand by them. What if there existed an alternative solution to the problem at hand which has not yet been considered? A solution which could aggregate a wide majority, or even a consensus. Finding and agreeing on such solution would be a worthwhile goal of the integrated system we are designing.

To do this we need to reframe the basic problem. Instead of finding the alternative among a finite set of options (thus asking a *closed* question: "which among those alternatives would you prefer?"), we should look for the answer to an *open* question: "what is the solution to this problem...?". When we reframe our question in this way, Voting Theory is not anymore the field to apply; we should, instead, use Genetic Algorithms[6]. Genetic Algorithms (GA), which are both a tool and the field that studies them, aim to look at the best solution in a space of potential alternative solutions. This space can be very wide, sometimes infinite, and it is assumed that not all possible solutions can be tested. The solution, instead, is found through further approximations. Each solution can be tested at a small (but never null) cost. Through those tests the algorithm receives a feedback on how well each solution satisfies the requirements that is trying to

meet (i.e. technically it calculates the *fitness* of the solution). Solutions that behave poorly (have a low fitness) are discarded; good solutions (solutions with a high fitness) are retained and modified (technically *mutated* and *recombined* - in the interest of brevity we will consider both operations here as *mutations*), to generate new possible solutions. Two elements are important in this process: the production and the evaluation of the solutions. Usually in GAs the evaluation and the mutations are done in an automatic way. Automatically calculating the fitness of a solution to a question in a political context is not just an unsolved problem, but probably an unsolvable one too. Similarly given a solution, finding alternative similar solutions is equally challenging. This is maybe why no one has suggested using GA in a political context. Recently a new type of Genetic Algorithm was proposed: a Human Based Genetic Algorithm [7][8][9]. In such algorithm humans beings provide the fitness of a solution, and given a set of solutions provide alternative solutions. For such algorithm to succeed we need a semi-continuous dialogue between the participants and the computer. The computer asking for possible solutions from each participant, then, after collecting those solutions, presenting them to the participants again, asking them to evaluate them. This evaluation is then used by the computer to assess the fitness of the existing solutions, to decide which solutions should be discarded and which should be kept. The surviving solutions are then fed back to the participants, asking them to produce new solutions, mutating these surviving ones. To establish such process in a political context we would need a strong participation between the citizen and the government. Exactly what the internet permits, and what many citizens are now demanding!

To test these ideas the authors have set up a website where users can ask questions, and participate in finding the solution to the questions asked. We are going to show some of the results later.

2.2 Selecting

Which solutions should be maintained from one generation to the next? The trivial (and sub-optimal!) strategy would be to define how many solutions should be kept, in advance, and then sum up the votes that each solution receives. The top n solutions would then be copied to the next generation. We can either let people vote for one solution each, or vote for all the solutions they support. A number of problems arise with either of these strategies. If the participants are allowed to write their own proposals, and only vote for one solution, the temptation will be high for each person to just vote for their own solution. Also if the solution a person prefers has a low probability to be selected, a person might switch their vote to a less desired outcome, but an outcome where his vote has a higher chance to be useful. This leads to the behavior of *strategic voting*, voting in a way that does not represents the voter real desires, to let an acceptable outcome be the result of the voting process. A better solution is to let everybody vote for all the solutions they agree to, a form of voting called *approval voting*[10]. The problem with this form of voting is that if we let everybody vote for each possible solution, then similar solutions will be supported by a group

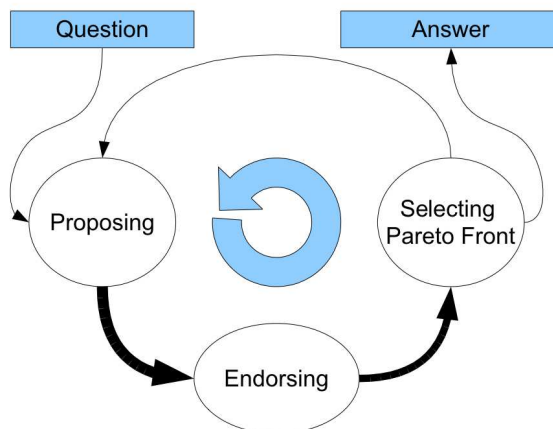


Fig. 1. Scheme of the genetic algorithm procedure. A question is asked. All the participants propose their answer. Then all the participants endorse the answers they support. Out of those proposals the Pareto Front is selected, and extracted. If those proposals achieve consensus, we have our answer. If not they are presented back to the community, asking the participants to be inspired by them to produce new proposals.

of people of similar size (often comprised of the same people). Thus if we take the n most popular solutions, they will probably be very similar, representing the same group of people. Instead people who did not agree with one of them might not agree with the others as well. Thus the resulting set of solutions will not represent the whole community. Instead a sub-community will be over-represented while other participants will be totally ignored. In Voting Theory this is often considered an unavoidable evil. But if we are not aiming to find which among a set of choices is better, but we want to set up a process that eventually leads us to find a proposal that has the widest possible support, ignoring some participants means to lose important information.

The logic behind the idea of summing up the votes is that each person is equally important. But when we are summing up the votes we are also losing information. If we have n people selecting the proposal to support among m alternatives, each of them sends 1 message among 2^m possible, thus each citizen sends $\log_2(2^m) = m$ bits of information. And considering that we have n citizens that are voting we are sending $n * m$ bits of information. But how much information is received? When we sum up the votes we end up knowing for each proposal how many people approve it. And since we are permitting each person to vote for multiple proposals, then the amount of information that we have received for each proposal will be $\log_2(n)$. Considering that we know this for each of the m proposals, the total amount of information received will be $m \log_2(n)$. In other words by summing up the votes we drop information from $n * m$ to $m \log_2(n)$. What went wrong? By summing up the votes we are ignoring any information which correlates one proposal to another. In other words we ignore information

of the kind: everybody who voted for proposal A , also voted for proposal B . But what would be the alternative?

Instead of using as fitness the sum of the votes that each proposal has received, we could consider fitness as a point in a multidimensional space. The number of dimensions will be the number of participants. The fitness will then be a boolean vector of size n , which we shall call the voting vector of the proposal. Such vector will have a 1 in the t position if the t participant has supported the proposal, and a 0 if not. A proposal that has reached consensus will be represented by a vector of all 1s. Then instead of selecting the proposal that has more votes (more 1s) we take the *Pareto Front* in this multidimensional space.

Pareto Fronts (PF) are particular subsets that are calculated starting from the concept of dominance. PF have been used in economy and in Multi-Criterion Decision Making[11]. Recently they have also been suggested as a tool in genetic algorithms and co-evolutionary learning[12]. To take the Pareto Front we must first define a function of dominance. We shall say that a proposal A *dominates* a proposal B if they have different voting vectors $\mathbf{v}_A \neq \mathbf{v}_B$, and in every dimension t , $v_A^t \geq v_B^t$. Since we have chosen to use only boolean vectors this is equivalent to saying that the set of people that support B must be a proper subset of the set of people that support A . Note that the procedure can be generalised, permitting each participant to evaluate each proposal by ranking it or evaluating each proposal by assigning it an integer. The Pareto Front will be the set of proposals that are not being dominated by any other proposal.

Let's look at some of the consequences of our selection choice: (1) the most popular proposal will always be in the Pareto Front; (2) for each participant at least one of the proposal he supports must be in the Pareto Front; as a corollary of (2) we have (3) each user who wants to have some particular proposal present in the Pareto Front must simply vote for only that proposal; but, vice-versa, (4) if a proposal is in the Pareto Front and is supported by only one person, that person cannot support any other proposal that is supported by any other participant. In other words nothing stops users to only vote for their own proposals to make sure they make it to the Pareto Front. But this asocial behavior will not go unspotted. Sometimes it is acceptable, (for example in the rare situation when a person really believes that their proposal is the only acceptable solution), but if repeated in time can lead to the other people ignoring that person and trying to reach a consensus without her.

Also note that we did not need to set the number of proposals that would be maintained from one generation to the other. This number is implicit in the definition of Pareto Front and will change from generation to generation. The only time a single proposal is selected is when everybody has voted for it. On the other hand, it is possible to have multiple solutions that all reach consensus if everybody votes for them. While keeping the Pareto Front unbounded might seem dangerous (what if too many proposals reach the next generation? Will the system explode?), recent results suggest that although no apparent limit is present, as long as the people vote following their values (and not in a random

way), the Pareto Front will still be bounded. More about this in the Results section.

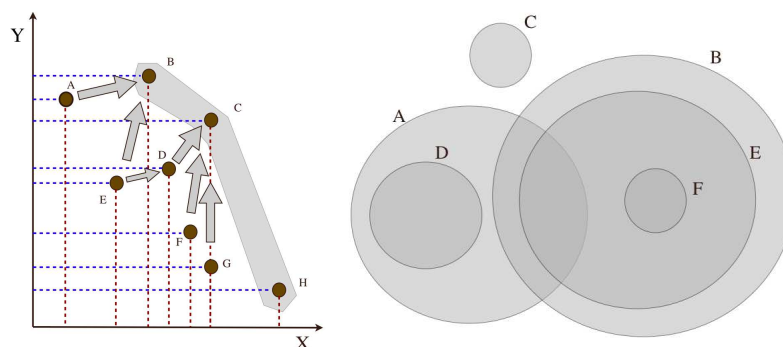


Fig. 2. **Left** Domination and Pareto Front in a 2 dimensional space. Each dot represents a possible solution. X and Y can either represent the Subjective Evaluation of those solutions by two participants or their Objective Evaluation (like it is being done in Multi-Criterion Decision Making[11, 15]). The Solutions B, C and H are not dominated by any other solution and thus represent the Pareto Front. **Right.** If the evaluation becomes boolean (support/ignore), then each solution can be represented as a set. Dominance becomes set inclusion, and the Pareto Front becomes the set of sets that are non dominated (in the image the sets A B, and C).

3 Testing the Idea

Applying those two ideas we coded a website, Vilfredo goes to Athens (available at <http://vilfredo.org>), to test them. The website permits the users to ask questions and participate in evolving proposals to the questions asked. From a formal point of view, once a question is asked it goes through a series of rounds called generations. Each generation is divided into a proposing phase and an evaluation phase. In the proposing phase the users are allowed to write their own proposals. In the evaluation phase they can evaluate all the proposals that have been submitted. This is done by clicking a button next to all the proposal each user likes, and ignoring the others. Thus there is no “I do not like” button. Ignored proposals are considered having being rejected. The Pareto Front of the proposal (as described above) is then extracted and presented to the users as representative of the whole community. If a consensus has been reached it is declared and the question is moved into the “solved questions” area. If instead a consensus has not been reached the users are invited to read the Pareto Front proposals and write new ones. In particular they are invited to try to bridge different ideas, and rewrite existent proposals, especially proposals they could not support, in a way they would be acceptable for them. Note that no further

aid or limits are given to the participants, other than presenting them with a blank slate and the list of the previous Pareto Front. They can ignore those suggestion, as well as use them.

Testing the website we soon discovered that it was necessary to fine tune the process to allow the procedure to work better. We shall list here the elements which we found important, with an explanation on how they affected the procedure.

- During the proposing phase, proposals being written are not revealed to the participants. This prevents users from copying each other's ideas. While copying each other's ideas could be seen as a positive element, we want to make sure that the ideas in the Pareto Front are more likely to be used than others.
- During the evaluation phase, proposals are presented in an anonymous format. No information about the author is presented. This forces users to rely on the actual content of a proposal, without being able to support something out of trust with the author. And, more importantly, this gives the possibility for everybody to have their proposal evaluated on an equal ground.
- When the evaluation is over, all the information about the authorship and who supported which proposal is revealed. This permits users to see which participants are only voting for their own proposals. Interestingly, we observed that this behavior drastically diminished after this information was made public.

Soon the website was used to bootstrap its own creation, with participants discussing, proposing and voting on the features of the website itself. Although the website was public, it did not attract a huge attention, and as such the number of users that have registered is around a hundred (137 on the 16 of May 2010) , and the number of different users that were participating at any one period of time, was less than 20. Many of the question did not attract a wide participation, ending up with only 2 or 3 participants agreeing on an answer. Discussing why this is so, goes beyond the scope of this paper. And of course further tests should be completed, possibly with the aid of workers in the psychology and sociological field to ascertain if the participants found the process satisfactory. For now we will instead focus on one the most popular question that was asked. The question was open to the public, but it was not publicised, which led to a very little turn over of participants during the process.

3.1 The Wall of Text Question

We shall now present an example of a question which was discussed: 'How should we handle the "wall of text" problem?'. The wall of text was the situation where a user proposal was so long that would slow down the process for everybody (since most people would wait to have read that proposal before voting any proposal). The question was presented and after 4 generations, and 61 proposals (of which 41 were original), the community reached a consensual agreement. It

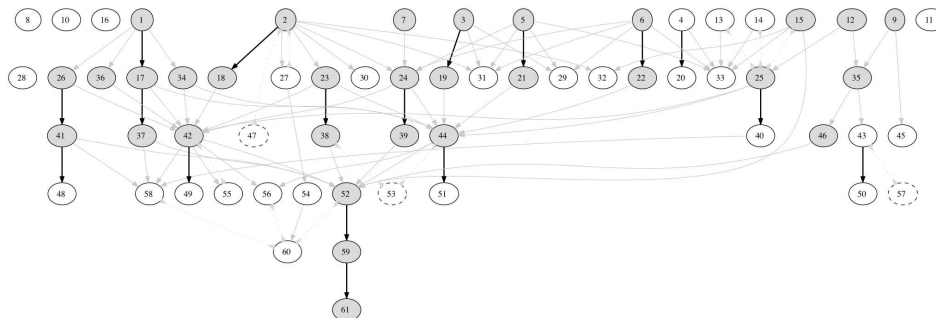


Fig. 3. Proposals submitted and voted in the “Wall of text” question. Sixty-one proposals were considered. After 4 generations a consensus was found. Each proposal is shown as a node. Proposals in the same generation are at the same height. Proposals that made it into the Pareto Front are copied to a new proposal in the next generation, and connected with a black arrow (making the actual number of original proposals to be 44). Grey arrows indicate a connection in the content (ideas that made it from one proposal to the next). Nodes with a dotted border are comments, (i.e. proposal not meant to be voted, but only presented by users to explain why a particular proposal is not acceptable. Making the actual number of real proposals 41). Dotted links represent a critique (the lower proposal criticize the higher proposal). By following the lines we can see that 29 proposals (with greyed background) out of 61 (or 19 out of 41) had their idea, or part of their idea, integrated in the final proposal.

should be noted that although the process seemed fast, each generation took two weeks time, of which one week in proposing and one week in voting. So the final result took about two months to complete. The final result had an abstract and the main text consisted of three points (an excerpt is presented here for brevity, the full data can be seen at [13]):

1. Proposals longer than 1,000 characters require an abstract. Abstracts are hard-limited to 500 characters. (These numbers could be tweaked.)
2. Give visual feedback about the quality of writing of the proposal body, perhaps in the form of a bar across to top of the input box which contains a gradation from green to red. [...] Metrics for the difficulty score could include: SMOG score; Length of proposal - weighted heavily; Number of other proposals by same author on same question [...]; Possibly “readability-votes” by other users; Possibly others;
3. Proposals are presented in order of their difficulty score. [...] And we confront it with some of the proposals presented in the first generation (the number refers to the actual proposal in Fig 2): (1) Writing an abstract of the proposal. (2) Shorter proposals appear in the first places of a ranking. There are two new buttons: I understand, I don’t understand. Negative understanding points sink the proposal in that ranking. (3) Prizes or reputation points for the succinct writer. (5) Rating systems. The idea being, that people will write more concisely to get a higher rating. (7) [...]include a Dismiss

button with one or two reason options. Select Reason for dismissal: Spam; Too Long; Irrelevant; Difficult to understand. Main point would be to get feedback to [the] proposer. It could be used with or without limiting character input. (12) It's fine as it is. No limit should there be on the length of proposals! (15) Start playing with length limits. Try a very restrictive one, such as 500-1000 characters and see how burdensome it is.[...] We see how no simple voting could have produced such result: One user suggesting to use 500 or 1000 characters limit; another to use an abstract; and a third to let proposals unbounded. The result included ideas from all three: proposals can be of any length (satisfying the third user), when they are of more than 1000 characters they need an abstract (thus satisfying the second user). Abstract which should not be longer than 500 characters (satisfying the first user). Note that the actual final proposal was much more complex as it had to satisfy more requirements.

4 Results

The main result that was observed was that the process worked better when the question asked was real and relevant. If the question was a test question, either unrealistic or such that the solution made no real difference in the life of the participants, it was harder to reach a consensus, the Pareto Front would be larger, people would get bored more easily and generally the result would just look random. When instead the question was real and relevant, the Pareto Front would never grow too much as participants would look for acceptable compromise between the winning proposals. From time to time the participants would try to recycle proposals that were eliminated, perhaps by integrating them inside winning proposals. Usually those actions were unsuccessful, but showed how the participants were trying to integrate their point of view inside the point of view of the community.

Researchers who work with Genetic Algorithms are used to running their experiments for a few hundred generations at a minimum. We could not run those experiments for such a long time. No willing volunteer would agree to participate. But in all the experiments that we made a consensus was found much sooner. Often after 3 or 4 generations. In particularly hard cases the consensus was reached after 10 or 11 generations. The answer to *which [what] is the meaning of life?* [14] could be found after only 6 generations (although the final solution was proposed on generation 4). The solution to the "wall of text" problem, presented above, needed 4 generations. And the result was unexpected by any of the participants, integrating elements from multiple people. Such good results hide a problem. When the result was hard to find, the participants would start to get bored, and often leave the question. Then the question would reach a smaller consensus. Rapid, but less meaningful. On the other hand we can assume that if the question was more important the participants would be less prone to leave. Leading to a situation where either they would find a consensus or they would

eventually enter in a loop. We should discuss this last possible result in the Discussion Section.

4.1 The Subjective Pareto Front as a subset of the Objective Pareto Front

An interesting question is: how would the system scale up? We are now going to present some mathematical results showing that under some basic assumptions, the size of the Pareto Front will be bounded, regardless of the number of people participating in the voting procedure. Let us assume that we have a question, with a number of diverse participants. We assume the participants will always vote according to their own subjective scale of values. We also assume that each person's value scale can be represented as a positive linear combination of some objective values that the participants share. In other words the participants share the same values but place different priorities upon them¹. We can now map every proposal according to those objective values, and given any participant, we can associate to that participant a line R , linear combination of those objective values. Each user will then order all the proposals according to his values as points on the line R .

Definition 1. *Given a point A and a line $R : y = ax$, we shall define A_R , as the point where R intersects the perpendicular line that is passing through A . We shall also refer to A_{R_x} (A_{R_y}) as the x (y) co-ordinate of the A_R point.*

Now we need to formally define the Pareto Front with respect to the subjective values of the participants. To do this we start by defining the concept of dominance.

Definition 2. *Given two distinct lines R, S (with $R : y = ax$ and $S : y = bx$) with $a, b > 0$. And given two distinct points $A(a_x, a_y)$ and $B(b_x, b_y)$, we will say that A dominates B respect to (R, S) if $A_R \geq B_R$, and $A_S \geq B_S$ and we shall write it $A >_{(R,S)} B$.*

In the field of Multi-Criterion Decision Making where Pareto Fronts are commonly used[11], the standard form of dominance considered is with respect to a basis of coordinates, (X, Y, \dots) . We too will say that a point A dominates a point B if $A >_{(X,Y)} B$

Definition 3. *The Pareto Front with respect to (R, S) will be the set of points that are non dominated respect to R and S .*

¹ This might seem a too strong presupposition, but a simple example will convince us that it is not so. One of the political discussions that most strongly polarises public opinion is the discussion on abortion. Everywhere around the world the people divide in a pro-choice group and a pro-life group. But the people in both groups have the same values: they both believe in freedom (thus choice) and life. They just diverge on the priority. No one wants an abortion, if it can be avoided, and no one wants to be told what she must do.

We need to prove that if two points $A(a_x, a_y)$ and $B(b_x, b_y)$ are in the Pareto Front with respect to the coordinates (X, Y) , then they can be in the Pareto Front with respect to (R, S) (with $R : y = ax$ and $S : y = bx$) with $a, b > 0$. We shall call the Pareto Front with respect to the (X, Y) coordinates the *Objective Pareto Front*, while the Pareto Front with respect to (R, S) the *Subjective Pareto Front*. In other words we are trying to prove that the Subjective Pareto Front is contained in the Objective Pareto Front.

First we will start with an existing result from the field in Multi-Criterion Decision Making [11] [15].

Proposition 1. *Let A dominate B ($A \succ_{(X,Y)} B$) then given any line $R : y = ax$, A_R will be farther from the origin than B_R ($A_{R_x} > B_{R_x}$ and $A_{R_y} > B_{R_y}$).*

Proof. Let us calculate A_{R_x} . To do this we calculate the line S perpendicular to R , and passing through A . Such line will be: $S = -\frac{1}{a}x + e$ and $A_{R_x} = a * ay + ax$. Thus $B_{R_x} = a * by + bx$. Now we need to prove that $A_{R_x} > B_{R_x}$. That is: $a * a_y + a_x > a * b_y + b_x$ but this is trivial since $a > 0$ and from the fact that $A > B$ we know that either (1) $ay > by$ and $ax = bx$ or (2) $ay = by$ and $ax > bx$ or (3) $ay > by$ and $ax > bx$. In all three cases $A_{R_x} > B_{R_x}$. And since $A_{R_y} = a^2 * ay + a * ax$ and $B_{R_y} = a^2 * by + a * bx$ then $A_{R_y} > B_{R_y}$ \square

Now we want to prove that if a point is in the Subjective Pareto Front it must be in the Objective Pareto Front.

Proposition 2. *Let A be a point in the Subjective Pareto Front then given two lines R, S ; A is in the Objective Pareto Front.*

Proof. The Thesis is equivalent to saying that if A is not in the Objective Pareto Front, it cannot be in the Subjective Pareto Front. Or, in other words, if it is being dominated with respect to (X, Y) it will also be dominated according to (R, S) . Without loss of generality let us consider two distinct points A, B . We want to prove that if $A \succ_{(X,Y)} B$, then $A \succ_{(R,S)} B$. But from Proposition 1, we know that $A_{R_x} > B_{R_x}$, and $A_{S_x} > B_{S_x}$. Thus the thesis. \square

Since the Subjective Pareto Front is a subset of the Objective Pareto Front, having more participants will not lead to an unbound growth of the Pareto Front, but just to a better approximation of the objective one. Which is, in itself, a positive result. In passing we note that this result also implies that when we need to estimate an Objective Pareto Front (a common occurrence in Multi-Criterion Decision Making [15]) we can simply ask an educated group of people to evaluate the options at hand, according to their personal preferences. And the result should naturally approximate the Objective Pareto Front.

5 Discussion

Two questions should be considered: how would such a system scale up and what should be done if a consensus is not found. Regarding the latter, we would

suggest that an evolutionary system should be used as a first means trying to reach consensus on an issue. If this fails, the partial results (i.e. the final Pareto Front) should be voted on in a more traditional way. But yet in a way that permits the participants to express multiple choices, like any Condorcet method. In this option still the initial evolutionary process would not go lost, as it has helped finding a set of optimal solution among which to chose.

But how much could such a system scale up and still be practical? Of course much work needs to be done in this regard. In the Results section we have shown how the Pareto Front will generally be bounded. Still the system might explode due to the number of proposals presented by the people. In our experience we noticed that it is always easier to get people to vote for proposal made by others then to write your own proposals. Thus we do expect the number of proposals to grow at a slower rate than the number of participants. But how much slower? In any case there will be a point where it will not be possible for everybody to evaluate the proposals from everybody else. But still the system could be implemented by distributing it upon a graph, with each node being a separate, parallel, genetic algorithm, and winning solution spreading among nodes. Also, even without solving the scaling problem the system could be used to let small to medium communities self-govern. Considering that probably less than 1 in 20 people will propose something, a community of a 1000 people could use the system in its current format. Also it could be used as a very democratic way to discuss something in a context where each participant's opinion is very important. For example, what would have happened if at the *2009 United Nations Climate Change Conference*, instead of having a more traditional voting system we had a Pareto Human Based Genetic Algorithm, with each nation being a participant? Surely no one would have protested that the discussion was being held in a non-democratic way.

Finally, we would like to point out how the system presented can be used in a different context all together, for example as basis for collaborative editing of a document, but that's another story.

6 Conclusions

We investigated an alternative way to reach an agreement on an open question. We looked at this as a possible aid in the decision making process, possibly for an e-government system. In our investigation we coded for a website where users can ask open questions and try to reach a consensus over them. In alternate phases each user can propose answers to the questions posed, or chose which answers he agrees on. By selecting the Pareto Front of the approved answers we extract a subset of the answers presented that fully represents the community. If no answer has reached unanimity, the Pareto Front is offered back to the community as an inspiration to write new proposals. Although the users are invited to try to integrate and bridge the proposals in the Pareto Front, no limit is posed over what each user can suggest. We present a trial case of the website on a particular question, and showed how the different proposals build up to

generate an final answer that everybody could approve. This was done through 4 generations, each divided into a proposing and a voting phase. More questions were tested, but few received the participation necessary to be considered as test cases of how the algorithm would behave. We also investigated the mathematics behind the extraction of a Pareto Front based on subjective evaluations of proposals and found under reasonable assumptions this to approximate the Pareto Front that would be extracted under an objective evaluation of such proposals. Thus suggesting that such subjective Pareto Front would not explode as the number of users increase. All together the algorithm proposed showed to be a possible alternative to more traditional forms of voting to reach an agreement on a solution. It's drawback seem to be how the increased participation risks lowering the number of people willing to participate...

References

1. Merriam Webster online dictionary, page, accessed 03.03.2010 <http://www.merriam-webster.com/dictionary/democracy>
2. Taylor and Pacelli. Mathematics and Politics: Strategy, Voting, Power and Proof. Springer Book (2008)
3. Data.gov <http://www.data.gov/> accessed 03.03.2010
4. Data.gov.uk <http://data.gov.uk/> accessed 03.03.2010
5. Vassilis Goulandris: Opengov.gr: The first 120 days of e-deliberation. available from <http://onlinepolitics.wordpress.com/2010/02/22/opengovgr-first-120-days-e-deliberation/> accessed 03.01.2010
6. John H. Holland,. Adaptation in natural and artificial systems. "An introductory analysis with applications to biology, control, and artificial intelligence", University of Michigan Press (Ann Arbor), 1975.
7. Alex Kosorukoff. Human based genetic algorithm. 2001 IEEE International Conference on Systems (2001)
8. Anne Defaweux, Tobias Grosche, Maria Karapatsiou, Alberto Moraglio, Alex Shenfield, Automated Concept Evolution. Technical Report Vrije Universiteit Brussel, Belgium (2003)
9. Speroni di Fenizio, Pietro and Anderson, Chris. Using Pareto Front for a Consensus Building, Human Based, Genetic Algorithm. Proceedings of the European Conference of Artificial Life, 2009.
10. Steven Brams, and Peter Fishburn. (2007). Approval Voting. Springer.
11. Matthias Ehrgott. Multicriteria optimization. Springer (2005)
12. Sevan G. Ficici Jordan B. Pollack Pareto Optimality in Coevolutionary Learning. Advances in Artificial Life: 6th European Conference, ECAL 2001, Prague, Czech Republic, September 10-14, 2001: Proceedings (2001)
13. Vilfredo goes to Athens. How should we handle the "wall of text" problem?, <http://vilfredo.org/viewhistoryofquestion.php?q=67&room=Vilfredo>, accessed 28.02.2010.
14. Vilfredo goes to Athens. Which is the meaning of Life?, <http://vilfredo.org/viewhistoryofquestion.php?q=14>, accessed 03.03.2010.
15. Evangelos Triantaphyllou, Multi-Criteria Decision Making Methods: A Comparative Study, Kluwer Academic Publishers, Boston, 2000.