Preliminary Results on Reputation Systems: Favoring More Balanced Profiles*

Jonathan Ben-Naim[1] and Dov Gabbay[2]

[1] Institut de recherche en informatique de Toulouse (IRIT-CNRS)
Université de Toulouse, 118 route de Narbonne
F-31062 Toulouse Cedex 9, France
bennaim@irit.fr
[2] Kings College, London
Strand, London WC2R 2LS
Bar-Ilan University
Ramat Gan, 52900 Israel
dov.gabbay@kcl.ac.uk

Abstract. In many large multi-agent systems (for example e-commerce applications), the users have to choose the agents to interact with. But this choice is difficult, in particular because the agents are too numerous. For the same reason, it is impossible to call a trusted arbiter that can judge every one of them. However, the agents provide a lot of information about their peers (for example feedbacks from past transactions). Consequently, the goal is to construct from this information an object helping the users to make decisions about the agents. Many reputation systems have been developed to achieve this goal. In the present sketch, we develop and begin to analyze, from an axiomatic point of view, new systems favoring those agents that have a more balanced profile. More precisely, the information available about the agents is modeled by a support graph, that is, the nodes represent the agents and an arrow means that the source agent supports the importance of the destination agent. The object constructed from this information is a ranking of the agents showing the relative importance of each agent. Our ranking methods are based on the following principle: for an agent, to reach a certain position in the ranking, it is necessary to pass a certain threshold both in the quantity of supporters and in their quality.

Keywords: reputation, ranking, axiomatization, multi-agent systems.

1 Introduction

In many large multi-agent systems, the users have decisions to make about the agents. In particular, a user has to choose those to interact with. But this choice is difficult, particularly because the agents are too numerous. For the same reason, it is impossible to call a trusted arbiter that can judge every one of them.

* The first author is supported by the French ANR-06-SETI-006 project "ForTrust".

However, the agents provide a lot of information about their peers. So, the goal is to construct, on the basis of this information, an object helping the users to make decisions. Many *reputation systems* have been developed to achieve this goal. They are called so because the object constructed typically gives a good account of the reputation of the agents.

Let us mention some important applications where a reputation system is needed. First of all, there are e-commerce applications, like Ebay. The agents are the buyers and the sellers and the information available about them takes the form of feedbacks from past transactions. Next, there are the webpages and their hypertext links. Indeed, a page can be seen as an agent and a link as the fact that the source page supports the importance of the destination page. Next, there are papers and their citations. Again, a paper can be seen as an agent and a citation as a positive opinion. As a last example, let us mention the bindings entity-key and the public key certificates. A binding $\langle E, K \rangle$ can be seen as an agent and a certificate written by E, signed with K, and supporting the validity of a binding $\langle E', K' \rangle$ can be seen as a positive opinion of $\langle E, K \rangle$ on $\langle E', K' \rangle$.

The present sketch contains *work in progress*. It is designed for a workshop intended primarily as a forum for interaction and discussion. Consequently, it is largely incomplete in several aspects, in particular bibliography, formal definitions, and results. Concerning the contributions, we develop two new reputation systems based on the following principle: those agents that have a more balanced profile are preferred. To properly introduce this principle, we need to specify what we mean by the terms reputation system and what particular kind of such systems we investigate.

By reputation system, we mean a system composed of two parts. The first part is an environment containing many agents that have the possibility of providing information about their peers. The second part is a method that, in any possible situation, constructs from this information an object helping the users to make decisions about the agents.

In the present work, the information available about the agents is modeled by a *support graph*. A node represents an agent and an arrow means that the source agent supports the importance of the destination agent. The object constructed from such a graph is a *ranking* of the agents showing the relative importance of each agent. By ranking, we mean a reflexive, transitive, and total binary relation $\preceq$ where $a \preceq b$ means that a is at least as important as b. The rank of an agent can be seen as a good representation of its reputation, since it is determined by what is said about that agent. Although this framework is very simple, it is sufficient to provide solutions for all the aforementioned applications.

When comparing two agents, we consider two criteria: the number of supporters and their importance. Our *ranking methods* are based on the following principle: the agents that receive a more balanced kind of support are preferred. In other words, we want to prevent the agents that have a big flaw in some criterion from getting a high position in the ranking. More precisely, for an agent, to reach a certain position in the ranking, it is necessary to pass a certain threshold both in the quantity of supporters and in their quality. We call this the *more-balanced-support principle.*

We develop two completely different methods that follow this principle. One is based on refinements, the other on improvements. But actually, we strongly believe that both methods always give the same ranking, that is, we believe that they are two different ways to define the same ranking function. If this is indeed the case, this would indicate that the ranking function in question is relatively central, because it would be motivated from two completely different points of view.

Next, it is crucial to analyze ranking methods from an axiomatic point of view. To our knowledge, the first axiomatic study of such methods was started in a paper of M. Tennenholtz [Ten04] and further developed in several papers written in collaboration with A. Altman, in particular [AT05], [AT06], [AT07a], [AT07b], and [AT08]. In the present sketch, we begin to analyze our methods. More precisely, we present two axioms that are among the most important for the kind of ranking methods considered: transitivity and strict transitivity. We are very confident that both of our methods satisfy these axioms, which would provide theoretical justifications for them.

2 Axioms

Recall that a support graph is a graph where a node represents an agent and an arrow means that the source agent supports the importance of the destination agent. In addition, recall that a ranking method is a function that transforms any support graph into a ranking of the agents showing the relative importance of each agent. We present two important properties for the kind of ranking methods considered. Actually, they could be seen as two parts of one property, but to improve readability, we present them separately:

- *Transitivity.* Consider two agents a and b, and suppose that the group A of all supporters of a is at least as important as the group B of all supporters of b, then a should be ranked at least as high as b. We consider that A is at least as important as B iff $B = \emptyset$ or there exists an injective function f from B to A such that for each element x of B, $f(x)$ is ranked at least as high as x.

- *Strict transitivity.* If the group A supporting a is strictly more important than the group B supporting b, then a should be ranked strictly higher than b. We consider that A is strictly more important than B iff (0) or (1) holds, where (0) is the following: $A \neq \emptyset$ and $B = \emptyset$, and (1) is that there exists an injective function f from B to A such that, first, for each element x of B, $f(x)$ is ranked at least as high as x, and second, $|B| < |A|$ or there exists x in B such that $f(x)$ is ranked strictly higher than x.

These axioms are discussed in detail in [AT08]. There exist several other desirable properties, but in the framework of this preliminary work, we investigate only those two.

3 A refinement-based ranking method

We develop a first ranking method based on the more-balanced-support principle. Note that a ranking on a set can be seen as an ordered partition of this set. By category, we mean an element of this partition. The idea is to begin with the trivial ranking R_0 that puts every agent in the same category, then we refine R_0 an infinite number of times, which gives R_1, R_2, etc. A final ranking is derived as follows: an agent is at least as important as another one iff this is the case in all the preliminary rankings.

In each refinement step, a new, finer, *view on the quality* of the agents becomes accessible. Our goal is to refine in such a way that, in the end, the following holds: for an agent, to reach a certain position in the final ranking, it is necessary to pass a certain threshold in every criterion. By criterion, we mean the quantity of supporters or any view on their quality.

Let us take a closer look at the first refinement step. When comparing two agents, a certain number of factors are considered: the quantity of supporters and every view on their quality. The influence of each factor is represented by a certain number of points. Concerning quantity, an agent naturally scores 1 point for each of its supporters. On the other hand, concerning the views on quality, no point can yet be scored. Indeed, in the first refinement step, these views are not yet accessible. They are precisely under construction. Therefore, they cannot have any influence on the first refinement.

So, each agent gets a certain score determined only by the quantity of supporters. If we were to refine R_0 on the basis of these scores, we would put in the top category those agents that have a maximal number of supporters, regardless of their importance. This does not correspond to our guideline. To counterbalance the absence of influence of quality, we limit the influence of quantity. More precisely, we cap the number of supporters considered in this step at 1, the minimal value such that the quantity of supporters still has some influence.

To summarize, if an agent has at least 1 supporter, it scores 1 point, and that is all. In other words, in the first refinement step, we consider two levels of quality: the level corresponding to a score of 0 and the level corresponding to a score of 1. The agents are evaluated on the basis of these two levels and ranked accordingly. The resulting ranking is called R_1.

Next, we enter the second refinement step. Thanks to the scores obtained in the first step, we have now access to a first, rough, view on the quality of the agents. More precisely, we can distinguish between a basic and a non-basic agent. By basic agent, we mean an agent having no supporter. But, we cannot yet distinguish between two non-basic agents of different importance. This first view on quality will influence the second refinement.

More precisely, if a supporter becomes more important in the first view on quality, then the agent supported naturally scores 1 point. By becoming more important in the first view, we mean getting a score of 1 in the first refinement step. In parallel, we lift a bit the limitation on the quantity of supporters, so this factor keeps having some minimal influence. More precisely, we increase the cap on the number of supporters considered by 1. In other words, if an agent has at least 2 supporters, it scores an additional point.

So, each agent gets a score determined by the quantity of supporters and the first view on their quality. If we were to refine R_1 on the basis these scores, we would put in the top category the agents having a maximal number of supporters that are important in the first view on quality, regardless of their importance in the next, finer, views that will become accessible later. This does not correspond to our guideline. Consequently, we limit the influence of the first view.

More precisely, the advancement of only one supporter is taken into account at a time. Therefore, we have to choose the order in which the supporters will be treated. We adopt the following strategy: the more a supporter is important, the earlier its advancement will be taken into account. More precisely, we introduce some linear ordering $l(R_1)$ on the agents that agrees with R_1. The choice of this linear ordering has no effect on the result. In the second refinement step, only the advancement of the best supporter of a in $l(R_1)$ will possibly allow a to score a point. In the next step, we will take into account the advancement of the second best supporter, and so on.

To summarize, if an agent a has at least 2 supporters, it scores 1 point. If the best supporter of a in $l(R_1)$ exists and gets a score of 1 in the first step, then a scores 1 point. In other words, for each category of R_1, three levels of quality are considered: the levels corresponding respectively to a score of 0, 1, and 2. The agents are evaluated on the basis of these three levels and ranked accordingly inside each category of R_1. The resulting ranking is called R_2.

Next, we enter the third refinement step. Thanks to the scores obtained in the second step, we have now access to a second, finer, view on the quality of the agents. This second view will influence the third refinement. More precisely, if a supporter becomes more important in this second view, then the agent supported scores a certain number of points. In addition, the more a supporter becomes important, the more the agent supported scores many points. More precisely, if a supporter makes x advances in the second view on quality, then the agent supported scores x points. By making x advances in the second view, we mean getting a score of x in the second refinement step. As previously, we limit the influence of this second view on the current refinement. In parallel, we lift a bit the limitations on the two other factors: the quantity of supporters and the first view on their quality.

To summarize, if an agent a has at least 3 supporters, it scores 1 point. If the second best supporter of a in $l(R_1)$ exists and gets a score of 1 in the first step, then a scores 1 point. If the best supporter of a in $l(R_2)$ exists and gets a score of x in the second step, then a scores x points. So, for each category of R_2 a certain number of levels of quality are considered, one for each possible score. The agents are evaluated on the basis of these levels, and ranked accordingly inside each category of R_2. The resulting ranking is called R_3. And so on.

More generally, in each refinement step, we have access to a new, finer, view on the quality of the agents. When a new view becomes accessible, we limit its influence on the current refinement. We do it in order to give time to the next, finer, views to become accessible and to influence significantly the final ranking. In parallel, we lift a bit the limitations on the quantity of supporters and on the already accessible views on their quality, so these factors keep having some minimal influence. As a result, for each agent and for each position in the ranking, the following holds: for that agent to reach that position, it is necessary to pass a certain threshold in the quantity of supporters and in every view on their quality, which corresponds to our guideline.

Here is for example a support graph containing 9 agents, and below are the five first rankings obtained from it by the refinement-based method:

Example 1.

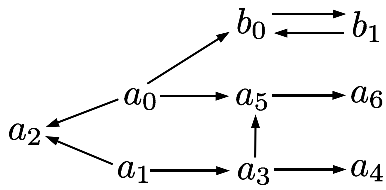

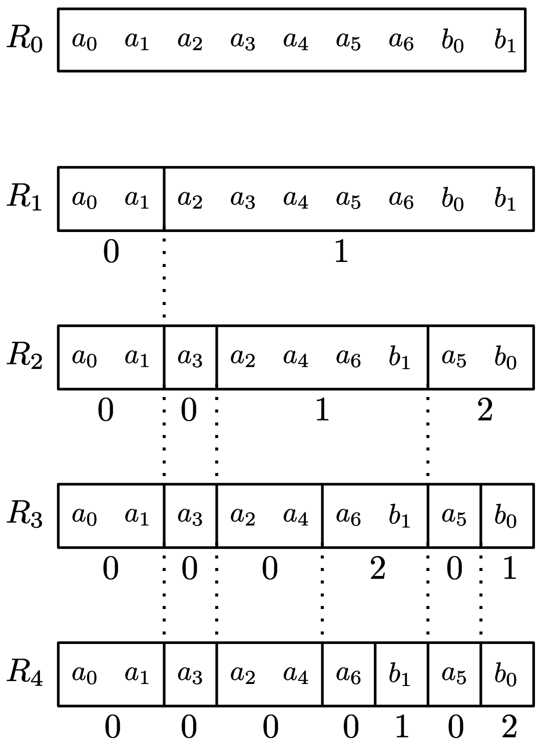

The initial ranking, denoted by R_0, puts every agent in the same category. We enter the first refinement step. The agents a_0 and a_1 have no supporter, so they get a score of 0. Actually, if an agent gets a score of 0 in some step, it will get a score of 0 in all the subsequent steps. The other agents do have supporters, so they all get a score of 1. These scores represent a first, rough, view on the quality of the agents. The resulting ranking is denoted by R_1.

We enter the second refinement step. The agent a_3 gets a score of 0. The agents a_2, a_4, a_6, and b_1 get a score of 1. More precisely, a_2 scores 1 point because it has 2 supporters. Concerning a_4, a_6, and b_1, they score 1 point because they have a supporter that makes an advance in the first view on the quality of the agents. The agents a_5 and b_0 both get a score of 2. Indeed, they score 1 point because they have two supporters and an additional point because they have a supporter that makes an advance in the first view on quality. These scores represent a second, finer, view on the quality of the agents. The resulting ranking is denoted by R_2.

We enter the third refinement step. The agents a_2, a_4, and a_5 get a score of 0. The agent b_0 gets a score of 1, because it has a supporter that makes an advance in the second view on quality. The agents a_6 and b_1 both get a score of 2, because they have a supporter that makes 2 advances in the second view on quality. These scores represent a third, finer, view on quality. We denote by R_3 the resulting ranking.

We enter the fourth step. The agent a_6 gets a score of 0. The agent b_1 gets a score of 1, because it has a supporter that makes an advance in the third view on quality. The agent b_0 gets a score of 2, because it has a supporter that makes 2 advances in the third view. The resulting ranking is denoted by R_4.

All the subsequent refinements leave R_4 unchanged, so R_4 is the final ranking obtained in this example by our refinement-based method.

4 An improvement-based ranking method

We develop a second, completely different, ranking method based on the more-balanced-support principle. Suppose a support graph is given. The idea is to imagine that the support links generate, step by step, two kinds of improvements: improvements in the importance of an agent and improvements in the importance of an improvement generated earlier. More precisely, we construct step by step from the support graph, an *improvement graph* where a node represents either an agent or an improvement, and an arrow means that the source element is an improvement in the importance of the destination element.

Initially, the improvement graph contains only a certain number of nodes, no link. These nodes represent exactly the same agents as those of the support graph. Next, for each agent a and each agent b, if a supports b, we introduce a new node p representing an improvement, and an arrow from p to b meaning that p is an improvement in the importance of b. In other words, the support a provides to b generates an improvement in the importance of b.

In the second step, for each agent a and each improvement p introduced in the first step, if a supports the agent (providing the support) that generated p, then we introduce a new improvement p' and an arrow from p' to p. In other words, the support a provides to an agent b generates an improvement in the importance of every improvement generated by b.

More generally, in the i^{th} step, for each agent a and each improvement p introduced in the $(i-1)^{\text{th}}$ step, if a supports the agent that generated p, then we introduce a new improvement p' and an arrow from p' to p.

In some sense, we unfold the support graph, that is, we replace a graph possibly containing cycles by a graph with no cycle, but possibly containing infinite chains.

Here is for example the improvement graph obtained from the support graph presented in Example 1:

Example 2.

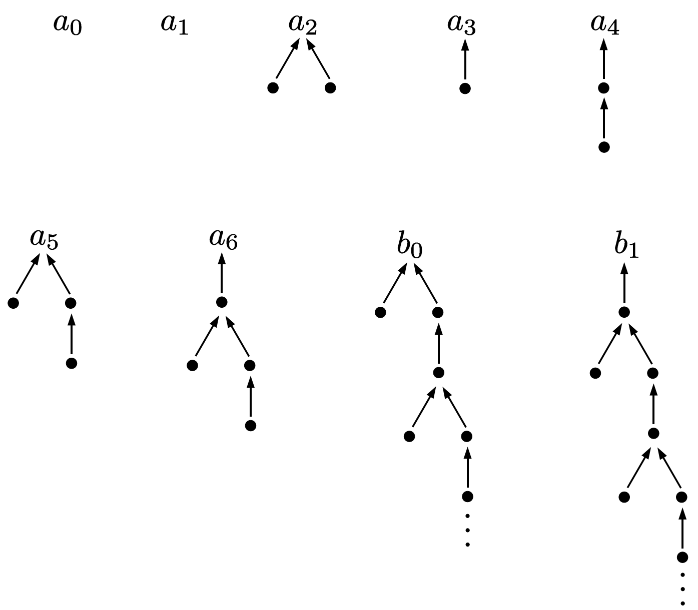

For each agent a, we extract from the improvement graph, the subgraph $S(a)$ of all improvements related (directly or indirectly) to a. This subgraph is called the *improvement structure* of a. The crucial point is that the form of $S(a)$ gives a good idea what kind of support a receives. More precisely, the more $S(a)$ is horizontally large, the more the supporters of a are numerous. The more it is vertically large, the more they are important. So, the idea is to rank the agents on the basis of the form of their improvement structures. More precisely, suppose we have at our disposal a ranking of the improvement structures showing the relative quality of the form of each structure. Then, we define that an agent a is at least as important as an agent b iff $S(a)$ is ranked at least as high as $S(b)$.

So, it remains to construct a ranking of the improvement structures. In accordance with the more-balanced-support principle, we favor those structures that have a more balanced form. In other words, for a structure, to reach a certain position in the ranking, it is necessary to pass a certain threshold both in the horizontal and in the vertical development. To meet this requirement, we adopt the following strategy: the rank of an improvement structure is determined by the rate at which it can be constructed under certain limitations. Let us explain this strategy.

A *construction* of an improvement structure $S(a)$ is a step-by-step development that begins with a and ends with $S(a)$, or is more and more approaching $S(a)$. In each step, two restrictions are imposed on the development of the structure under construction. First, only improvements that are directly connected to the existing elements can be added. In other words, it is impossible to add (in one step) a chain of improvements. This limits the vertical development. The second restriction is the following: for each existing element, at most one new improvement can be introduced and attached to it. This limits the horizontal development.

In more precise terms, a construction of $S(a)$ is an infinite sequence of graphs $\langle S_0, S_1, \ldots \rangle$ such that the three following points hold: first, S_0 consists only of a; second, the union of all S_i's is identical to $S(a)$; and third, S_{i+1} can be obtained from S_i by introducing, for each node p of S_i, at most one new node p' and a link from p' to p.

For example, consider the improvement graph given in Example 2. Here is a construction of $S(a_6)$:

Example 3.

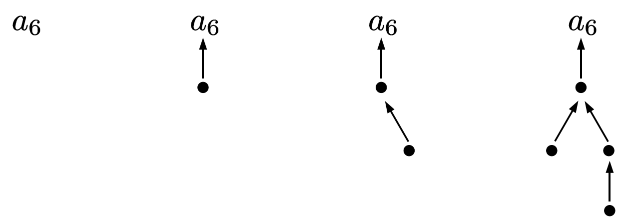

Intuitively, because of the development restrictions, the following holds: for an improvement structure, to be constructible at a certain rate, it is necessary to be sufficiently developed both horizontally and vertically. In other words, it is necessary to be sufficiently balanced. So, we rank the improvement structures on the basis of the rate of their constructions. More precisely, suppose we have at our disposal a ranking of the constructions showing the relative intensity of the rate of each construction. Then, we define that the form of a structure S is at least as good as the form of a structure S' iff some construction of S is ranked at least as high as every construction of S'.

So, the last step is to construct a ranking of the constructions. A construction $\langle S_0, S_1, \ldots\rangle$ defines a certain *sequence of increases*, that is, the sequence $\langle n_0, n_1, \ldots\rangle$ such that every n_i is equal to the number of new elements introduced in S_i. For example, the sequence of increases defined by the construction of $S(a_6)$ given in Example 3 is $\langle 1, 1, 1, 2, 0, 0, \ldots\rangle$.

We consider that the rate of a construction $\langle S_0, S_1, \ldots\rangle$ is at least as high as the rate of a construction $\langle S'_0, S'_1, \ldots\rangle$ iff both constructions define the same sequence of increases, or $\langle S_0, S_1, \ldots\rangle$ defines a lexicographically greater sequence, that is, in some step i, more improvements are introduced in S_i than in S'_i, and in every previous step j, the same number of improvements are introduced in S_j and S'_j.

To summarize, we define a ranking showing the relative intensity of the rate of each construction, from which we define a ranking showing the relative quality of the form of each improvement structure, from which we finally define a ranking showing the relative importance of each agent. Put differently, an agent a is at least as important as an agent b iff there exists a construction of $S(a)$ such that the rate of this construction is at least as high as the rate of every construction of $S(b)$.

For example, consider the support graph introduced in Example 1. For every agent a, we have chosen one of the highest-ranked construction of $S(a)$, and we have displayed below the sequence of increases defined by this construction. Finally, the ranking of the agents obtained from the support graph of Example 1 by the improvement-based method is also displayed below. It is identical to the one obtained by the refinement-based method.

Example 4.

a_0	1	0	0	0	0	0
a_1	1	0	0	0	0	0
a_2	1	1	1	0	0	0
a_3	1	1	0	0	0	0
a_4	1	1	1	0	0	0
a_5	1	1	2	0	0	0
a_6	1	1	1	2	0	0
b_0	1	1	2	1	2	1
b_1	1	1	1	2	1	2

a_0 a_1	a_3	a_2 a_4	a_6	b_1	a_5	b_0

5 Conclusion

To conclude, we have developed two new ranking methods based on the following principle: for an agent, to reach a certain position in the ranking, it is necessary to pass a certain threshold both in the quantity of supporters and in their quality. We are very confident that both of our methods satisfy transitivity and strict transitivity, which would provide theoretical justifications for them. In addition, we strongly believe that these two methods are actually two different ways to define the same ranking function. If this is the case, this would mean that the ranking function in question is relatively central, because it would be motivated from two completely different points of view. We are working on these conjectures and on a complete axiomatization of our ranking methods.

References

[AT05] Alon Altman and Moshe Tennenholtz. On the Axiomatic Foundations of Ranking Systems. In *Proceedings of 19th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 917–922, 2005.

[AT06] Alon Altman and Moshe Tennenholtz. Quantifying Incentive Compatibility of Ranking Systems. In *Proceedings of the 21th National Conference on Artificial Intelligence (AAAI)*, volume 1, pages 586–591, 2006.

[AT07a] Alon Altman and Moshe Tennenholtz. An Axiomatic Approach to Personalized Ranking Systems. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1187–1192, 2007.

[AT07b] Alon Altman and Moshe Tennenholtz. Incentive Compatible Ranking Systems. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1–8, 2007.

[AT08] Alon Altman and Moshe Tennenholtz. Axiomatic Foundations for Ranking Systems. *Journal of Artifical Intelligence Research (JAIR)*, 31:473–495, 2008.

[Ten04] Moshe Tennenholtz. Reputation Systems: An Axiomatic Approach. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 544–551, 2004.