

Chapter 1

A voting system with a fixed point

C. de Kerchove

*Department of Mathematics, Namur University, Namur, Belgium,
E-mail: cristobald.dekerchove@unamur.be*

P. Van Dooren

*Department of Mathematical Engineering, Université catholique de Louvain, Louvain-la-Neuve, Belgium,
E-mail: paul.vandooren@uclouvain.be*

A voting system is presented that is based on an iterative procedure that converges to a unique fixed point. The votes are casted by m raters regarding the reputation of n items, are organized as a $m \times n$ voting matrix X , which is possibly sparse when each rater does not evaluate all items. From this matrix X , a unique rating of the considered items is finally obtained via an iterative procedure which updates the reputations of the n items as well as that of the m raters. For any rating matrix the proposed method converges linearly to the unique vector of reputations. Some applications of this voting system will be presented.

1.1 Introduction

One of the most influential changes in our generation is undoubtedly the internet and its use for the communication of information. It is via the internet that most individuals are looking for information or are provided with information when signing up for some *information channel*. Typical tools for the finding of information are search engines, such as Google and Yahoo, whereas Facebook, CNN, various newspapers and several entertainment channels are examples of information channels. Many of these channels use databases of *opinions* gathered from a pool of arbitrary users and are based on votes that are ultimately used to rate the objects that users are

interested in. These objects could be books, hotels, restaurants, movies or touristic locations: one may e.g. refer to Amazon to find out about popular books, Booking or Trivago to inquire about hotels, MovieLens to look into movies, or Tripadvisor to investigate about various touristic hotspots. The list of such interactive sites is continuously growing. These actions can all be interpreted as a form of voting; but the honesty or reliability of the *raters* cannot always be verified.

A rater on the MovieLens database may give random ratings to movies he/she has not even seen, or a dishonest voter on Tripadvisor may post biased opinions just to favor his or her "friends".

Clearly these websites can only benefit from their rating system being as trustworthy as possible. At first sight, this looks like an impossible task since one cannot verify the honesty of all raters.

However, the coherence of the ratings provided by isolated raters can at least be checked against that of the *average* opinion. This is the approach that we propose here. We will actually try to achieve two simultaneous goals. The first such objective is to assign a *reputation* to each of the evaluated items, and the second one is to assign a grading of *reliability* or *trust* to each of the raters who evaluated the items.

We establish a clear difference between the *reputation* of an item, that is, between what is generally said or believed about the quality of characteristics of an item, and the *reliability* of a rater, which is our expectation that the rater gives a fair or relevant evaluation to the item in question.

We illustrate the need of such a voting system by recalling a voting scandal at the 2002- Winter Olympics. To this effect, we quote the following account from Wikipedia [8] (see also [1, 7]):

At the 2002 Winter Olympics held in Salt Lake City, allegations arose that the pairs' figure skating competition had been fixed. The controversy led to two pairs teams receiving gold medals: the original winners Elena Berezhnaya and Anton Sikharulidze of Russia and original silver-medalists Jamie Salé and David Pelletier of Canada. The scandal was one of the causes for the revamp of scoring in figure skating to the new ISU Judging System. [...]. The ISU Judging System replaced the previous 6.0 system in 2004. This new system was created in response to the 2002 Winter Olympics figure skating scandal, in an attempt to make the scoring system more objective and less vulnerable to abuse.

During this event, two of the judges favored the Russian team with

scores that were quite different from the averages scores of the remaining voters. The press and the public immediately, openly criticized the results of the voting, and the reaction was so strong that the president of the International Olympic Committee announced the decision that, though Russia would be allowed to keep the Gold Medal, the Canadian team, to whom the silver medal had been awarded as a result of the competition, would also get a gold medal.

The new ISU system is actually quite complex and tailored to the specific case of figure skating events. We present a different voting system, based on a fixed point iteration and that turns out to be the solution of an optimization problem. We delve into the advantages of this system and illustrate its use in several applications.

1.2 Reputation and trust

Various measures of reputation have been proposed in recent years under the names of reputation, voting, ranking or trust systems, among others and they deal with a number of contexts ranging from the classification of football teams to the reliability of each individual in peer to peer systems. Surprisingly enough, the most used method for measuring reputation on the Web, amounts simply to averaging the votes. In that case, the reputation is, for instance, the average of scores represented by stars in YouTube, or the percentage of positive transactions in eBay. Such a method, then, implicitly trusts evenly each rater of the system. Besides this method, many other algorithms exploit the structure of networks generated by the votes: raters and evaluated items are nodes connected by votes, as illustrated in Figure 1.1.

There are many different ways of defining trust or reputation and each of them has advantages and shortcomings. We refer here to [2] for a short survey on these ideas and the principles they are based on. Obviously the choice of a specific reputation system depends on subjective properties that are just accepted. For example, in the averaging method mentioned above, it is tacitly agreed that every rater is taken into account in the same manner, whereas the PageRank algorithm is based on the principle that a random walk over the network is a good model for the navigation of a web surfer. The fundamental assumption underlying the method we present here is the following:

Raters diverging often from other raters' opinions ought to be taken less




FIGURE 1.1: Network and matrix of votes from raters to objects and from raters to other raters

into account than the remaining raters.

This principle is the basis of our filtering process and implies that all votes are taken into account, but with a continuous validation scale, in contrast for instance, to the direct deletion of outliers. Moreover, the weight of each rater depends on the distance between his/her votes and the reputation of the objects he/she evaluates: typically weights of random raters and outliers decrease during the iterative filtering. The main criticism to be raised against this method is that it discriminates marginal evaluators.

Votes, raters and objects can appear, disappear or change, making the system dynamical. This is for example the case when we consider a stream of news like in [19]: news sources and articles are ranked according to their publications over time. Nowadays, most sites driven by raters involve dynamical opinions. For instance, the blogs, the site Digg and the site Flickr are good places to exchange and discuss ideas, remarks and votes about various topics ranging from political election to photos and videos. We will see that our proposed system allows for the consideration of evolving voting matrices and provides time varying reputations.

1.3 Weighted averages of votes

A natural way of tackling the problem of unreliable or unfair raters in reputation systems is to assign a weight to the evaluations of the raters.

Hence the range of weights corresponds to a continuous scale of validation of the votes. These weights change via an iterative refinement that is guaranteed to converge to a reputation score for every evaluated item and to a reliability score for every rater. At each step the reliability of a rater is calculated according to some distance between his/her given evaluations and the reputations of the items he/she evaluates. This distance is interpreted as the *belief-divergence*. Typically, a rater diverging too much from the group should be distrusted to some extent. The same definition of distance appears in [5, 6, 9] and is used for the same purpose. The strength of the reputation system we will describe here is that it can be applied to any static network of raters and items and that it converges to a unique fixed point. Moreover, our reputation system can also be extended to dynamical systems with time-dependent votes.

We describe our approach for a static system with m raters and n objects to be rated. The entry X_{ij} represents the vote of rater $j \in \{1, \dots, m\}$ for item $i \in \{1, \dots, n\}$, the matrix $X \in [a, b]^{n \times m}$ is the *voting matrix*. Each vote is in the positive real interval $[a, b]$, and the vector $\vec{x}_j$, the j -th column of X , is the vector of votes of rater j .

The graph of votes and raters can be represented by an adjacency matrix $A \in \{0, 1\}^{n \times m}$ where $A_{ij} = 1$ if object i is evaluated by rater j , and is equal to 0 otherwise. For the sake of simplicity, it is first assumed that every rater evaluates all items. Then the item's reputation vector $\vec{r}$ is defined as the weighted sum of the votes

$$\vec{r} := X \frac{\vec{w}}{\vec{1}_m^T \vec{w}}, \quad (1.1)$$

where $\vec{1}_m$ is the vector of all ones. Since this is a convex combination of the vectors $\{\vec{x}_1, \dots, \vec{x}_m\}$, it follows that $\vec{r} \in [a, b]^n$. The rater's weight vector $\vec{w}$ depends on the discrepancy with the other votes, interpreted as belief-divergence, which we define as

$$\vec{w} := G(\vec{r}) := \vec{1}_m - k\vec{d}, \quad \text{where} \quad \vec{d} := \frac{1}{n} \begin{bmatrix} \|\vec{x}_1 - \vec{r}\|_2^2 \\ \vdots \\ \|\vec{x}_m - \vec{r}\|_2^2 \end{bmatrix}, \quad (1.2)$$

and k is a positive parameter. Clearly, as k tends to zero, $\vec{w}$ tends to $\vec{1}_m$, and $\vec{r}$ tends to the average of the votes. Increasing k corresponds to more stringent discrimination toward outliers.

We proved in [2] that for $k < 1/b$ the vector $\vec{w}$ is always positive, and that this implies then that there exists a unique pair of vectors $\vec{r}, \vec{w}(\vec{r})$ satisfying both (nonlinear) equalities (1.1) and (1.2). Moreover the nonlinear iteration given by

$$\vec{w}^0 := \vec{1}_m, \quad \vec{r}^{t+1} := X\vec{w}^t / (\vec{1}_m^T \vec{w}^t), \quad \vec{w}^{t+1} := G(\vec{r}^{t+1}), \quad \text{for } t = 0, 1, 2, \dots$$




FIGURE 1.2: Representation of steps $\vec{r}^t$ of the nonlinear iteration in the unit box $[0, 1] \times [0, 1]$. The sequence $E(\vec{r}^t)$ decreases with t and converges to $E(\vec{r}^*)$.

converges to a unique fixed point $[\vec{w}^*, \vec{r}^*]$ that satisfies equations (1.1) and (1.2).

The uniqueness of the solution is established via the definition of the cost function $E(\vec{x}) = -\frac{1}{2k} \vec{w}(\vec{x})^T \vec{w}(\vec{x})$ that is minimized for $\vec{x}$ equal to the reputation vector $\vec{r}^*$. Moreover each step given by the nonlinear iteration resulting from formulas (1.1) and (1.2) corresponds to the steepest descent direction

$$\nabla_{\vec{r}} E(\vec{r}^t) = -\frac{1}{\alpha^t} (\vec{r}^{t+1} - \vec{r}^t),$$

of the cost function with step size $\alpha^t := \frac{n}{21_m^T \vec{w}^t} \geq \frac{n}{2m}$. It is shown in [2] that the corresponding steepest descent iteration

$$\vec{r}^{t+1} := \vec{r}^t - \alpha^t \nabla_{\vec{r}} E(\vec{r}^t)$$

converges to the fixed point $\vec{r}^*$ of the iteration and that $\vec{r}^*$ is the unique minimum of the cost function $E(\vec{x})$ for $\vec{x}$ in the hypercube $[a, b]^m$. This is illustrated in Figure 1.2 when only two objects are considered, that is, in the particular case when $m = 2$, and where the voting interval is $[0, 1]$.

Therefore the solution should not only be viewed as the fixed point of a nonlinear iteration, but it can also be interpreted as the minimizer of $-\vec{w}^T \vec{w}$ (and hence, as the maximizer of the 2-norm of $\vec{w}$). It therefore remains to be shown that there is a *unique minimizer* $\vec{r}^*$ in the imposed constrained set $\vec{r} \in \mathcal{H} := [a, b]^n$. This was again analyzed in [2], where it is shown that the function $E(\vec{r})$ is convex in $\mathcal{H}$, *provided* that the positive parameter

k is smaller than the value $\frac{1}{b}$. This indeed guarantees that the weighting vector $\vec{w}$ is strictly positive, from which it then also follows that the function $E(\vec{r})$ is strictly convex in $\mathcal{H}$. In Figure 1.3 we show (for a 2-dimensional vector $\vec{r}$ and the set $\mathcal{H} = [0, 1]^2$), a plot of four possible configurations of the function $E(\vec{r})$ depending on the value of k . In (a) we chose $k < 1$ which implies that the energy function is convex, in (b), (c) and (d), we gradually increase k which introduces saddle points and local minima and maxima, and eventually makes the function concave.




FIGURE 1.3: The function $E(\vec{r})$ for increasing values of k . In (a), $k < 1$ and there is a unique minimum. In (b), (c) and (d), k is increased and saddle points and local maxima and minima are observed.

1.4 Extensions

Here we briefly mention three different extensions that have been analyzed in [2].

Sparse votes

In reputation systems like Amazon, Tripadvisor and the Movielens

Database, the votes are clearly sparse since most raters do not give their opinion about all objects. If we assume that the entry X_{ij} of the voting matrix X is set to 0 when the entry A_{ij} of the adjacency matrix is 0, then $A \circ X = X$, where the symbol “ $\circ$ ” is used to denote the elementwise product of two matrices of the same dimensions (also called the *Hadamard* product). This property turns out to be crucial for the derivation of the sparse voting scheme explained in [2] and is based on a fixed point idea. The formulas (1.1) and (1.2) are now replaced with the following expressions :

$$\vec{r} = \frac{[X\vec{w}]}{[A\vec{w}]} \quad (1.3)$$

$$\vec{w} = G(\vec{r}) := \vec{1}_m - k\vec{d}, \quad \text{where} \quad \vec{d} = \begin{bmatrix} \frac{1}{n_1} \|\vec{x}_1 - \vec{a}_1 \circ \vec{r}\|_2^2 \\ \vdots \\ \frac{1}{n_m} \|\vec{x}_m - \vec{a}_m \circ \vec{r}\|_2^2 \end{bmatrix}, \quad (1.4)$$

where $\vec{a}_j$ is the j -th column of the adjacency matrix A , and n_j is the j -th element of the vector $\vec{n}$ containing the numbers of votes for each item, i.e., $\vec{n} = A^T \vec{1}_n$, whereas the scalar n denotes the total number of items. We point out that $\frac{[\cdot]}{[\cdot]}$ is the componentwise division of two vectors of the same dimension, which implies that every item is evaluated by at least one rater with nonzero weight. It is easy to verify that when the matrix A is the matrix of all ones, one retrieves the formulas for the dense voting matrix. Moreover, the nonlinear iteration

$$\vec{w}^0 := \vec{1}_m, \quad \vec{r}^{t+1} := \frac{[X\vec{w}^t]}{[A\vec{w}^t]}, \quad \vec{w}^{t+1} := G(\vec{r}^{t+1}), \quad \text{for } t = 0, 1, 2, \dots$$

converges to a unique fixed point $[\vec{w}^*, \vec{r}^*]$ that satisfies equations (1.3) and (1.4). We refer again to [2] for the proofs of these assertions.

Time-varying votes

This extension makes it possible to also consider dynamical votes where the rating matrix changes over time. Clearly, votes and web users are constantly evolving on the Web, therefore it appears necessary to develop also dynamical reputation systems. In this scenario, we consider discrete sequences of votes and adjacency matrices such as

$$\{X^s, s = 1, 2, 3, \dots\} \quad \text{and} \quad \{A^s, s = 1, 2, 3, \dots\},$$

that evolve on a discrete time axis, and we again assume that $A^s \circ X^s = X^s$ for every s . The iteration then becomes

$$\vec{w}^0 := \vec{1}_m, \quad \vec{r}^{t+1} := \frac{[X^{s+1}\vec{w}^t]}{[A^{s+1}\vec{w}^t]}, \quad \vec{w}^{t+1} := G_{s+1}(\vec{r}^{t+1}), \quad \text{for } t = 0, 1, 2, \dots,$$

where

$$G_{s+1}(\vec{r}^{s+1}) := \vec{1}_m - k \begin{bmatrix} \frac{1}{n_1^{s+1}} \|\vec{x}_1^{s+1} - \vec{a}_1^{s+1} \circ \vec{r}^{s+1}\|_2^2 \\ \vdots \\ \frac{1}{n_m^{s+1}} \|\vec{x}_m^{s+1} - \vec{a}_m^{s+1} \circ \vec{r}^{s+1}\|_2^2 \end{bmatrix}.$$

The convergence issue is clearly more delicate here but in the case of periodic votes it is e.g. shown in [2] that under mild conditions, the iteration converges also to a “fixed” periodic limit cycle.

Other discriminant functions

The scalar function

$$g(d) = 1 - kd$$

links the belief-divergence $\vec{d}$ to the weights $\vec{w}$ by an affine function. A similar idea was already present in [6], [9] and [5], but using different scalar functions, namely

$$g(d) = \frac{1}{d}, \quad g(d) = \frac{1}{\sqrt{d}}, \quad \text{and} \quad g(d) = e^{-kd}.$$

The motivation for using these more complex functions is that the corresponding minimization problem has a statistical interpretation, but the use of these functions also makes the problem of characterizing the fixed points harder. We refer to [2] for a further discussion on this issue.

1.5 A worked example

We illustrate the sparse extension of the method described in Section 1.3 with an experiment involving a data set (supplied by the GroupLens Research Project) of 100,000 ratings given by 943 users on 1682 movies. The votes ranged from 1 to 5 and the movies were selected such that each rater voted on at least 20 of the 1682 movies. This corresponds to a very sparse voting matrix but since every voter has a sufficient overlap with other raters, the computation of the divergence between raters remains sufficiently relevant. In order to test the robustness of our reputation system, we added 237 spammers giving always a vote of 1 except for their preferred movie, which they rated with a vote of 5.

Let $\vec{r}^*$ and $\tilde{\vec{r}}^*$ be respectively the reputation vector before and after the addition of these spammers. We expect that such behavior will be penalized by decreasing the spammer’s weights. Figure 1.4 illustrates the effect of

adding spammers for the two different methods we want to compare: firstly for our method using weighted averages, where the total perturbation can be measured by the distance $\|\vec{r}^* - \tilde{\vec{r}}^*\|_1 = 267$, and secondly by taking the unweighted average, where it is clearly seen that all reputations tend to be diminished. The distance is then given by $\|\vec{r}^* - \tilde{\vec{r}}^*\|_1 = 638$ and, as expected, this is greater than in the previous case, since spammers receive as much weight as the other raters.




FIGURE 1.4: X-Axis: the sorted movies according to their reputations before the addition of spammers. Y-Axis: their reputations according to our algorithm (Top) and to the average (Bottom).

Let us now look at the evolution of the weights during the iterations. The distribution of weights is shown in Figure 1.5 after one step, two steps and after convergence. Clearly the spammers receive eventually a much smaller weight than the original voters, and the method converges in very few steps.

This shows that our method could also be used to characterize outliers. Spammers could be detected by setting a threshold on the rater's weights w_i . In our example it follows from the last plot in Figure 1.5 that most raters with a weight below 0.6 are spammers and that most of those with a weight larger

than 0.6 are not. Dismissing completely the raters with a low converged weight can thus be viewed as a method to eliminate spammers. However, we prefer to take them into account, though with a reduced weight.




FIGURE 1.5: X-Axis: the weights of the raters. Y-Axis: the density after one iteration (Top), after two iterations (Left), and after convergence (Right). In dark grey: the spammers. In white: the original raters. In light grey: overlap of both raters.

1.6 Concluding remarks

Several other examples can be found in [2] and [3]. It is shown there that the voting system in competitions like the Eurosong contest suffers from so-called cultural voting (see [4]) and that this deficiency can be corrected by employing our filtering techniques. We point out that our technique could also be applied to websites such as Booking.com, Amazon.com, Tripadvisor and so on, where rankings are being offered based on anonymous votes. Bringing order to votes on the Web is certainly a promising topic that requires further investigations. Refining votes and hence reputations is one way to achieve that aim.

In conclusion, the main issues in reputation systems are, in our view, the relevance of the measure, the robustness against different sort of attackers, the application of the method to any sorts of data and the easiness for users to understand the measure.

It is surely tricky to determine how relevant a measure is in the context of voting. In our case we accept the idea of *belief-divergence* as a basis for cal-

culating the rater's weight, even if it implies the disqualification of marginal users. Nevertheless the parameter d allows us to quantify the degree of discrimination. Moreover the exponent in equation (1.2) can be chosen to be higher than 2, but then the uniqueness of the fixed point is no longer guaranteed. Several fixed points, however, could be interpreted as different opinion groups. In that way a marginal group would maintain its reputation if the number of its members is large enough. It makes sense to allow several opinions for a same movie, but providing one intrinsic value for each item should be more relevant in most contexts (such as the reputation of sellers and buyers on E-bay).

A durable reputation system must be robust. Smart cheaters who understand the system well enough to take advantage of it are certainly the most effective spammers. The way to proceed according to our approach is simple: they need to establish their reliability by correctly evaluating a group of items and then with that trust, they can rate some target items. In order to significantly change the reputation of these target items, they must have a number of coordinated evaluations that is larger than the one of honest raters. Such cheaters can thus be easily disqualified by looking at their coordinated ratings of one or several items. Unfortunately, this sort of spammer requires an extra process, similar to the procedure used by Google to detect spam farms that create thousands of links to boost their page ranks

The constantly increasing size of web data sets requires algorithms that are not too time consuming. Typically, a linear complexity in the number of votes is ideal. This is especially true for dynamical reputation systems, where the frequency of updates is high. In addition to this efficiency requirement, the method must be applicable to any "sparse" data. In general, the network resulting from votes between raters and objects is not complete, i.e., each object is not evaluated by all raters. These two points - complexity and sparse data - are a must if one wants a widespread use of a reputation system.

It is also desirable that reputation systems be able to cope with time-dependent data. More recent opinions may be considered more valuable than older ones, especially in the case of timely items such as news, arts, fashion, etc. The approach we described above can be easily extended to incorporate dynamical systems with time-dependent votes (see [2]) but the convergence to a fixed point is then replaced by the tracking of a time-varying trend.

Last but not least, the method must be understandable by those to whom it is directed, namely, by the users. Indeed, users cannot be confident in a voting if the measure of reputation looks like a black box. Although our method is more complicated than a simple E-bay-like system in which all ratings have the same weight, it remains relatively simple. In addition, users like the voting system to be transparent. For example, a record of the voting

history and comments from the raters helps users to develop their own opinions.

Note that all figures in this chapter are taken from Chapter 7 of the first author's Ph.D. thesis, Cristobald de Kerchove [3].