Consistent Continuous Trust-Based Recommendation Systems

Ola Rozenfeld and Moshe Tennenholtz

Technion – Israel Institute of Technology, Haifa 32000, Israel,
Microsoft R&D Labs, Israel
ola.rozenfeld@gmail.com,
moshet@ie.technion.ac.il

Abstract. The goal of a trust-based recommendation system is to generate personalized recommendations from known opinions and trust relationships. Prior work introduced the axiomatic approach to trust-based recommendation systems, but has been extremely limited by considering binary systems, while allowing these systems to be inconsistent. In this work we introduce an axiomatic approach to deal with consistent continuous trust-based recommendation systems. We introduce the model, discuss some basic axioms, and provide a characterization of a class of systems satisfying a set of basic axioms. In addition, as it turns out, relaxing some of the axioms leads to additional interesting systems, which we examine.

1 Introduction

Many online systems offer their users access to a vast variety of products or services, which drives the need for high quality personalized recommendations. Often, these systems utilize social network structures of their users, as well as the users' opinions of one another and the products, in order to improve the quality of recommendations. Google's page ranking system [18] uses links to represent "voting" for web pages, Amazon and eBay's reputation systems (e.g. [19]) aggregate feedbacks that users leave for transactions , and the Epinions trust/reputation system (e.g. [17]) aggregates explicit trust/distrust links between its users. In recent years, recommendation and reputation systems became the focus of intensive research (e.g. [15, 20, 8, 23, 11]).

In this work we focus on the setting where there is one item of interest, and various users have rated this item. A user wishes to "predict" her own rating of the item by consulting her friends, who, in turn, might consult their friends and so on. There exist many automated recommendation systems that fit this general framework; however, this raises the question of comparing the relative merits of these systems. There exist two main approaches to studying and comparing recommendation/reputation systems: the experimental approach and the axiomatic approach. The experimental approach evaluates the performance of a given system on a particular set of users and ratings (along various performance metrics); the obvious advantage of this approach is that it provides a practical estimate of the real accuracy of the system. However, the results of a particular system on different data might be different; we need a general understanding of the properties of different systems *before* we decide which one we want to implement for a particular setting. The axiomatic approach, which has a long history in social choice theory, aims to achieve these goals.

In [6] the authors use the axiomatic approach in the restricted case of recommendation systems in which the setting is an annotated directed graph, where some of the nodes are labeled by votes of $+$ and $-$. In that model a node represents an agent, an edge directed from a to b represents the fact that agent a trusts agent b, and a subset of the nodes are labeled by $+$ or $-$, indicating that these nodes have already formed positive or negative opinions about the item under question. Based on this input, a recommendation system must output a recommendation for each unlabeled node. In this paper we extend this study to the general case, where trust values and votes are continuous, allowing users to express a range of recommendations, instead of just a binary "like/dislike". This extension is not only a technical one, since it allows considering *consistent* trust-based recommendation systems, as explained below.

One of the most desired properties of a recommendation system is that it should be consistent with its own recommendations. Namely, if a system comes with a particular recommendation to agent p, based on other agents' observations and the trust network, then if this recommendation turns out to be correct, this should not change any recommendations for the other agents. It is easy to see that the simplified binary setting does not allow for such consistency. Hence, we view the axiomatic study of consistent continuous trust-based recommendation systems as essential for the understanding/analysis of desirable realistic systems.

1.1 Overview of results

In section 2 we introduce our model of continuous trust-based recommendation systems, where weights are non-negative real numbers, and votes are normalized to the interval $[-1, 1]$, In this framework we define our model axioms, which are minimal requirements we may wish to have, as well as more elaborated axioms dealing with the notion of consistency; these are further extended to several monotonicity axioms. In section 3 we introduce an axiom termed Independence of Irrelevant Stuff [IIS], and show that this axiom simplifies matters by allowing to consider only the way votes are locally aggregated. We show a natural recommendation system that satisfies our axioms, the Random Walk system, which is an extension of the system studied in previous work. Much emphasis in this paper is given to transformations of recommendation systems which preserve their desired properties; in particular, in section 4 we introduce two transformations that allow to tweak an existing recommendation system in order to give more/less weight to radical votes, and to define the meaning of different trust values. This brings us in section 5 to a theorem which fully characterizes the set of recommendation systems which satisfy a set of five natural axioms; namely, this characterization shows that all systems satisfying these axioms must be a modification of the Random Walk system under the two transformations above, and that each such modified Random Walk system satisfies the axioms. We also show that all five axioms are essential. By relaxing the axioms, we get some interesting new recommendation systems; systems which incorporate discount factors are discussed in section 6, and systems which incorporate the aggregated network vote are discussed in section 7.

1.2 Additional Related work

There are several ways to study recommendation systems. Standard evaluation tools include simulations and field experiments (e.g. [7, 19, 13]). In addition, one may also consider computational properties of suggested systems. As far as axiomatic studies are concerned, our work builds on previous work on axiomatizations of ranking systems. The literature on the axiomatic approach to ranking systems deals with both global ranking systems [1, 2, 22, 9, 23, 4, 5] and personalized ranking systems [7, 10, 3, 16]. Personalized ranking systems are very close to trust-based recommendation systems. In such systems, agents rank some of the other agents. Then an aggregated ranking of agents, personalized to the perspective of a particular agent, is generated based on that information. However, previous studies on the axiomatic approach have not been concerned with situations where the participants share reviews or opinions on items of interest which are external to the system. Many existing recommendation system are based on collaborative filtering (CF), which is a completely different approach than the trust-based systems considered in this paper. Combining trust-based and CF approaches is a direction of recent research [21].

2 Model and Axioms

The trust-based recommendation setting can be modeled formally as follows:

A *voting network* is a directed annotated graph $G = (N, V, E, v, w)$ where N is a set of agents, $V \subseteq N$ is a set of voters (the agents who have an informed opinion on the item of interest), $v : V \rightarrow [-1, 1]$ gives the vote of each voter, which represents his opinion on the item (normalized to $[-1, 1]$), $w : E \rightarrow \Re^{\geq 0}$, where $w(x, y)$ represents the trust of agent x in agent y. We will denote by $\overline{V}$ the set of non-voters $N \setminus V$.

Example of a voting network appears in Fig.1.

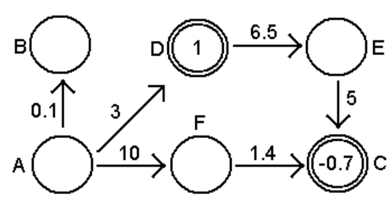

Fig. 1. Here, voters are designated by a double circle ($V = \{C, D\}$). The votes v appear inside the circle; the trust weights appear near the appropriate edge.

A *recommendation system* R takes as input a voting network G and has as output the recommendations $r(u) \in [-1, 1]$ for every $u \in \overline{V}$ ($R(G) : \overline{V} \rightarrow [-1, 1]$). We denote $R(G)$ by R_G; for notation simplicity, we extend the definition of R_G to N by setting $R_G(u) = v(u)$ for all $v \in V$ (so that $R_G(u)$ now means recommendation or vote of u).

This model is the natural extension of the model in [6] for continuous trust and vote values. First, we restrict ourselves to recommendation systems that satisfy the following basic conditions:

1. **Anonymity:** Isomorphic graphs correspond to isomorphic recommendations. Formally, for a permutation π of N such that $\pi(G) = G'$ (where $\pi(G)$ stands for applying π on v and w appropriately), it holds that $R(G') = \pi(R_G)$
2. **Neutrality:** The system is a priori indifferent towards positive or negative opinions – switching the signs on votes will cause switched signs of recommendations. Formally,
$R(-G) = -R_G$, where $-G = (N, V, E, -v, w)$.
3. **No-edge equals zero trust:** Adding edges with weight 0 does not affect recommendations.
4. **No-node equals orphan non-voter:** Adding non-voters with no incoming or outgoing edges does not affect any existing recommendations.
5. **Continuity:** The recommendation of a node v is a *continuous* function of the votes and the trust in every point, except possibly for points where all v's outgoing trust equals 0.

We call this set of conditions **Model Axioms (MA)**. In essence, we want these requirements as part of our model. Requirements 1 and 2 are natural, and appeared in the same form in the binary model of [6] as well (here, we renamed Symmetry to Anonymity in order to be consistent with social choice literature). In 3 we wanted to express that trust value of 0 is equivalent to no trust at all, and in 4 we wanted to express our implicit assumption that N represents only those agents who are in some sense informed (either know other agents or tried the product); implicitly, there is an infinite amount of agents in the system, but the recommendation system should only take the informed ones into account. (An alternative way to model the requirements 3 and 4 was to fix the set of agents to the set of natural numbers, $\mathbb{N}$, and require the input trust weights to induce a finite graph). Finally, in 5, we want the system to be "stable" in a sense: we don't want small changes in the input to cause big changes in the output. However, the case where a node v is a sink is a reasonable exception to this rule, since trusting no one at all and having trust in someone (no matter how small) are fundamentally different situations. For example, in many natural systems it holds that, if a node trusts a single voter with a vote of v, then its recommendation will also be v, no matter what is the weight of the trust link. However, if a node does not trust anyone at all, its recommendation is 0. We don't want these cases to violate our continuity requirement.

In the following we implicitly assume MA; in particular, when we speak of paths in a voting network G we will refer only to edges of positive weight; parents (or predecessors) of a node u ($Pred(u)$) are nodes z for which $w(z, u) > 0$; similarly, children (or successors) of a node u ($Succ(u)$) are nodes z for which $w(u, z) > 0$. However, we admit that the requirement MA[4] is debatable; we will show some natural recommendation systems that violate it.

In this work, we want to concentrate on the consistency requirement:

Node Consistency (NC): Let u be a non-voter. If we turn u into a voter with the vote

$R_G(u)$, no recommendations will change. Formally, for any voting network G and agent $u \in \overline{V}$, let $G' = (N, V \cup \{u\}, E, v \cup \{(u, R_G(u))\}, w)$. Then, for all $z \in \overline{V}$, $R_{G'}(z) = R_G(z)$.

Intuitively, consistent recommendation systems rely on their own predictions – they don't change their prediction for an agent if their predictions for other agents turned out to be correct. Their output is stable, in a sense: if the output is added to the input, nothing changes.

Consistency seems a very natural requirement; however, many systems used in real life are not consistent. Consider, for example, the Majority system: the recommendation of a node is 1 if the majority of his neighbors cast a positive vote, -1 if the majority of his neighbors cast a negative vote, and 0 otherwise. This system does not rely on its own predictions, since only voter neighbors are counted (it violates node consistency). A similar problem occurs in all binary systems – the Random Walk and Personalized Page Rank systems, as defined in [6], are not consistent because of similar considerations: intuitively, a binary system which has to commit to recommendations of 0, 1 or -1, cannot be very sensitive to nuances in the input (in the sense that many different inputs map to the same output), and therefore has little chances of being consistent. [1] However, this particular problem is somewhat artificial: we can redefine the systems to fit appropriately into the continuous setting. For example, the Personalized Page Rank system (PPR) [12] can be adapted as follows: the recommendation of a node v is the expected vote value of a voter that can be reached from a random walk starting from v, with a restarting probability α (in the binary setting, it was defined as the sign of the above value). It is easy to see that even this continuous system is not consistent. Consider the graph in Fig.2: the PPR recommendation of node B is 1; but if B becomes a voter with a vote of 1, the PPR recommendation of A will become 0.5 instead of 0.2.

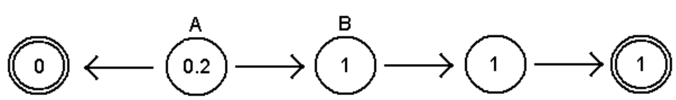

Fig. 2. Here, the votes and recommendations of PPR appear inside the circles; $\alpha = 0.5$; the trust weight is 1 for all edges.

Thus far, we spoke of consistency in terms of node values (votes). However, our system has two kinds of input: nodes and edges (votes and trust). Below is one intuitive requirement concerning trust values:

Edge Consistency (EC): If we increase trust in an agent with a value (recommendation or vote) equal to ours, our value will not change. Formally, for any voting network G, agent $u \in \overline{V}$, edge $(u, z) \in E$, and $w' > w(u, z)$ let $G' = (N, V, E, v, w \setminus \{((u, z), w(u, z))\} \cup \{((u, z), w')\})$. Then, $R_G(u) = R_G(z) \quad \Rightarrow \quad R_{G'}(u) = R_G(u)$.

Consistency in itself is too weak a requirement: it only tells us how the system should behave if its output is fed back into the input, without any changes. For example, a recommendation system that always gives a recommendation of 0 to all agents

[1] It would be interesting to translate this intuition into a formal impossibility proof.

is consistent, as well as the AVE system that gives the recommendation of the average vote to all agents. Now we would like to express the intuition about how the output of a recommendation system should change if the input changes. The input of the system comes in two parts: votes and trust. Therefore, as before, we formulate two requirements: what happens when a vote changes, and what happens when trust changes.

Node Monotonicity (NM): For any voting network G and agent $u \in N$, let $G' = (N, V \cup \{u\}, E, v \setminus \{(u, R_G(u))\} \cup \{(u, r')\}, w)$. Then:

1. If $r' = R_G(u)$, then for every $z \in N$, $R_{G'}(z) = R_G(z)$
2. For every parent z of u,
   $sgn(R_{G'}(z) - R_G(z)) = sgn(r' - R_G(u))$

Edge Monotonicity (EM): For any voting network G, agent $u \in \overline{V}$, edge $(u, z) \in E$, and $w' > w(u, z)$ let $G' = (N, V, E, v, w \setminus \{((u, z), w(u, z))\} \cup \{((u, z), w')\})$. Then, $sgn(R_{G'}(u) - R_G(u)) = sgn(R_G(z) - R_G(u))$.

Here, $sgn(x)$ is defined for $x \in \Re$ as:

$$sgn(x) = \begin{cases} \frac{x}{|x|} & x \neq 0 \\ 0 & otherwise \end{cases}$$

Note that NM implies NC and EM implies EC (follows from the definitions). Intuitively, by node monotonicity we demand the following: if we change a parent's vote, the child's vote must change appropriately (this is what makes the monotonicity requirement strict). The trivial system of recommending 0 to everyone satisfies NC, but not NM. Edge monotonicity states that increasing trust in a given agent, all other things being equal, should bring our recommendation strictly closer to the vote of that agent [2] (note that due to NC, it doesn't matter if we restrict z to belong to V in the definition of EM). Note also that AVE system satisfies NM and EC, but not EM (since it ignores the trust network).

The first question of interest for us is: which recommendation systems satisfy NM and EM? The Random Walk system (RW) is defined as follows: intuitively, to get a recommendation for a non-voter u, we perform a random walk in G starting from u, where in each step we move from node z to a random neighbor z', choosing z' with a probability proportionate to the trust of z in z', relatively to the overall trust of z. If the walk reaches a voter, the value of the walk is the vote value. Otherwise, the value of the walk is 0. The recommendation is the expected value of a walk over all random walks. Formally, for a voting network G, let $S \subseteq \overline{V}$ be the set of non-voters that cannot reach any voter. For each $u \in N$, create a variable $r_u \in [-1, 1]$. Solve the following from r_u:

$$r_u = \begin{cases} 0 & u \in S \\ v(u) & u \in V \\ \frac{\sum_{z \in Succ(u)} w(u,z) r_z}{\sum_{z \in Succ(u)} w(u,z)} & otherwise \end{cases}$$

The recommendations $RW(G)$ are defined as $RW(G)(u) = r_u$.

Fig 3 shows the result of RW on the example network from Fig.1.

[2] The Positive Response axiom in the binary setting of [6] can be derived from EM in our setting.

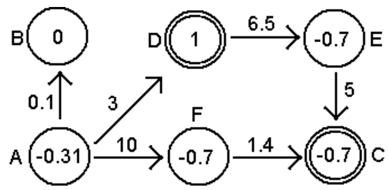

Fig. 3. Here, the votes and the recommendations of RW appear inside the circles.

This system is a direct adaptation of the Random Walk system from [6] to the continuous setting (in the binary setting, the recommendation of u was -1, 0 or 1 according to the sign of r_u, which is neither continuous nor consistent). It is easy to see that the RW recommendation system satisfies MA, NM and EM.

3 IIS Axiom

The following axiom makes things easier, because it allows us to restrict attention to "local" functions:

Independence of Irrelevant Stuff (IIS) encompasses the following two requirements:

1. Let $z \in N$ be a node not reachable from node u. Then for the subgraph G' in which node z and all its associated edges were removed, $R_{G'}(u) = R_G(u)$.
2. Let $G = (N, V, E, v, w)$ and $e \in V \times N$ an edge leaving a voter. Then for the subgraph $G' = (N, V, E \setminus \{e\}, v, w \setminus \{(e, w(e))\})$ in which e has been removed, $R_{G'} = R_G$.

The first requirement demands that a path of trust exist between a node u and a given voter in order for u to be influenced in any way by that voter – a condition consistent with what we expect from trust-based recommendation systems: we want to ignore non-trusted opinions. The second requirement captures the intuition that an agent trusts his own opinion infinitely more than opinions of others. It is easy to see that RW satisfies IIS; in section 7, we will present some systems that satisfy MA+NM+EM, but not IIS[1]. [3]

Proposition 1. *A recommendation system R that satisfies MA, IIS and NC can be written as: $R_G(u) = F\begin{pmatrix} v_1, \ldots, v_n \\ w_1, \ldots, w_n \end{pmatrix}$ for each non-voter u, where the v_i's are the recommendations/votes of children of u, w_i's are the appropriate weights (formally: for each z_i s.t. $(u, z_i) \in E$, $v_i = R_G(z_i)$, $w_i = w(u, z_i)$), and F is an aggregate function of $v_1, \ldots, v_n, w_1, \ldots, w_n$ (also denoted $F(\overrightarrow{v}, \overrightarrow{w})$) that satisfies the following:*

1. *$F(\overrightarrow{v}, \overrightarrow{w}) \in [-1, 1]$ for all $v_1, \ldots, v_n \in [-1, 1]$ and $w_1, \ldots, w_n \in \Re^{\geq 0}$.*

[3] We conjecture that IIS[2] can be derived from MA+NM+EM, but we have yet to find a proof.

2. $F\begin{pmatrix} v_1,\ldots,v_n \\ w_1,\ldots,w_n \end{pmatrix} = F\begin{pmatrix} v_{\pi(1)},\ldots,v_{\pi(n)} \\ w_{\pi(1)},\ldots,w_{\pi(n)} \end{pmatrix}$, *where π is any permutation of* $\{1,\ldots,n\}$.
3. $F(\overrightarrow{v},\overrightarrow{w}) = -F(-\overrightarrow{v},\overrightarrow{w})$
4. $F(\overrightarrow{v},\overrightarrow{0}) = 0$
5. $F\begin{pmatrix} v_1,v_2,\ldots,v_n \\ 0,w_2,\ldots,w_n \end{pmatrix} = F\begin{pmatrix} v_2,\ldots,v_n \\ w_2,\ldots,w_n \end{pmatrix}$
6. *F is continuous everywhere except possibly at points where* $\overrightarrow{w} = \overrightarrow{0}$

Proof. First we will show that $R_G(u)$ depends only on the weights of outgoing edges from u and the respective recommendations of children of u. Suppose that this is not the case. Then, there exist two voting networks, G_1 and G_2, non-voters u in G_1 and v in G_2, such that the local vicinity of u and v under R are the same, but u and v receive different recommendations. Formally, if u has k neighbors $u_1,\ldots,u_k$ in G_1, then v has k neighbors $v_1,\ldots,v_k$ in G_2, and there exists a permutation π of $1,\ldots,k$ s.t. for every $1 \le i \le k$ $w(u,u_i)_{G_1} = w(v,v_{\pi(i)})_{G_2}$ and $R_{G_1}(u_i) = R_{G_2}(v_{\pi(i)})$, but $R_{G_1}(u) \neq R_{G_2}(v)$. Applying NC k times on G_1 results in a voting network G'_1 where the nodes $u_1,\ldots,u_k$ are voters with votes $R_{G_1}(u_1),\ldots,R_{G_1}(u_k)$ respectively, and the recommendation of u is unchanged: $R_{G_1}(u) = R_{G'_1}(u)$. Similarly, we apply NC k times on G_2 to get G'_2 where $v_1,\ldots,v_k$ are voters with votes $R_{G_2}(v_1),\ldots,R_{G_2}(v_k)$ respectively, and the recommendation of v is unchanged: $R_{G_2}(v) = R_{G'_2}(v)$. In particular, $R_{G'_1}(u) \neq R_{G'_2}(v)$. We now repeatedly apply IIS[2] to G'_1 to remove all outgoing edges from $u_1,\ldots,u_k$, until we get a network $\widehat{G_1}$ where $u_1,\ldots,u_k$ are sinks; by IIS[1], $R_{G'_1}(u) = R_{\widehat{G_1}}(u)$. Similarly, we apply IIS[2] to G'_2 to get $\widehat{G_2}$ where $v_1,\ldots,v_k$ are sinks; $R_{G'_2}(v) = R_{\widehat{G_2}}(v)$, and therefore $R_{\widehat{G_1}}(u) \neq R_{\widehat{G_2}}(v)$. In $\widehat{G_1}$, only the nodes $u_1,\ldots,u_k$ are reachable from u; therefore, from IIS[1], all other nodes and their associated edges can be removed to create $\widetilde{G_1}$, consisting only of $u,u_1,\ldots,u_k$, where $R_{\widetilde{G_1}}(u) = R_{\widehat{G_1}}(u)$; similarly, applying IIS[1] to $\widehat{G_2}$ we get $\widetilde{G_2}$ consisting only of nodes $v,v_1,\ldots,v_k$, where $R_{\widetilde{G_2}}(v) = R_{\widehat{G_2}}(v)$. Therefore, we still have $R_{\widetilde{G_1}}(u) \neq R_{\widetilde{G_2}}(v)$. But $\pi \cup \{(u,v)\}$ shows that $\widetilde{G_1}$ and $\widetilde{G_2}$ are isomorphic graphs; therefore, from Anonymity, we should have $R_{\widetilde{G_1}}(u) = R_{\widetilde{G_2}}(v)$ – contradiction.

Therefore, there exists an underlying aggregate function F of votes and weights satisfying 1. Properties 2,3,5 follow directly from MA[1,2,3] respectively; property 4 follows from IIS[1], Neutrality and MA[3]: from IIS[1] and Neutrality we have that recommendations of all sinks are 0, and by MA[3] adding edges with weight 0 does not change recommendations. Finally, property 6 of F follows from MA[5]. □

4 Scaling Votes and Trust

We have already seen one recommendation system that satisfies MA, NM, EM and IIS; the following propositions will enable us to create many more such systems. All of them introduce transformations on recommendation systems that preserve (some of) their desired properties.

More formally, let T be a transformation on the space of recommendation systems. We say that T *preserves a set of properties* X if, for every recommendation system R

satisfying *all* the properties in X, the recommendation system $T(R)$ satisfies all these properties as well.

The following transformation can be used to give more (or less) weight to "radical" opinions (those closer to the ends of the $[-1,1]$ scale):

Proposition 2. *Let $f : [-1,1] \rightarrow [-1,1]$ be a continuous monotone anti-symmetric onto function. Then, the transformation $T_f(R) = f^{-1} \circ R \circ f$ (meaning: apply f to all the votes in G, then apply R, then apply f^{-1} on the resulting recommendations) preserves $\{MA, NM, EM\}$ and $\{MA, NM, EM, IIS\}$*

Proof. Anonymity is obviously preserved, since $T_f(R)$ does not rely on node names anywhere; preservation of Neutrality of $T_f(R)$ follows from anti-symmetry of f; preservation of MA[3,4] follows directly from the definition of $T_f(R)$; MA[5] follows from continuity of f, and IIS follows directly from the definition of $T_f(R)$.

NM: w.l.o.g. suppose we change a non-voter u in G with the recommendation of $r = T_f(R)_G(u) = f^{-1}(R_{f(G)}(u))$ into a voter u' in G' with a vote $r' > r$ (abusing notation, $f(G)$ here means the network G where f is applied to all voters). Applying f to both sides of $r' > r$, we get $R_{f(G')}(u) > R_{f(G)}(u)$; therefore, due to NM of R, for each parent v of u, $R_{f(G')}(v) > R_{f(G)}(v)$. Due to monotonicity of f^{-1}, this implies $T_f(R)_{G'}(v) > T_f(R)_G(v)$.

Similarly, EM follows directly from monotonicity of f^{-1}. Note that we did not rely on IIS anywhere except in the proof of preservation of IIS itself, meaning that T_f preserves $\{MA, NM, EM\}$ separately as well. □

Note that a continuous monotone function is also one-to-one. If we want to use a function f which is not onto, we need an additional assumption in order for the transformation to preserve the desired properties (without an additional assumption, f^{-1} might not be well defined). We call this assumption Conservativeness:

Conservativeness: A recommendation system R is said to satisfy Conservativeness if for every voting network G and non-voter u with a single positive weighted outgoing edge to a voter z (w.l.o.g. $v(z) \geq 0$), $0 \leq R_G(u) \leq v(z)$.

Conservativeness states that a recommendation for an agent cannot be more radical than the vote of someone he trusts – the value can only be altered in the direction of uncertainty. We can show that conservative systems satisfy the natural betweenness property – the recommendation of a node falls between the votes of his neighbors, possibly skewed towards uncertainty:

Proposition 3. *Let R be a recommendation system that satisfies MA, NC, EM, IIS and Conservativeness. For a voting network G and $u \in N$, we define $R_{Succ(u)} = \{R_G(z) | z \in Succ(u)\}$. Then, for every voting network G and every $u \in \overline{V}$,*
$\min\{R_{Succ(u)} \cup \{0\}\} \leq R_G(u) \leq \max\{R_{Succ(u)} \cup \{0\}\}$
Moreover, when $|R_{Succ(u)}| > 1$, the inequalities are strict:
$\min\{R_{Succ(u)} \cup \{0\}\} < R_G(u) < \max\{R_{Succ(u)} \cup \{0\}\}$

Proof. We use induction of the amount of neighbors of u: we start with u being a sink, and add neighbors one by one, sorted by increasing order of votes, showing that in each step the recommendation of u stays strictly between the previous recommendation and the new added vote. (We rely on NC and IIS in order to treat voter and non-voter neighbors similarly – these assumptions can be relaxed). With the addition of the first vote, betweenness follows directly from Conservativeness. Now, for the induction step, suppose the current recommendation of u is r, and we add a trust link of weight w to a node with a recommendation/vote $r' > r$. According to EM, the recommendation of u should increase; suppose, for contradiction, that now betweenness does not hold, and the new recommendation of u, $r'' > r'$. Since this trust link was not the first we added, $R(u)$ changed continuously; meaning, that there exists $0 < w' < w$ s.t. when we add a link of weight w' from u to r', the resulting recommendation is exactly r' (the middle value theorem). But then, by EC, increasing trust from w' to w should not affect recommendations – contradiction. Note that we needed the Conservativeness assumption only because in MA[5] we don't require R to be continuous in the point with no outgoing trust links at all – if R is continuous everywhere, this assumption is not needed. □

The RW system obviously satisfies Conservativeness. If we begin with a conservative system, we can show a wider family of transformations that preserve our desired properties:

Proposition 4. *Let $f : [-1, 1] \to [-1, 1]$ be a continuous monotone anti-symmetric function. Then, the transformation $T_f(R) = f^{-1} \circ R \circ f$ preserves $\{MA, NM, EM, IIS, Conservativeness\}$.*

Proof. (Sketch) From IIS and prop. 3 above, $T_f(R)$ is well defined, even if f is not onto; therefore, MA, NM, EM and IIS are preserved by prop. 2. Preservation of Conservativeness follows directly from monotonicity of f. □

It is an interesting open question whether the Conservativeness requirement is indeed necessary for the proof – meaning, whether there exist recommendation systems satisfying MA, NM, EM and IIS, but not Conservativeness.

The following transformation can be used to specify what contributes more to the decision – many neighbors with small trust values or few neighbors with larger trust values (given that the total trust is the same in both cases):

Proposition 5. *Let $g : \Re^+ \to \Re^+$ be a continuous monotone increasing function, with $g(0) = 0$. Then, the transformation $T_g(R) = R \circ g$ (meaning: apply g to all the weights in G, then apply R) preserves $\{MA, NM, EM\}$ and $\{MA, NM, EM, IIS\}$.*

Proof. (Sketch) The proof follows directly from properties of R and g; the assumption $g(0) = 0$ is required in order for MA[3] to hold. □

5 Axiomatic Characterization Of RW Systems

In this section we wish to characterize the family of recommendation systems that results from the closure of RW under the transformations T_f and T_g, for any f and g. In order to do that, we need to define some additional properties of recommendation systems. We start with some intuition:

One desirable requirement is the transitivity of trust relations (appeared originally in [6]). Let G be a graph, v a node, and A, B two disjoint sets of nodes. We say that v *trusts A more than B* if the recommendation of v in a voting network on G where all nodes in A vote 1, all nodes in B vote -1, and no other nodes vote, is positive.

Transitivity (TR): We say that a recommendation system R is transitive if, for every graph G, node v and disjoint sets A, B, C, if v trusts A more than B and v trusts B more than C, then v trusts A more than C.

[6] showed, for the binary setting, that Transitivity is impossible to reconcile with IIS and their versions of monotonicity. This still holds in the continuous case:

Proposition 6. *The requirements MA[1,2,3], IIS, either one of NM or EM, and Transitivity are inconsistent.*

Proof. Consider the graph in Fig.4.

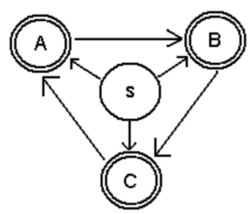

Fig. 4. Given MA[1,2,3], IIS and either one of NM or EM, Transitivity cannot hold.

Suppose our system satisfies Anonymity, Neutrality, IIS and either one of NM or EM; we prove that s trusts A more than B: suppose that A votes 1, B votes -1, and C doesn't vote. From IIS, the recommendation of C depends only on the vote of A. If we assume NM, then the recommendation of C has to be positive: if A voted 0, the recommendation of C would be 0 because of Neutrality; then, as A increases the vote from 0 to 1, the recommendation of C should increase. Similarly, if we have EM, then C should receive a recommendation of 0 if the edge (C, A) was not present (Neutrality+MA[3]); therefore when the edge (C, A) is added, by EM the recommendation of C should increase. So in both cases we see that C receives a positive recommendation. From IIS, the edge (B, C) can be discarded. If the edge (s, C) was not present, s would receive a recommendation of 0 (IIS+Neutrality); then, from EM, adding an edge (s, C) should strictly increase the recommendation. In case of NM, the reasoning is similar: if C was changed into a voter, the recommendation of s would not change (NC); then, the edge (C, A) can be discarded; now, if the vote of C was 0, then s would receive a recommendation of 0 due to Neutrality, because the graph would be symmetric; when a vote of C is increased, so does the recommendation of s.
Similarly, we can show that s trusts B more than C and C more than A, therefore TR does not hold. [4] □

[4] The example is from [6], but in their setting it couldn't be used for an impossibility proof.

So the Transitivity requirement needs to be weakened, if we want any hope of accommodating it. In the binary setting, [6] had the axiom of Trust Propagation: if u trusts v with a weight k, and v trusts (only) nodes $v_1, \ldots, v_k$ with a weight of 1, then the edge (u, v) can be replaced with k edges (u, v_i), with no recommendations being affected.

The direct translation of this axiom into the continuous setting would have the weight on the edge (u, v) equal to the sum of the weights on (v, v_i), but we found it to be too restrictive. We did not want to limit the way the trust weights are interpreted (does a trust link of weight 2 "matter" exactly as much as two trust links of weight 1?). Therefore, the axiom we chose retains the spirit of Trust Propagation, while avoiding to specify the exact way of translating the trust weights:

Separability: Let $\oplus$ be an operator $\oplus : \Re^{\geq 0} \times \Re^{\geq 0} \to \Re^{\geq 0}$ satisfying:

1. $w \oplus 0 = w$
2. Associativity: $(w_1 \oplus w_2) \oplus w_3 = w_1 \oplus (w_2 \oplus w_3)$
3. $\oplus$ is a continuous, strictly monotone increasing function in both variables

We say that a recommendation system R is *separable* if, for any voting network G, if a node v trusts only nodes v_1, v_2 with respective trust weights of w_1, w_2, and a node u trusts v with a weight of $(w_1 \oplus w_2)$, then the edge (u, v) can be replaced with edges (u, v_1),(u, v_2) with weights w_1, w_2, and no recommendations will change (see Fig. 5).

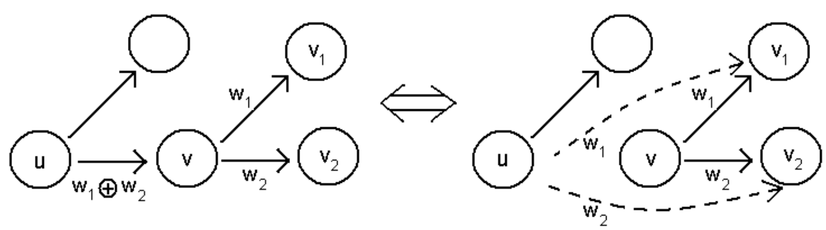

Fig. 5. Separability axiom.

The RW system obviously satisfies Separability, with $\oplus = +$.

Next, we have a more strict variant of Conservativeness:

Neighborhood Consensus (NCS): A recommendation system R is said to satisfy Neighborhood Consensus if for every voting network G and non-voter u with a single positive weighted outgoing edge to a voter z, $R_G(u) = v(z)$.

Note that NCS+EM imply that if all neighbors of a node agree on their votes, the recommendation for that node equals these votes (hence the name). Despite being more strict than Conservativeness, this property is very natural and appears in many social choice systems. The RW system also satisfies this requirement.

Now we can formulate the characterization result:

Theorem 1. *A recommendation system R can be defined as the unique solution of the system of equations:*

$$r_u = \begin{cases} 0 & \textit{no path exists from } u \textit{ to a voter} \\ v(u) & u \in V \\ f^{-1}\left[\frac{\sum_{z \in Succ(u)} g(w(u,z)) f(r_z)}{\sum_{z \in Succ(u)} g(w(u,z))}\right] & otherwise \end{cases}$$

where f is a continuous monotone anti-symmetric function, and g a continuous monotone increasing function, with $g(0) = 0$, if and only if R satisfies the following conditions: MA, IIS, NM, NCS and Separability.

Proof. $\Rightarrow$ Follows from the properties of RW and propositions 2 and 5 (since the recommendation system above is exactly $T_f(T_g(RW))$).
$\Leftarrow$ From prop. 1 we have that $R_G(u) = F\begin{pmatrix} v_1, \ldots, v_n \\ w_1, \ldots, w_n \end{pmatrix}$ for each non-voter u, where the v_i's are the recommendations/votes of children of u, w_i's are the appropriate weights, and F is an aggregate function of $v_1, \ldots, v_n, w_1, \ldots, w_n$.

From prop. 1 and additional properties of R, we also know that the aggregate function F satisfies:

1. $F\begin{pmatrix} v, \ldots, v \\ w_1, \ldots, w_n \end{pmatrix} = v$ (from NCS)
2. $F(\overrightarrow{v}, \overrightarrow{w})$ is continuous wherever $\overrightarrow{w} \neq \overrightarrow{0}$ (from MA[5])
3. $F(\overrightarrow{v}, \overrightarrow{w})$ is strictly monotone increasing in v (from NM)
4. $F\begin{pmatrix} v_1, \ldots, v_n \\ w_1, \ldots, w_n \end{pmatrix} = F\begin{pmatrix} F\begin{pmatrix} v_1, v_2 \\ w_1, w_2 \end{pmatrix} & , v_3, \ldots, v_n \\ w_1 \oplus w_2 & , w_3, \ldots, w_n \end{pmatrix}$
   where $\oplus$ is an operator $\oplus : \Re^{\geq 0} \times \Re^{\geq 0} \to \Re^{\geq 0}$ satisfying:
   (a) $w \oplus 0 = w$
   (b) Associativity: $(w_1 \oplus w_2) \oplus w_3 = w_1 \oplus (w_2 \oplus w_3)$
   (c) $\oplus$ is a continuous, strictly monotone increasing function in both variables

   (from Separability)

Using induction and properties of $\oplus$, the last property of F can be generalized into: For every $r < n, v_1, \ldots, v_n, w_1, \ldots, w_n$

$$F\begin{pmatrix} v_1, \ldots, v_n \\ w_1, \ldots, w_n \end{pmatrix} = F\begin{pmatrix} F\begin{pmatrix} v_1, \ldots, v_r \\ w_1, \ldots, w_r \end{pmatrix} & , v_{r+1}, \ldots, v_n \\ w_1 \oplus \ldots \oplus w_r & , w_{r+1}, \ldots, w_n \end{pmatrix}$$

Details: the basis, $r = 2, n > r$ follows from Separability. We will assume the generalization holds for all $r' < r$ and $n' < n$, and we will show it holds for r, n:

$$F\begin{pmatrix} v_1,\ldots,v_n \\ w_1,\ldots,w_n \end{pmatrix} = F\begin{pmatrix} F\begin{pmatrix} v_1,\ldots,v_{r-1} \\ w_1,\ldots,w_{r-1} \end{pmatrix} & ,v_r,\ldots,v_n \\ w_1 \oplus \ldots \oplus w_{r-1} & ,w_r,\ldots,w_n \end{pmatrix} =$$

$$= F\begin{pmatrix} F\begin{pmatrix} F\begin{pmatrix} v_1,\ldots,v_{r-1} \\ w_1,\ldots,w_{r-1} \end{pmatrix} & ,v_r \\ w_1 \oplus \ldots \oplus w_{r-1} & ,w_r \end{pmatrix} & ,v_{r+1},\ldots,v_n \\ w_1 \oplus \ldots \oplus w_r & ,w_{r+1},\ldots,w_n \end{pmatrix} =$$

$$= F\begin{pmatrix} F\begin{pmatrix} v_1,\ldots,v_r \\ w_1,\ldots,w_r \end{pmatrix} & ,v_{r+1},\ldots,v_n \\ w_1 \oplus \ldots \oplus w_r & ,w_{r+1},\ldots,w_n \end{pmatrix}$$

[14] has shown that all the aggregate functions $\overrightarrow{v}, \overrightarrow{w}$ satisfying the above properties can be written in the form

$$F\begin{pmatrix} v_1,\ldots,v_n \\ w_1,\ldots,w_n \end{pmatrix} = f^{-1}\left[\frac{g(w_1)f(v_1)+\ldots+g(w_n)f(v_n)}{g(w_1)+\ldots+g(w_n)}\right]$$

where $f : [-1,1] \rightarrow [-1,1]$ is a continuous monotone anti-symmetric function and $g : \Re^+ \rightarrow \Re^+$ is a monotone increasing continuous function, with $g(0) = 0$; we only need to use this result to complete our proof. □

We note that each axiom in the above theorem is necessary for the result to hold; namely, there exist examples of systems satisfying all but one of these axioms which could not be written in the above form. Below we shortly discuss each example.

Relaxing MA is not interesting – in particular, it makes no sense to consider systems which are not anonymous. Relaxing NCS is the focus of the next section – in particular, applying prop. 7 below on RW can be used to create examples of systems satisfying MA, NM, IIS and Separability, but not NCS. Section 7 discusses relaxations of IIS: applying prop. 10 there on RW results in systems satisfying MA, NM, NCS and Separability, but not IIS. Relaxing the NM requirement is trivial: a system that always gives recommendations of 0 will do. Relaxing Separability – we apply the following transformation on the graph G: for each node v with outgoing edges $e_1,\ldots,e_k$ and appropriate weights $w_1,\ldots,w_k$, let $t_v = w_1+\ldots+w_k$. We change the weights of $e_1,\ldots,e_k$ to $w_1^{t_v},\ldots,w_k^{t_v}$, respectively (we raise all weights to the power of t_v). Let G' be the graph resulting from the transformation. Our system returns the result of RW on G'. It is easy to see that the system satisfies MA, NM, NCS and IIS, but not Separability.

6 Discount Factors

The following transformation introduces a property that many people will find intuitive and crucial for a recommendation system: a discount factor. Indeed, one of the biggest objections to the RW system (or, indeed, to any system satisfying NCS and IIS) is that it ignores the length of the trust chains – an opinion of someone who is trusted by a friend of a friend of a friend means exactly as much as an opinion of a friend. Various “multiplicative” discount systems exist, which are based on the intuitive idea that our recommendation is the weighted average of our neighbors, multiplied by some discount value $0 < \alpha < 1$; It is easy to see that these systems violate the very principle behind Edge Consistency: if we increase trust in an agent who reached exactly the same opinion on the product as we have, our opinion should not change. In multiplicative discounted trust systems, the opinion will always change in such situation, because the discounted (rather than the real) value is taken into account. In order for EC to hold, the discount weights must be absolute rather than relative:

Proposition 7. *Let $w_0 > 0$. Let $T_{w_0}(R)$ be a recommendation system defined as follows: we add a new voter t with a vote of 0 to the network, and make every node point to t with weight w_0. Formally, for a voting network $G = (N, V, E, v, w)$, let $G' = (N \cup \{t\}, V \cup \{t\}, E' = E \cup \{(u,t) | u \in N\}, v \cup \{(t, 0)\}, w \cup \{(e, w_0) | e \in E' \setminus E\})$. For a node $u \in V$, $T_{w_0}(R)_G(u) = R_{G'}(u)$. Then, the transformation T_{w_0} preserves $\{MA, NM, EM, IIS\}$; if R does* not *satisfy IIS, then MA[4] might be violated by $T_{w_0}(R)$; everything else still holds.*

Proof. Since all changes in votes or edges in G result only in corresponding changes in G', the preservation of NM and EM is immediate. MA[1,2,3] follow from the definition of $T_{w_0}(R)$. The IIS requirement is necessary in proving MA[4]: a new non-voter gets a recommendation of 0 by both systems, and because R is IIS, no other recommendation is affected by its addition. The resulting recommendation system is continuous everywhere – continuity in all points is achieved by simply avoiding the point $\overrightarrow{w} = \overrightarrow{0}$, since there are now no non-voter sinks in G'. Finally, no unreachable nodes were made reachable by the transformation of the network, and no links from voters were taken into account, therefore IIS is preserved. □

However, the distinction of relative vs. absolute discount factors can also be viewed as a shortcoming of EM. Indeed, it all boils down to how do we want our system to behave on the following simple example: a single non-voter v, pointing to a voter u with a vote of 1. Suppose our system has discounts, so that the recommendation of v is r, with $r < 1$. What would be a reasonable behavior if we now add an edge from v to a voter x with a vote r', with $r < r' < 1$? From one point of view, the recommendation of v should increase, since v has now increased his trust in someone who reached a higher opinion on the product than v. This is the point of view on which we based our formulation of EM. But on the other hand, it can also be said that the recommendation of v should decrease, since v has now increased trust in someone with a lower opinion on the product than everyone v trusted so far! One formal expression of such intuition is as follows:

Edge Monotonicity 2 (EM2): For any voting network G and a non-sink agent $u \in \overline{V}$, let G' be a graph resulting from G when u increases his trust in agent v. Then, if $R_G(v) \geq R_G(x)$ for every neighbor x of v, then $R_{G'}(v) \geq R_G(v)$; if at least one of these inequalities is strict, then $R_{G'}(v) > R_G(v)$.

So we have two valid, but conflicting intuitions: the first (EM) is satisfied by absolute discount systems, while the second (EM2) is satisfied by multiplicative discount systems. The only systems which conform to both of these natural requirements are the systems satisfying NCS, i.e. having no discounts at all – such as the systems characterized in the previous section.

We also note that if we replace EM by EM2 in propositions 2, 4 and 5, the results still hold.

7 Including The Aggregated Network Vote

IIS is a debatable requirement; even in trust-based systems, it seems natural to consider opinions of unrelated people. For example, suppose I don't know anyone in the system; still, if all 1000 people in the system rated the item at 1, it seems plausible that I should receive a positive recommendation. The following transformation allows us to create recommendation systems that take the aggregated network vote into account, thereby breaking IIS.

Proposition 8. *Let R be a recommendation system that satisfies MA, NM, EM, IIS and NCS and let $w_A > 0$. Let $T_{w_A}(R)$ be a recommendation system defined as follows: for a voting network $G = (N, V, E, v, w)$, let G' be a network with an additional non-voter t, who points to every node in the graph, and every node in the graph points back to him with a trust of w_A. Formally, $G' = (N \cup \{t\}, V, E' = E \cup \{(u, t)|u \in N\} \cup \{(t, u)|u \in N\}, v, w \cup \{(e, w_A)|e \in E' \setminus E\})$. Then, for a node $u \in \overline{V}$, define $T_{w_A}(R)_G(u) = R_{G'}(u)$. Then, the recommendation system $T_{w_A}(R)$ satisfies MA, NM and EM; $T_{w_A}(R)$ does not satisfy IIS and NCS; if R did* not *satisfy NCS, then $T_{w_A}(R)$ might violate MA[4] (but everything else still holds).*

Proof. All the desired properties except for MA[4] follow directly from the properties of R, since all changes in votes or edges in G map to corresponding changes in G'. Extra care should be taken to ensure that MA[4] still holds, since new edges are added to G' when a new non-voter is added to G. Here we require the NCS assumption: since R satisfies MA[4], the addition of a new non-voter z will not change recommendations; because of NCS, when the edge (z, t) is added, the recommendation of z will equal the recommendation of t; furthermore, since the new node and the associated edge are unreachable from anywhere else, no recommendations will change (IIS). Now, due to EC of R, adding the edge (t, z) will not change the recommendation of t; and again, due to IIS of R, no other recommendations will change as a result. □

As we see from the proofs of prop. 8 and 7, trying to enjoy both worlds (having discounts and integrating aggregated network votes) seems to necessarily violate our MA[4] requirement: intuitively, when a new uninformed member is added to the network, his recommendation must be the aggregated vote of the network (call it $\mathcal{V}$); because of the discounting, his opinion is closer to 0 than $\mathcal{V}$; since his opinion is integrated back into the network (due to EM), the integrated opinion of the network also changes, thereby affecting existing recommendations.

The above criticism of the IIS requirement does not necessarily imply that it should be discarded; rather, it can be used to identify a shortcoming of the current formulation of IIS. The problem lies in the implication that an uninformed agent should always receive a recommendation of 0, regardless of the votes and trust that exist in the rest of the system. On the other hand, we don't want to add implicit trust links to agents who already have trust in other agents; after all, the intuition behind trust-based recommendation systems is that an agent can specify exactly which agents he personally trusts. Consider the following variant of IIS:

Independence of Irrelevant Stuff (Weakened) (IISW): For any voting network G and agent $u \in \overline{V}$, let $v_1, \ldots, v_k$ be the neighbors of u, with trust values of $w_1, \ldots, w_k$. If $k \geq 1$ and $v_i \in V$ for all $1 \leq i \leq k$, then the recommendation of u depends solely on $v_1, \ldots, v_k, w_1, \ldots, w_k$.

In a sense, this weaker requirement still conforms to the intuition behind the original IIS; in fact, proposition 1 holds for IISW almost just as well as for IIS:

Proposition 9. *A recommendation system R that satisfies MA, IISW and NC can be written as: $R_G(u) = F\begin{pmatrix} v_1, \ldots, v_n \\ w_1, \ldots, w_n \end{pmatrix}$ for each non-voter u which is not a sink, where the v_i's are the recommendations/votes of children of u, w_i's are the appropriate weights, and F is an aggregate function of $v_1, \ldots, v_n, w_1, \ldots, w_n$ that satisfies the following:*

1. $F(\overrightarrow{v}, \overrightarrow{w}) \in [-1, 1]$ *for all* $v_1, \ldots, v_n \in [-1, 1]$ *and* $w_1, \ldots, w_n \in \Re^{\geq 0}$.
2. $F\begin{pmatrix} v_1, \ldots, v_n \\ w_1, \ldots, w_n \end{pmatrix} = F\begin{pmatrix} v_{\pi(1)}, \ldots, v_{\pi(n)} \\ w_{\pi(1)}, \ldots, w_{\pi(n)} \end{pmatrix}$*, where* π *is any permutation of* $\{1, \ldots, n\}$.
3. $F(\overrightarrow{v}, \overrightarrow{w}) = -F(-\overrightarrow{v}, \overrightarrow{w})$
4. $F\begin{pmatrix} v_1, v_2, \ldots, v_n \\ 0, w_2, \ldots, w_n \end{pmatrix} = F\begin{pmatrix} v_2, \ldots, v_n \\ w_2, \ldots, w_n \end{pmatrix}$ *for any* $n > 1$.
5. *F is continuous everywhere except possibly at points where* $\overrightarrow{w} = \overrightarrow{0}$

Proof. The proof follows the proof of prop. 1, only now we can use IISW to reach the result already after the first transformation, from G_1, G_2 to G_1', G_2'. Note also that now we need u to be non-sink in order to be able to use IISW. □

So, a more natural way of integrating the aggregated network vote into our recommendation system would use it only for agents who have no specified trust (breaking IIS, but not IISW):

Proposition 10. *Let $w_A > 0$, and let T_{w_A} be a transformation defined as follows: for a voting network $G = (N, V, E, v, w)$, let $S \subseteq N$ be the set of all* sinks*, and let G' be as in prop. 8, but with only sinks pointing to the new node t. Formally, $G' = (N \cup \{t\}, V, E' = E \cup \{(u, t) | u \in S\} \cup \{(t, u) | u \in N\}, v, w \cup \{(e, w_A) | e \in E' \setminus E\})$. For a node $u \in \overline{V}$, $T_{w_A}(R)_G(u) = R_{G'}(u)$. Then, T_{w_A} preserves $\{MA, NM, EM, IISW, NCS\}$.*

Proof. (Sketch) The proof resembles the proof of prop. 8; one should take extra care with showing EM, since the case of a new edge leaving a previous sink must now be checked separately. In this case EM follows directly from NCS assumption. □

Note that if R did *not* satisfy NCS, then $T_{w_A}(R)$ might violate not only MA[4], but EC as well, as illustrated by the example in Fig. 6. Here, if we increase the trust of A in B by ε, the recommendation of A will decrease, because the new edge (A, B) takes place of the implicit edge (A, t), and then discounting takes place.

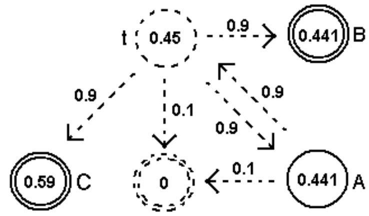

Fig. 6. The recommendations are given by the system $T_{w_A}(T_{w_0}(RW))$, where $w_A = 0.9$ and $w_0 = 0.1$.

We should also note that adding a discount (edge to vote 0) *only* to the new node t is also possible – if we make this adjustment in propositions 8, 10, the results still hold. The resulting recommendation systems will take the discounted aggregated network vote into consideration; a result that cannot be achieved by decreasing the w_A constant.

7.1 Considering Opinions of Voters

The recommendation systems suggested by prop. 9 and prop. 10 include a new "network node" into the graph – a new node which trusts all agents equally. It could be argued that the network node does not fully consider the existing trust of the system; in particular, it disregards the opinions of voters. Indeed, all the systems we have seen so far exhibit this property; on one hand, this property was natural enough to include it as part of IIS axiom (IIS[2]); on the other hand, it can be argued that the intuition that led us to formulate this requirement is still present in the weaker version of IIS, IISW. The intuition was that an agent's own opinion weights infinitely more to him than the opinion of his neighbors. IISW still has this property: any non-sink agent's recommendation is still a local function. However, if an agent trusts no one in the system, then why shouldn't he try to aggregate the trust of voters as well as any other data in the system?

An intuitive way to modify the transformation T_{w_A} in prop. 10 above to include links from voters is by aggregating them into the outgoing edges of the system node, t: instead of making t point to all nodes in the graph with equal weight, its outgoing edges can aggregate the trust that already exists in the system – for example, by pointing to each node with the summary weight of its incoming edges.

Consider the following recommendation system, R: for any voting network G, let G' be the graph that results from G by adding a new node t with the following set of outgoing edges: $\{(t, u, w) | w = \sum_{e=(t,u) \in E} w(e)\}$ and the following set of incoming edges:
$\{(u, t, 1) | u \in S\}$. We then define $R_G(u) = RW_{G'}(u)$.

Proposition 11. *R does not satisfy Edge Consistency.*

Proof. Consider the network in Fig.7 below:

With the edge (A, B) not present, $R(A) = -0.5$. According to EC, if we increase the trust (A, B) from 0 to some positive value, the recommendation of A should not change. But it is easy to see that $R(A)$ decreases as a result, since the new edge will be added to the new node t as well, causing the recommendation of A to decrease. □

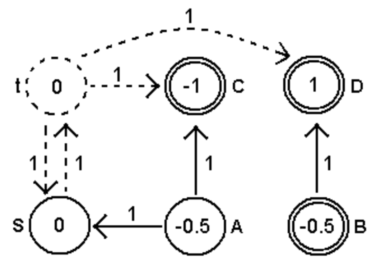

Fig. 7. Here, the votes and the recommendations of R appear inside the circles. The solid lines represent the original graph, G; the dashed lines represent the additions of G' to G.

But, in a sense, the above violation of EC does not conform to the intuition of what EC represents. After all, the recommendation of S decreased as a result of the added trust from A to B; since A trusts S, it is reasonable that A's recommendation decreases as well. So, in a way, this explains why the intuition of EM implicitly assumed IIS – EM makes little sense if IIS doesn't hold. Consider the following weakened version if EM:

Edge Monotonicity Weakened (EMW): For any voting network G, agent $u \in \overline{V}$, edge $(u, z) \in E$, and $w' > w(u, z)$ let $G' = (N, V, E, v, w \setminus \{((u, z), w(u, z))\} \cup \{((u, z), w')\})$. Then, if it holds that for every neighbor x of u, $R_{G'}(x) = R_G(x)$, then $sgn(R_{G'}(u) - R_G(u)) = sgn(R_G(z) - R_G(u))$.

EMW is simply EM restricted for the cases where adding trust does not change local recommendations. Note that EMW implicitly assumes IISW, in exactly the same way that EM implicitly assumed IIS - it makes sense only if the recommendations of u depend only on u's local vicinity.

Proposition 12. *For any recommendation system R which satisfies IISW, NCS and Node Consistency, EMW is equivalent to the following: for any voting network G, agent $u \in \overline{V}$ such that every neighbor of u is a voter, edge $(u, z) \in E$, and $w' > w(u, z)$ let $G' = (N, V, E, v, w \setminus \{((u, z), w(u, z))\} \cup \{((u, z), w')\})$. Then $sgn(R_{G'}(u) - R_G(u)) = sgn(R_G(z) - R_G(u))$.*

Proof. The alternative formulation of EMW is a particular case of the original – if all neighbors x of u are voters, then, obviously, $R_{G'}(x) = R_G(x)$. Suppose that the alternative formulation holds for R. Let G be a voting network, with agent $u \in \overline{V}$, edge $(u, z) \in E$, and $w' > w(u, z)$; let $G' = (N, V, E, v, w \setminus \{((u, z), w(u, z))\} \cup \{((u, z), w')\})$. Suppose that for every neighbor x of u, $R_{G'}(x) = R_G(x)$. We consider two separate cases: if u is a sink in G, then because of NCS $R_{G'}(u) = R_{G'}(z) = R_G(z)$, and therefore $sgn(R_{G'}(u) - R_G(u)) = sgn(R_G(z) - R_G(u))$. Otherwise, consider the networks $\overline{G}$ and $\overline{G'}$ that result from G and G' when all the neighbors of u are changed into voters, with votes equal to their respective recommendations. From NC, $R_{\overline{G}}(u) = R_G(u)$ and $R_{\overline{G'}}(u) = R_{G'}(u)$. From IISW, we can eliminate all nodes except for u and his neighbors, and get corresponding networks $\widetilde{G'}$ and $\widetilde{G}$, so that the recommendations of u remain unchanged. But $\widetilde{G'}$ results from $\widetilde{G}$ by increasing trust of u in v, when all neighbors of u are voters; therefore, we can use the alternative formulation of EMW to get $sgn(R_{\widetilde{G'}}(u) - R_{\widetilde{G}}(u)) = sgn(R_{\widetilde{G}}(z) - R_{\widetilde{G}}(u))$, and therefore $sgn(R_{G'}(u) - R_G(u)) = sgn(R_G(z) - R_G(u))$. □

It is easy to see that the recommendation system R in prop. 11 satisfies $IISW$ and EMW. We can now formulate a more general result:

Proposition 13. *Let $\oplus$ be an operator $\oplus : \Re^{\geq 0} \times \Re^{\geq 0} \rightarrow \Re^{\geq 0}$ satisfying:*

1. $w \oplus 0 = w$
2. *Associativity:* $(w_1 \oplus w_2) \oplus w_3 = w_1 \oplus (w_2 \oplus w_3)$
3. $\oplus$ *is a continuous, strictly monotone increasing function in both variables*

Let $T_\oplus$ be a transformation defined as follows: for a voting network $G = (N, V, E, v, w)$, let $S \subseteq N$ be the set of all sinks, and let G' be the graph resulting from G when a new node t is added to $\overline{V}$, with incoming edges (u, t) weighted 1 from every $u \in S$, and outgoing edges (t, u) for every $u \in N$, weighted $\oplus_{(z,u)\in E} w(z, u)$. For a node $u \in \overline{V}$, define $T_\oplus(R)_G(u) = R_{G'}(u)$.
Then, if a recommendation system R satisfies MA, IIS, NM, EM and NCS, then $T_\oplus(R)$ satisfies MA, IISW, NM, EMW and NCS.

Proof. (Sketch) MA[1,2,3] follow directly from the definition of $T_\oplus$. MA[4] follows from the combination of NCS, NC, EC and IIS, using the same reasoning as in proof of prop. 8. MA[5] follows from continuity of R and $\oplus$. IISW, NM and EMW follow directly from IIS, NM and EM of R. □

8 Conclusions and Further Work

This paper explored continuous trust based recommendation systems. Our main requirement was consistency - we focused on recommendation systems which rely on their own predictions. First, we showed the "basic" such system – Random Walk – and then we introduced a few transformations that customize recommendation systems according to specific needs: give more or less weight to radical opinions, tweak the importance of having multiple trust links of small weights vs. few trust links of large weights, tweak the "damping factor" (discount) of trusted opinions, and consider the aggregated opinion of the network. All along we used the axiomatic approach – we started by formalizing our requirements from recommendation systems, and then identified families of recommendation systems that conform to these requirements.

All the recommendation systems in this paper were some version of Random Walk applied to a transformed voting network. A central question is whether this focus is justified. In section 5 we showed that if we want our system to conform to some additional requirements, then, indeed, we have no choice but to use a generalized version of Random Walk. But what if we consider the most general possible recommendation system? Are there any continuous consistent recommendation systems *at all* that are not fundamentally based on Random Walk? We strongly suspect that the answer is negative; however, the concept of being "fundamentally based on Random Walk" is not easy to formalize. Attempts of full characterization of consistent continuous recommendation systems are in the focus of our current research.

References

1. A. Altman and M. Tennenholtz. On the axiomatic foundations of ranking systems. In *Proc. 19-th International Joint Conference on Artifcial Intelligence*, pages 917–922, 2005.
2. A. Altman and M. Tennenholtz. The pagerank axioms. In *Proceedings of the 6-th ACM conference on Electronic Commerce*, 2005.
3. A. Altman and M. Tennenholtz. An axiomatic approach to personalized ranking systems. In *Proc. 20-th International Joint Conference on Artifcial Intelligence*, 2007.
4. Alon Altman and Moshe Tennenholtz. Quantifying incentive compatibility of ranking systems. In *Proc. of AAAI-06*, 2006.
5. Alon Altman and Moshe Tennenholtz. Incentive compatible ranking systems. In *Proceedings of the 2007 International Conference on Autonomous Agents and Multiagent Systems (AAMAS-07)*, 2007.
6. R. Andersen, C. Borgs, J. Chayes, U. Feige, A. Flaxman, A. Kalai, V. Mirrokni, and M. Tennenholtz. Trust-based recommendation systems: an axiomatic approach. In *Proceedings of WWW-08*, 2008.
7. Paolo Avesani, Paolo Massa, and Roberto Tiella. A trust-enhanced recommender system application: Moleskiing. In *SAC '05: Proceedings of the 2005 ACM symposium on Applied computing*, pages 1589–1593, New York, NY, USA, 2005. ACM Press.
8. Y. Bakos and C. N. Dellarocas. Cooperation without enforcement? a comparative analysis of litigation and online reputation as quality assurance mechanisms. MIT Sloan School of Management Working Paper No. 4295-03, 2003.
9. Allan Borodin, Gareth O. Roberts, Jeffrey S. Rosenthal, and Panayiotis Tsaparas. Link analysis ranking: algorithms, theory, and experiments. *ACM Trans. Inter. Tech.*, 5(1):231–297, 2005.
10. Alice Cheng and Eric Friedman. Sybilproof reputation mechanisms. In *P2PECON '05: Proceeding of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems*, pages 128–132, New York, NY, USA, 2005. ACM Press.
11. R. Dash, S. Ramchurn, and N. Jennings. Trust-based mechanism design. In *Proceedings of the Third International Joint Conference on Autonomous Agents and MultiAgent Systems*, pages 748–755, 2004.
12. T. Haveliwala. Topic-sensitive pagerank. In *WWW 02: Proceedings of the 11th international conference on World Wide Web*, pages 517–526, 2002.
13. D. Houser and J. Wooders. Reputation in auctions: Theory and evidence from ebay. Working Paper, University of Arizona, 2000.
14. T. Kitagawa. On some class of weighted means. In *Proceedings of the Physico-Mathematical Society of Japan*, volume 16, 1934.
15. Jon M. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)*, 46(5):604–632, 1999.
16. Raph Levien. *Attack Resistant Trust Metrics*. PhD thesis, University of California, Berkeley, 2002.
17. P. Massa and P. Avesani. Controversial users demand local trust metrics: An experimental study on epinions.com community. In *Proc. of AAAI-05*, pages 121–126, 2005.
18. L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web. Technical Report, Stanford University, 1998.
19. P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. *The Economics of the Internet and E-Commerce, M.R. Baye*, 2002.
20. P. Resnick, R. Zeckhauser, R. Friedman, and E. Kuwabara. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.

21. Ajit P. Singh, Asela Gunawardana, Chris Meek, and Arun C. Surendran. Recommendations using absorbing random walks. In *North East Student Colloquium on Artificial Intelligence (NESCAI)*, 2007.
22. Giora Slutzki and Oscar Volij. Ranking participants in generalized tournaments. *International Journal of Game Theory*, 33(2):255–270, 2005.
23. M. Tennenholtz. Reputation systems: An axiomatic approach. Proceedings of UAI-04, 2004.