A Propagation and Aggregation Algorithm for Inferring Opinions in Structural Graphs

Nardine Osman, Carles Sierra, and Angela Fabregues

Artificial Intelligence Research Institute (IIIA-CSIC), Barcelona, Spain

Abstract. This paper is concerned with the issue of reputation and the computation of group opinion. We argue that entities may receive both objective and subjective opinions, and distinguishing between the two is crucial for achieving more precise measures. Additionally, we argue that the group opinion about an entity α is not only influenced by the opinions that α receives (whether objective or subjective), but by the reputation of other entities that α is related to. As such, we propose a method that permits the propagation and aggregation of opinions in structural graphs, allowing the inference of more precise reputation measures through the description of both objective and subjective group opinion.

Keywords: Trust, reliability, reputation, opinion propagation & aggregation

1 Introduction

There is a common understanding that reputation represents group opinion. Existing work has mainly focused on the direct opinions an entity (whether a person, a peer, an agent, or even an item) receives. This paper, however, introduces the concept of inferred opinions, based on relations between entities. Consider the example of having a poorly reputable football team that just started hiring well known players. Naturally, one would say that such a move is reasonable since it increases the team's chances, or *expectations*, of winning games, and hence, increase the team's reputation. These indirect opinions are highlighted when the entities are related. In other words, the reputation of different entities may influence each other when these entities are related. For example, having one entity being a part of another implies a *propagation* of opinions between these two entities. Hence, to consider indirect influential opinions, one should have a clear definition of the relations that link entities together, since opinions may only propagate along such relations. For the time being, we focus on the simple 'part of' relation. This results in the construction of a structural graph, which we define in the following section.

In addition to the introduction of the notion of opinion propagation in structural graphs, we also introduce the distinction between objective and subjective opinions. For example, the team being highly rated by some magazine should not be as influential in comparison with the team's actual performance, such as losing major games. We categorise opinions accordingly:

– **Objective opinions.** These opinions may not be falsified.[1] For example, if one ping pong player wins against another, then this may be interpreted as the former player stating that the latter is weaker, and vice versa.
– **Subjective opinions.** These are divided into two further subclasses.
  • *Direct subjective opinions.* For example, a scientific paper may win an award, or a scientific paper may receive good reviews from experts in the field. These may be viewed as direct subjective opinions.
  • *Influential opinions of related entities.* We say opinions may propagate from a parent node in the structural graph to its children nodes, and vice versa. For example, the final opinion about a scientific paper may be influenced by the final opinion about the conference where it has been accepted. Similarly, the opinion about a conference may also be influenced by the opinions about the papers it accepted.

The rest of this paper is divided as follows. Section 2 provides the basic definitions that the proposed algorithm is based on. The main model is then introduced by Section 3, which illustrates how group opinion may be calculated and how reputation is defined accordingly. The entire algorithm is summarised by Section 4. Some preliminary results are presented by Section 5, before Section 6 closes with a brief conclusion.

2 Basic Definitions

In what follows, we provide a clear definition of the structural graph that is needed for the propagation of opinions (Section 2.1), what direct individual opinions are (Section 2.2), and what group opinion is (Section 2.3).

2.1 The Structural Graph

We define a graph whose nodes represent entities that may form or receive opinions as follows:

Definition 1.

$$SG = \langle N, G, O, E, A, T, \mathcal{P}, \mathcal{F} \rangle$$

where

– *N is the set of nodes, or entities,*
– *G is the set of agents, peers, people, or even entities that may form opinions about $\alpha \in N$ (we note that G and N may or may not intersect, depending on the different fields of application),*
– *O is the set of direct opinions, whose elements are defined shortly,*
– *$E = \{e_1, ..., e_n\}$ is the evaluation space for O, where terms e_i account for terms like 'bad', 'good', 'very.good', etc.,*

[1] Objective information is not usually referred to as 'opinions'. However, to compute reputation, which is generally defined as group opinion, we propose to interpret all sources of information that could influence reputation as opinions.

– *A is the set of attributes (e.g. $\{strength, quality, \dots\}$) that opinions address,*
– *T represents calendar time,*
– *$\mathcal{P} \subseteq N \times N$ specifies which nodes are part of the structure of which others,*
– *$\mathcal{F} : R \times N \times A \times T \to O$ is a relation that links a given agent, node, attribute, and time to an opinion.*

2.2 Individual Opinions

We define an opinion $o^t_\alpha(\beta) \in O$ as follows:

Definition 2.

$$o^t_\alpha(\beta) = \{e_1 \mapsto v_1, ..., e_n \mapsto v_n\}$$

where,

– *$t \in T$, $\alpha \in G$, and $\beta \in N$,*
– *$\{e_1, ..., e_n\} = E$,*
– *$v_i \in [0,1]$ represents the value assigned to each element $e_i \in E$, with the condition that $\sum_{v_i} v_i = 1$.*

In other words, $o^t_\alpha(\beta)$ represents the opinion that an entity α may hold about entity β at time t; and the opinion is specified as a discrete probability distribution over the evaluation space E.[2] We note that the opinion one holds with respect to another may change with time, hence various instances of $o^t_\alpha(\beta)$ may exist for the same α and β but with distinct ts.

2.3 Group Opinion

This paper distinguishes between two different types of group opinion, based on the categorisation of opinions presented by Section 1:

– Group opinion based on objective opinions: $\mathbb{D}^t_\alpha(e_i)$
– Group opinion based on subjective opinions: $\mathbb{P}^t_\alpha(e_i)$

[2] We note that this paper only considers a discrete evaluation space E. In our proposed algorithm, as the equations of this paper illustrate, the main operations that are carried out over the probability distributions are the application of the addition $+$, subtraction $-$, multiplication $\times$, and division $/$ operators. These operators could easily be applied to continuous distributions as well. However, Equation 2 requires entropy measures. While some entropy measures (such as the minimum relative entropy) are usually very hard to calculate, calculating the entropy of a distribution is feasible for both discrete and continuous cases. Hence, we believe continuous probability distributions may be used, if necessary. Furthermore, discrete probability distributions already provide more information than that provided by the ranges of opinion values of existing methods.

The first is an aggregation of objective opinions only, while the second is an aggregation of both objective and subjective opinions resulting in a final subjective measure. Hence, like individual opinions, group opinion is defined as a probability distribution that represents the probability that entity α is e_i at time t (or has the reputation of being e_i at time t).

But why do we distinguish between these two different group opinions? Consider reputation in the field of football. Teams may play against each other. The results of these games may be viewed as having direct opinions being formed by one team about the other. For example, Barcelona winning Real Madrid 6–2 may be interpreted as Barcelona forming an opinion about Real Madrid being very weak and Real Madrid forming an opinion about Barcelona being very strong. This is one example of direct opinions in the field of football. Now consider that Real Madrid is either ranked high by some magazine or starts recruiting highly reputable players (at least higher than what they already have). In such a case, we might agree that this should increase the overall expectation of the team's performance, which could be viewed as increasing group opinion. Nevertheless, we believe such opinions are subjective, and their reliability cannot be matched to that of the objective group opinion (such as having Barcelona winning every single game for the last couple of years). Hence, we find differentiating between objective and subjective group opinion to be crucial.

Furthermore, we say subjective group opinion should lose its value with time, and move towards the objective one. For example, if Real Madrid kept on recruiting highly reputable players but failed to actually win their games, then the final reputation measure should always move towards the objective group opinion, i.e. the results of their games. Hence, we say, although subjective measures are important to describe the current group opinion, a purely objective measure is also needed. With time, and with the lack of new information, subjective group opinions should move towards objective ones. This is the notion of decay: everything loses its value with time. Similarly, objective group opinion would also decay towards the flat probability distribution (the distribution describing the state of complete ignorance), although at a presumably much slower rate.

The following section illustrates how these different measures may be calculated, highlights the links between them, and elaborates on the notion of decay.

3 The Proposed Model

This section focuses on the computation of group opinions, both objective ($\mathbb{D}$) and subjective ($\mathbb{P}$), in Sections 3.1 and 3.2, respectively. Reputation is then defined by Section 3.3.

3.1 The Default Opinion $\mathbb{D}$

We say the default opinion of an entity α is the group's opinion about α that is based on objective opinions only. The group's opinion is calculated by considering all the objective opinions expressed in the past, taking into account the certainty of each of these opinions.

Assessing an objective opinion. Assume that β at time t gives the following opinion about α: $o^t_\beta(\alpha) = \{e_1/v_1, \ldots, e_n/v_n\}$. We need to consider how much *value* this opinion has, based on how reliable is β in giving opinions about α.

In this model, we will consider that the overall reliability of any opinion is the reputation value of the entity expressing the opinion, which changes along time. This reputation value $\mathcal{R}$ is defined later on by Section 3.3. However, in this section, we use this reputation value (which we view as an indication of the reliability of the opinion $o^t_\beta(\alpha)$) to modify the opinion value. The basic idea is that the more reliable an opinion is, the closer the final value is to the original opinion, and the less reliable an opinion is, the closer the final value is to the flat (uniform) distribution $\mathbb{F}$ (where $\mathbb{F} = \frac{1}{|E|}$). Thus, we define the distribution representing β's final view about α at time t as follows:

$$\mathbb{O}^t_\beta(\alpha) = \mathcal{R}^t_\beta \times o^t_\beta(\alpha) + (1 - \mathcal{R}^t_\beta) \times \mathbb{F} \tag{1}$$

The certainty of an opinion. The group's opinion is based on an aggregation of individual opinions. However, the certainty of each of these individual opinions is crucial. We say, the more uncertain an opinion is then the smaller its effect on the final group opinion is. The maximum uncertainty is defined in terms of the flat distribution $\mathbb{F}$. Hence, we define this certainty measure as follows:

$$\mathcal{I}(\mathbb{O}^t_\beta(\alpha)) = \mathcal{H}(\mathbb{O}^t_\beta(\alpha)) - \mathcal{H}(\mathbb{F}) \tag{2}$$

where $\mathcal{H}(\mathbb{X})$ represents the entropy of a probability distribution $\mathbb{X}$. In other words, the certainty of an opinion is essentially the difference in entropy between the opinion and the flat distribution.

Calculating $\mathbb{D}$. Again, we note that an entity can give opinions on another one at different moments in time. So let us define by $T_\beta(\alpha) \subseteq T$ the set of time points in which β has given opinions about α. The default group opinion $\mathbb{D}^t_\alpha$ about α at time t is then calculated as follows:

$$\mathbb{D}^t_\alpha = \frac{\displaystyle\sum_{\beta \in G} \sum_{t' \in T_\beta(\alpha)} \mathbb{O}^{t' \to t}_\beta(\alpha) \cdot \mathcal{I}(\mathbb{O}^{t' \to t}_\beta(\alpha))}{\displaystyle\sum_{\beta \in G} \sum_{t' \in T_\beta(\alpha)} \mathcal{I}(\mathbb{O}^{t \to t'}_\beta(\alpha))} \tag{3}$$

where, $\mathbb{O}^{t' \to t}_\beta$ represents the decayed value of $\mathbb{O}^{t'}_\beta$, and is discussed shortly.

This equation essentially states that the default group opinion is an aggregation of all $\mathbb{O}^{t' \to t}_\beta(\alpha)$ that represent the view of every entity β that has formed an opinion of entity α at time t. However, different views are given different weights, depending on the certainty $\mathcal{I}(\mathbb{O}^{t' \to t}_\beta(\alpha))$ of these views.

Initialising $\mathbb{D}$. When an entity α is first introduced or created at time t, there is no information what so ever about this entity yet. Hence, its initial probability distribution is the flat distribution $\mathbb{F}$ that accounts for the maximum ignorance (i.e. the maximum entropy): $\mathbb{D}^t_\alpha(e_i) = \mathbb{F}(e_i) = \frac{1}{|E|}$. Along time, and as objective direct opinions are formed, this probability gets updated following Equation 3.

Decaying $\mathbb{D}$ (and $\mathbb{O}$). Like any other type of information, the default group opinion is expected to lose its value with time. For example, assume that a given player has played a lot of games and gained a high default opinion; however, for a very long time, this player has never played again. What can one say about the player's default opinion at the present time? Naturally, its glorious history does not necessarily mean that the player still has those old skills. Hence, we say that with time, $\mathbb{D}$ loses its value (very) slowly by decaying towards the flat probability distribution $\mathbb{F}$ according to the following equation:

$$\mathbb{D}^{t' \to t}_\alpha = \Lambda(\mathbb{F}, \mathbb{D}^{t'}_\alpha) \tag{4}$$

where Λ is the *decay function* satisfying the property that $\lim_{t\to\infty} \mathbb{D}^t_\alpha = \mathbb{F}$. In other words, Λ is a function that makes $\mathbb{D}^t_\alpha$ converge to $\mathbb{F}$ with time. One possible definition for Λ could be: $\mathbb{D}^{t' \to t} = (\mathbb{D}^t - \mathbb{F})\nu^{\Delta_t} + \mathbb{F}$, where $\nu \in [0, 1]$ is the decay rate, and $\Delta_t = 1 + (t - t')/\kappa$, where κ determines the pace of decay.

Single opinions are pieces of information and as such they also decay along time. $\mathbb{O}^{t' \to t}_\beta$, which represents the decayed value of opinion $\mathbb{O}^{t'}_\beta$ at time t, is then similarly defined:

$$\mathbb{O}^{t' \to t}_\beta = \Lambda(\mathbb{F}, \mathbb{O}^{t'}_\beta) \tag{5}$$

3.2 The Inferred Opinion $\mathbb{P}$

While the default opinion $\mathbb{D}^t_\alpha$ represents the *objective* direct opinions of group members, the inferred opinion $\mathbb{P}^t_\alpha$ represents the final subjective opinion which is influenced by: objective direct opinions and subjective (both direct and propagated) opinions.

Calculating $\mathbb{P}$. How $\mathbb{P}$ is calculated differs with the different types of opinions triggering this calculation. The different cases are presented below.

1. **Subjective opinions.** If an entity is influenced by subjective opinions (whether direct or not), then its $\mathbb{P}^t_\alpha$ value is calculated accordingly:

$$\mathbb{P}^t_\alpha = \zeta \, \mathbb{P}^{t' \to t}_\alpha + (1 - \zeta) \, \mathcal{X} \tag{6}$$

   where ζ is generally based on the reliability of α and $\mathcal{X}$ describes the new subjective opinion. This equation implies that when α is highly reputable, the effect of $\mathcal{X}$ is minimal, and vice versa. The exact values of ζ and $\mathcal{X}$ are dependent on the type of the subjective opinion, which we outline below:
   (a) *Direct subjective opinions.* In this case, we say $\zeta = (\mathcal{R}^t_\alpha)^{\mathcal{R}^t_\beta}$ and $\mathcal{X} = o^t_\beta(\alpha)$. In other words, if an entity β forms an opinion about an entity α, then $\mathcal{X}$ takes the value of β's new opinion $o^t_\beta(\alpha)$. ζ would mainly be based

on the reliability of α, but is also influenced by the reliability of β since different entities should have different strength in affecting α. We note that $\mathcal{R} \in [0,1]$, as illustrated by Section 3.3.
In some cases, however, β may be a foreign entity to the structural graph. Examples of this case are when a paper wins a award, or a magazine ranks football players. We assume that it is hard to know the reputation of foreign sources and their effect on α. In such cases, the default value is $\mathcal{R}^t_\beta = 1$. Alternatively, the user may be free to assign a different reliability measure to $\mathcal{R}^t_\beta \in [0,1]$.

   (b) *Influential opinions of related entities.* In this case, we say $\zeta = (\mathcal{R}^t_\alpha)^{f(d_\alpha)}$ and $\mathcal{X} = \mathbb{P}^t_\beta$, where $f(d_\alpha) = (\mathcal{R}_\beta + d_\alpha - 1)/d_\alpha$. In other words, if a neighbouring node β (whether it was a parent or a child node) had its $\mathbb{P}^t_\beta$ value modified, then this should affect α's $\mathbb{P}^t_\alpha$ value. Again, the more reliable α is, then the smaller the effect of β should be. Nevertheless, the effect of β on α should also be influenced by the number of neighbouring nodes that α has (defined as d_α, or the degree of α). The larger this number, the smaller the effect of one neighbouring node is, and vice versa. We note that in this case, $d_\alpha \in [1, \infty]$. And the function $f(d_\alpha) = \mathcal{R}_\beta$ when $d_\alpha = 1$, and $\lim_{d_\alpha \to \infty} f(d_\alpha) = 1$.

2. **Objective opinions.** Objective opinions should have a stronger effect than subjective ones. In comparison with Equation 6, $\mathbb{P}^t_\alpha$ should now be calculated by giving more weight to the new objective opinion, as illustrated below:

$$\mathbb{P}^t_\alpha = \frac{\mathcal{R}^t_\alpha \ \mathbb{P}^{t' \to t}_\alpha + \mathcal{R}^t_\beta \ o^t_\beta(\alpha)}{\mathcal{R}^t_\alpha + \mathcal{R}^t_\beta} \tag{7}$$

   Note that unlike Equation 6, even if α was fully reliable ($\mathcal{R}^t_\alpha = 1$), the new objective opinion of β is still accounted for by taking into consideration the reliability of β with respect to that of α (and vice versa).

Initialising $\mathbb{P}$. Similar to the default group opinion $\mathbb{D}$, we say $\mathbb{P}^t_\alpha(e_i) = \mathbb{F}(e_i)$, where t is the time α is first introduced. Along time, this probability is updated according to the section above, either as opinions about α are formed by others, or as neighbouring entities have their $\mathbb{P}$s updated, influencing that of α.

Decaying $\mathbb{P}$. The value of $\mathbb{P}$ is a subjective value, as it is influenced by subjective opinions. For example, the reputation of a team changes as it changes its team members, since opinions about new team members influence the opinion about the team. However, such information is subjective, and what really matters at the end is whether the team is actually capable of winning with this new group of team members or not. For this reason, we believe that with time, the subjective opinion $\mathbb{P}$ should decay at a reasonable rate towards a more stable and objective opinion: the default opinion $\mathbb{D}$. This is expressed by the following equation:

$$\mathbb{P}^{t' \to t}_\alpha = \Lambda(\mathbb{D}^t_\alpha, \mathbb{P}^{t'}_\alpha) \tag{8}$$

where Λ is the *decay function* that has been introduced earlier by Equation 4.

3.3 Reputation and Reliability

As illustrated earlier, an essential point in evaluating the opinion of a given entity is how reliable ($\mathcal{R}^t_\beta$) it is. The idea behind the notion of reliability is very simple: an entity that is considered very good in a certain field is usually considered to be very good as well in assessing how others are in that field. This is based on the *ex cathedra* argument. An example of a current practice following the application of this argument is the selection of members of committees, advisory boards, etc.

But how is reputation calculated? Given an evaluation space E, it is easy to see what could be the 'best' opinion about someone: the 'ideal' distribution, or the 'target', which is defined as $\mathbb{T} = \{e_n \mapsto 1\}$. Given a 'target' distribution $\mathbb{T}$, the reputation of an entity β may then be defined as the distance between the current default opinion $\mathbb{D}^t_\beta$ and the ideal distribution $\mathbb{T}$, as follows:

$$\mathcal{R}^t_\beta = 1 - \textsc{emd}(\mathbb{D}^t_\beta, \mathbb{T}) \tag{9}$$

where EMD is the earth movers distance that calculates the distance (whose range is $[1, 0]$) between two probability distributions [1]. [3][4] As time passes and opinions are formed, the reputation measure evolves along with the default opinion. We note that at any moment in time, the measure $\mathcal{R}^t_\beta$ can be used to rank the different entities.

4 The Algorithm

The proposed model of Section 3 illustrates how opinions may be inferred through the propagation (Equation 6) and aggregation (Equations 3, 6, and 7) of individual opinions in structural graphs. Algorithm 1 summarises this model.

We note that this algorithm runs locally for a given node $\alpha \in N$. The algorithm is invoked every time α receives a direct opinion $o^t_\beta(\alpha)$, or its neighbouring node β updates its $\mathbb{P}$ value. We assume α saves all its computed $\mathbb{O}$ values (the value of the direct opinions it has received, following Equation 1) as well as its latest $\mathbb{P}$ and $\mathbb{D}$ values. The algorithm then proceeds by following the equations of the previous section in a straight forward manner.

5 Results

As illustrated by Figure 1, real life applications fall into different categories, based on whether they make use of objective opinions, subjective opinions, or both; or whether they make use of structural graphs or not. For example, Chess or Ping Pong are games with individual players, and the scores of the matches may be interpreted as objective opinions. The Diplomacy game is an example

[3] One important aspect to apply EMD is to determine what the distance between the terms in E is. That is the matrix $D = \{d_{ij}\}_{i,j \in [1,n]}$. The distance is certainly domain dependent, and can possibly be learned.

[4] Naturally, other distance measurements may also be used.

Algorithm 1 Updating node α's reputation $\mathcal{R}$ and inferred opinions $\mathbb{D}$ and $\mathbb{P}$

Require: N to represent the nodes of the structural graph
Require: $G = \{\alpha, \beta, \dots\}$ a group of agents that may form opinions about nodes
Require: $E = \{e_1, \dots, e_n\}$ an evaluation space
Require: $t \in T$ to represent calendar time
Require: $o^t_\beta(\alpha)$ to represent the direct opinion that $\beta \in G$ holds about $\alpha \in N$
Require: $T_\beta(\alpha) \subseteq T$ to represent the set of time points in which β has given opinions about α
Require: $\mathbb{X}^t$ to represent the value of the probability distribution $\mathbb{X}$ at time t
Require: $\mathbb{X}^{t' \to t}$ to represent the decayed probability distribution $\mathbb{X}^{t'}$ at time t, following Equations 4, 5 and 8
Require: $get\_opinion(o^t_\beta(\alpha))$ to represent α's receipt of the direct opinion $o^t_\beta(\alpha)$
Require: $get\_neighbour\_update(\mathbb{P}^t_\beta)$ to represent α's receipt of the neighbouring node β's updated $\mathbb{P}$ value
Require: $obj(o^t_\beta(\alpha))$ to represent that the direct opinion $o^t_\beta(\alpha)$ is an objective one
Require: $\mathcal{R}^t_\beta$ to represent β's known reputation at time t
Require: $\mathcal{I}(\mathbb{O})$ to represent the certainty of the opinion $\mathbb{O}$, following Equation 2
Require: d_α to represent the degree of the node α
Require: $f(d_\alpha, \mathcal{R}_\beta) = (\mathcal{R}_\beta + d_\alpha - 1)/d_\alpha$ which we simply refer to as $f(d_\alpha)$ when β is obvious
Require: $\textsc{emd} : 2^{\mathbb{P}(E)} \times 2^{\mathbb{P}(E)} \to [0, 1]$ which calculates the earth-mover distance between two probability distributions

$\mathbb{F}(e_i) = \frac{1}{n}, \quad \forall\, e_i \in E$
$\mathbb{T} = \{e_n \mapsto 1\}$
when $get\_opinion(o^t_\beta(\alpha))$ **do**
  if $obj(o^t_\beta(\alpha))$ **then**
    $\mathbb{O}^t_\beta(\alpha) = \mathcal{R}^t_\beta \times o^t_\beta(\alpha) + (1 - \mathcal{R}^t_\beta) \times \mathbb{F}$

$$\mathbb{D}^t_\alpha = \frac{\sum_{\beta \in G} \sum_{t' \in T_\beta(\alpha)} \mathbb{O}^{t' \to t}_\beta(\alpha) \cdot \mathcal{I}(\mathbb{O}^{t' \to t}_\beta(\alpha))}{\sum_{\beta \in G} \sum_{t' \in T_\beta(\alpha)} \mathcal{I}(\mathbb{O}^{t \to t'}_\beta(\alpha))}$$

    $\mathcal{R}^t_\alpha = 1 - \textsc{emd}(\mathbb{D}^t_\alpha, \mathbb{T})$
    $\mathbb{P}^t_\alpha = \frac{\mathcal{R}^t_\alpha \mathbb{P}^{t' \to t}_\alpha + \mathcal{R}^t_\beta o^t_\beta(\alpha)}{\mathcal{R}^t_\alpha + \mathcal{R}^t_\beta}$
  else
    if $\beta \in N$ **then**
      $\gamma = \mathcal{R}^t_\beta$
    else
      $\gamma = 1$
    end if
    $\mathcal{R}^t_\alpha = 1 - \textsc{emd}(\mathbb{D}^{t' \to t}_\alpha, \mathbb{T})$
    $\mathbb{P}^t_\alpha = (\mathcal{R}^t_\alpha)^\gamma \cdot \mathbb{P}^{t' \to t}_\alpha + (1 - (\mathcal{R}^t_\alpha)^\gamma) \cdot o^t_\beta(\alpha)$
  end if
end when
when $get\_neighbour\_update(\mathbb{P}^t_\beta)$ **do**
  $\mathcal{R}^t_\alpha = 1 - \textsc{emd}(\mathbb{D}^{t' \to t}_\alpha, \mathbb{T})$
  $\mathbb{P}^t_\alpha = (\mathcal{R}^t_\alpha)^{f(d_\alpha)} \cdot \mathbb{P}^{t' \to t}_\alpha + (1 - (\mathcal{R}^t_\alpha)^{f(d_\alpha)}) \cdot \mathbb{P}^t_\beta$
end when

of individual players whose reputation is based on the subjective opinions of other team members. In Football, however, one may view a team as being composed of players and sometimes one player may play in different teams (based on the league), giving rise to the notion of a structural graph. Additionally, opinions about team players may sometime be subjective, such as being ranked by some magazine. Scientific publications may be viewed as an example that uses structural graphs (conference proceedings are composed of papers, papers are composed of sections, etc.), and opinions on scientific publications by other researchers in the field are subjective.

	uses structural graphs	**doesn't use structural graphs**
uses objective opinions	Football	Chess Ping Pong
uses subjective opinions	Football Publications	Diplomacy

Fig. 1. Categorised applications

We choose the Chess example for experimentation, because there exists an official ranking and predicting algorithm for Chess (ELO [2]) that we can compare to ours. Hence, for a given dataset that specifies the real outcome of games, we run the ELO algorithm and our proposed one to compute the reputation of players and predict the outcome of future games accordingly. We then compare the predicted outcome of each of the algorithms to the real one. Initially, we ran several experiments over real Chess data. However, we noticed that the performance of both the ELO mechanism and ours was similar. For instance, in one experiment, our algorithm performed 2.3% better than ELO. Looking at the results, it seemed that players in the same tournament are more or less of the same experience, and hence, reputation. For this reason, the final results of games were a little bit random, and hence, the performance of both ELO and this paper's proposed algorithm was similar.

We then moved on to simulated data. We created two players, A and B, that played against each other over a number of years. A was initially a 'bad' player and it lost around 80% of its games during the years 1992-1998. However, after 1998, A stopped playing for a while, and it resumed playing in 2004. Its performance dramatically improved over the years 2004-2010. In general, our proposed algorithm performed better than ELO by 3.5%. We note that the results are still preliminary, as they simulate two players who play around 30 matches each. Figure 2 plots the distance between the real results and the predicted results of both ELO and our algorithm. The distance is measured using the earth mover's distance function, EMD; hence, the maximum distance possible is 1, and the minimum is 0. However, as illustrated by Figure 2, the main difference is highlighted in the year 2004, when the ELO algorithm performs very poorly compared to ours, since our algorithm's decay function allows a better prediction when behaviour changes with time.

As such, we conclude that our proposed algorithm is essentially useful in applications where the quality being assessed (behaviour of humans, performance of agents, quality of papers, etc.) could actually change with time.

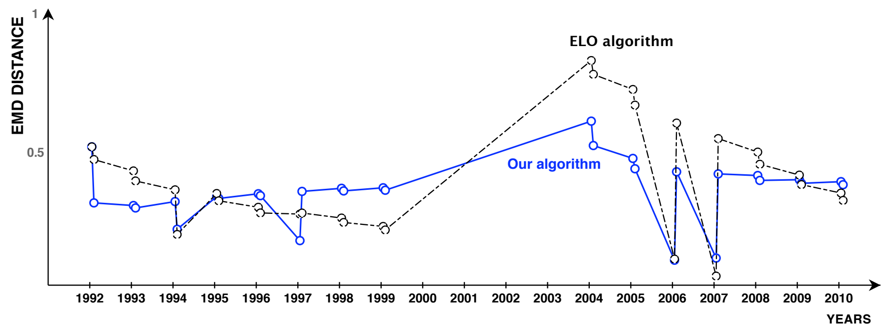

Fig. 2. The distance between predicted results and real results

6 Conclusion

This paper has proposed a model that allows agents to infer objective and subjective group opinion through the propagation and aggregation of opinions in structural graphs. It provides a clear distinction between objective and subjective opinions and their effect on group opinion. Additionally, the paper introduces the concept of opinion propagation between related entities.

In comparison with existing research, the research carried out by [3–5] studies the dynamics of opinion formation by focusing on the effect of social relations on how peoples' opinions may influence each other in a social network. The influence that one reviewer agent may have on another's subjective opinion is an interesting issue. Aggregation mechanisms, such as those presented by [6], may help in defining the appropriate aggregation method based on whether subjective opinions are dependent on each other or not. Repage [7], ReGreT [8], and SUNNY [9] provide mechanisms for computing the confidence in a reviewer based on the social relations. In this paper, we follow the *ex cathedra* argument which states that an agent's reputation could be used as an indication of its reliability in assessing others in its field. This fits perfectly in our equations that are concerned with objective opinions (Equations 1 and 7). However, again, when aggregating subjective opinions, social network analysis may be useful in contributing to the reliability of those opinions.

Concerning the propagation of opinions, we note that numerous research has addressed similar issues, such as [10–13]. PageRank [12] and Hits [13] calculate the relevance of web pages by analysing their position in the network and how they links to each other. Similarly, SARA [10] and CiteRank [11] present algorithms on how reputation may propagate based on who is citing whom. Their reputation propagates along citation links. This paper, on the other hand, focuses on the propagation of reputation along the structural links by focusing on the composition of entities and using the *part of* relation as an indication to the flow of opinions from one entity to another. Research work on ontology-based recommender systems, such as [14, 15], makes use of the clustering or classification of information and uses machine learning and data mining techniques for ranking and recommending entities. One may draw similarities between the taxonomies used by such systems and that of the structural graph of this document; although the propagation mechanism of this paper is unique in both its algorithm and semantics.

Acknowledgement

This work has been supported by both the LiquidPublications project (`project.LiquidPub.org`), funded by the Seventh Framework Programme for Research of the European Commission under FET-Open grant number 213360, and the Agreement Technologies project (`www.agreement-technologies.org`), funded by CONSOLIDER CSD 2007-0022, INGENIO 2010.

References

1. Peleg, S., Werman, M., Rom, H.: A unified approach to the change of resolution: Space and gray-level. IEEE Transactions on Pattern Analysis and Machine Intelligence **11**(7) (1989) 739–742
2. Elo, A.: The Rating Of Chess Players, Past & Present. 1st edn. Arco Pub. (1978)
3. Wu, F., Huberman, B.A.: Social structure and opinion formation. http://arxiv.org/abs/cond-mat/0407252 (July 2004)
4. Blondel, V.D., Guillaume, J.L., Hendrickx, J.M., Kerchove, C., Lambiotte, R.: Local leaders in random networks. Physical Review E **77**(3) (2008) 036114
5. Klimek, P., Lambiotte, R., Thurner, S.: Opinion formation in laggard societies. EPL (Europhysics Letters) **82**(2) (2008) 28008+
6. Sierra, C., Debenham, J.: Information-based reputation. In Paolucci, M., ed.: First International Conference on Reputation: Theory and Technology. (2009) 5–19
7. Sabater-Mir, J., Paolucci, M., Contexs, R.: Repage: REPutation and imAGE among limited autonomous partners. Journal of Artificial Societies and Social Simulation (JASSS'06) **9**(2) (2006)
8. Sabater, J., Sierra, C.: Social regret, a reputation model based on social relations. ACM, SIGecom Exchanges **3.1** (2002) 44–56
9. Kuter, U., Golbeck, J.: Sunny: A new algorithm for trust inference in social networks using probabilistic confidence models. In: Proceedings of the 22nd national conference on Artificial intelligence (AAAI'07), AAAI Press (2007) 1377–1382
10. Radicchi, F., Fortunato, S., Markines, B., Vespignani, A.: Diffusion of scientific credits and the ranking of scientists. http://arxiv.org/abs/0907.1050 (Sep 2009)
11. Walker, D., Xie, H., Yan, K.K., Maslov, S.: Ranking scientific publications using a simple model of network traffic. http://arxiv.org/abs/physics/0612122 (2006)
12. Page, L., Brin, S., Motwani, R., Winograd, T.: The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab (1999)
13. Kleinberg, J.M.: Authoritative sources in a hyperlinked environment. Journal of the ACM **46**(5) (1999) 604–632
14. Bouza, A., Reif, G., Bernstein, A., Gall, H.: Semtree: Ontology-based decision tree algorithm for recommender systems. In: Proceedings of the Seventh International Semantic Web Conference (ISWC2008). (October 2008)
15. Ziegler, C.N., Schmidt-Thieme, L., Lausen, G.: Exploiting semantic product descriptions for recommender systems. In: Proceedings of the 2nd ACM SIGIR Semantic Web and Information Retrieval Workshop, Sheffield, UK (July 2004)