On the Epistemic Value of Reputation

The Place of Ratings and Reputational Tools in Knowledge Organization

Submission for the Eleventh International ISKO Conference 2010
Paradigms and conceptual systems in KO
February 23, 2010 – February 26, 2010

Organized by the Italian chapter of ISKO and the Faculty of Philosophy of the Sapienza University of Rome

Submitted by:

Gloria Origgi
Institut Jean Nicod
Ecole Normale Supérieure/ École des hautes études en sciences sociales
29, rue d'Ulm
F-75005 Paris
email: gloria.origgi@ehess.fr
www: http://www.institutnicod.org
tel: +33 (0) 1 443 22 6464
fax: +33 (0) 1 443 22 699

Judith Simon
Institut Jean Nicod
Ecole Normale Supérieure
29, rue d'Ulm
F-75005 Paris
email: judith.simon@ens.fr
www: http://www.institutnicod.org
tel: +33 (0) 1 443 22 6464
fax: +33 (0) 1 443 22 699

Gloria Origgi (Institut Jean Nicod/CNRS, Paris)
Judith Simon (Institut Jean Nicod, Paris)

On the Epistemic Value of Reputation
The place of ratings and reputational tools in knowledge organization

Abstract: In this paper we want to explore the epistemological relevance and value of reputation understood as evaluative social information. Using reputation to classify and assess an agent or an item can be epistemologically useful in the absence or - as is especially relevant today - overabundance of information. However, in order to be and remain epistemically useful and ethically just it has to be open to constant scrutiny and revision. We will introduce a model of rational consensus as an example for the rational use of reputation for epistemic purpose before analyzing different reputational tools on the Web. We will conclude our paper with a critical comment on the potential danger of using social information to evaluate information and knowledge claims, resp. to warn from epistemic injustices on the Web and elsewhere.

1: Introduction

What is that scarlet piece of tissue in the shape of an ***A*** sewn on Hester Prynne's gown in Nathaniel Hawthorne's masterpiece *The Scarlet Letter*? Is it a symbol of her sin, a "badge of shame", an indelible sign of her community's contempt? Is it a cruel reminder of her past, a succinct history of her misdeeds? Imagine that in the same colonial New England village, you do not have just a badge for the poor Hester, but each member of the community wears a letter that represents some past records of its owner. We can also imagine sets of identical badges worn by members of the community who have similar records: sinners, heroes, drunkards....Imagine that the elders of the community have the right to attach these labels to the villagers. Their judgments, based on their purported wisdom, become an easy way for the villagers to dispose of a basic classification of social types within the community that will allow them to manage their relations with others, to make inferences and predictions about their behavior, that is, construct a basic "social map" that will help them orient in their society. Morally this may be questionable, but epistemologically it can be useful.

We want to explore in this paper, the *epistemic value* of this type of social information, that is, reputation, while being aware of the *ethical and political problems* that might come with using it for epistemic purpose. Using the judgment on past records to classify an agent or an item can be epistemologically useful in the absence or - as is especially relevant today - overabundance of information. But it has to be and remain open to constant scrutiny and revision to be epistemically useful and ethically just.

2: Reputation as Evaluative Social Information

Reputation is a special kind of social information: it is social information about the value of people, systems and processes that release information. We want to explore here the relationship between this special form of social information - that implies an evaluative stance - and the processes of knowledge organization and evaluation. More precisely, we want to argue not only that (1) we make use of other people's reputations to evaluate information, but also (2) within systems, like the Web, that make possible the easy and dynamic organization and re-organization of knowledge, our own rankings may determine new content and generate new categories.

Reputation is the informational track of our past actions, it is the credibility that an

agent or an item earn through repeated interactions. We would like to defend an epistemological perspective according to which relying on reputational cues is an efficient way of shaping the too rich informational landscape around us by creating new relevant categories. Experts and authorities not only bloom where information is scanty, but also, and most crucially, in an information-dense world in which filtering out relevant information is our prominent cognitive activity. The epistemological enquiry we are advocating here implies that reputation and rating systems are an essential ingredient of collective processes of knowledge and play a *cognitive* role in extracting information. In an information-dense environment, where sources are in constant competition to get attention and the option of the direct verification of the information is often simply not available at reasonable costs, evaluation and rankings are epistemic tools and cognitive practices that provide an inevitable shortcut to information. We assume that there is no ideal knowledge that we can adjudicate without the access to previous evaluations and adjudications of others. No Robinson Crusoe's minds that investigate and manipulate the world in a perfect solitude. Our modest epistemological prediction is that the higher is the uncertainty on the content of information, the stronger is the weight of the opinions of others in order to establish the quality of this content.

Of course, this opens the epistemological question of the epistemic value of these rankings and reputation mechanism, that it, to what extent their production and use by a community changes the ratio between truths and falsities produced by that community and, individually, how an awareness of rankings should affect a person's beliefs. After all, rankings introduce a bias in judgment and the epistemic superiority of a biased judgment is in need of justification.

3: Rational Model for the Epistemic Use of Social Information

To illustrate how reputation understood as social information that comes with an evaluative stance can *rationally* be used for epistemic purpose, we introduce a formal model of rational consensus. In "Rational Consensus in Science and Society" Keith Lehrer and Carl Wagner develop their formal theory of consensus that rests upon the employment of consensual probabilities, utilities and weights and is meant to provide a model for rational decision making processes in science and society more generally ((Lehrer and Wagner 1981))). To our mind, this model is actually a model of how to quantify and use reputation for epistemic purpose.

Lehrer & Wagner argue that for decision making processes to be rational, it is central that *all* relevant information for the topic of concern has to be used ((Lehrer and Wagner 1981)). However, this spectrum of available information - for instance concerning disputes on scientific theories - is not limited to experimental information, but should also include the opinions of experts on other experts in the field. Lehrer calls this second type of information *social information* ((Lehrer 1990)) – and we call it reputation, i.e. social information that comes with an evaluative stance.

To illustrate how this social information might be used for epistemic purposes, Lehrer uses the so-called "expert dilemma" as a scenario. The expert dilemma describes the frequently encountered situation in which a decision has to be made despite the fact that evidence for answering a question is inconsistent and different experts recommend different options. An example would be whether or not to release a new medication or vaccine before all clinical trials are completed when facing the threat of an epidemic.

The basic question of Lehrer & Wagner ((Lehrer and Wagner 1981)) is the following: If scientific dissent is prevailing, but suspension of judgment is not an option, how should the conflicting information be used to reach a consensual conclusion? "Consent on the reputation of the experts in order to decide on the issue" could be the motto of their approach. Social information is used here as a crucial factor to decide on content information.

Using reputation as a decisive factor for factual matter rests upon the assumption that each expert in a certain community might be more or less reliable or competent with respect to the specific question at stake. If that is the case, it would most rational to include *each expert's answer weighted by his competence regarding the issue.* And the best way to assess the competence of each expert would be to use the aggregated reputation judgment of all other experts because they are most likely in the best position to judge the competence of their peers.

Lehrer & Wagner develop a quite complex mathematical model that describes an iterative and collective process to reach quantitative values for the reputation of each scientist ((Lehrer and Wagner 1981)). The basic idea however, is quite simple. The first step in this model consists in each expert giving a weight to all other experts summarizing all his information about the other's expertise and reliability concerning the issue at stake, in other words: he gives a quantitative indicator of what he considers to be the reputation of the scientist with respect to topic at hand. In a second step, the average reputation values for each scientist are calculated with a specific algorithm and then laid open. Then in the second round, each expert has to reassess the reputation value he has given to all other members of the community, i.e. she has the chance to revise his or her judgment taking into account the average weights which the other members of the community have given to their fellows. Similarly to Delphi-studies in the social sciences, this process is then ideally repeated until finally a consensual weight for each member of a community is achieved ((Linestone and Turoff 2002)).

The idea is that, if you are less secure about the reputation of a certain researcher, you might tend more towards the group average in your second vote. If you are very sure about the reputation of someone, however, you will not let yourself be influenced by this average. If everyone acts this way, that is considered to be most rational, then the consensus that is finally achieved is considered to be the most rational consensus. Crucially, once these consensual weights are achieved, they can be applied to answering the question of concern by weighting each member's vote on the issue with their consensual personal weight of reputation.

So, what should be obvious is that reputational cues, i.e. social information about other people that is evaluative, are *being used* – and that they *are useful*. Clearly, not all epistemic usage of reputation cues has to follow such a formal method. Quite on the contrary, ratings and other reputational tools might be used in a variety of different ways on the Web and our everyday life more generally. Nonetheless, Lehrer & Wagner's model delivers a clear example of the potential that reputation understood as social information from an evaluative stance, can have for epistemic tasks ((Lehrer and Wagner 1981)).

4: Reputational Tools on the Web

What the Web makes possible today is an algorithmic treatment of methods of gathering social information to extract knowledge. Ratings and rankings on the Web are

the result of collective human registered activities with artificial devices. However, the control of the heuristics and techniques that underlie this dynamics of information may be out of sight or incomprehensible for the users who find themselves in the very vulnerable position of relying on external sources of information through a dynamic, machine-based channel of communication whose heuristics and biases are not under their control. Thus, the reputational tools that are proliferating on the Web should be scrutinized by epistemically responsible users who do not want to accept too naïvely the outcome of a process they do not control.

The role of these *reputational tools* to filter information is getting more and more central in our Web-based epistemic practices ((Origgi 2009), (Origgi 2007)). And even more explicitly, we state that those systems that embody an access to others' judgments and rankings are rapidly outperforming, in terms of reliability, the random aggregation of multiple judgments and preferences on which many systems were based, as it is shown by the growing impact of the Web 2.0 on our epistemic practices. A growing number of examples of architectures on the Web show how these rankings work to produce new arrangements of information.

The Web 2.0 has provided the underlying networking structure to share ranked preferences. If you take the Web of the early years of 2000, one of the main feature that attracted much attention and criticism was the possibility to "customize" information for each user in order to fit each one's special needs and purposes of navigation. The endless potential of re-organization of the new, dynamic, information architectures based on the aggregation of chunks of contents according to specific rules (in contrast with the rigid tree-structures of the first-generation of web pages) opened the opportunity to create and organize "content on demand". News websites, online stores, search-engines, etc thus started to provide "My-" features to the users, that is, easily arranged customized pages with targeted news and other information for the users, personalized lists of products, personalized recommendations etc. This gave rise to a series of positive expectations and negative warnings, such as the risk of neglecting other people's points of views and perspectives by concentrating only on personally relevant information (cf. (Sunstein 2002)). Now, thanks to the social Web, these systems are evolving into systems of shared preferences, in which people can rely on someone else's preferences and ranking to construct their own categorization of information. Examples of this preferences-sharing are website such as Del.icio.us in which you can share your bookmarks with other people, or Flickr, in which, for each uploaded photo, not only you can see who uploaded it, but also who are the profiles that added it among their favorite pictures. Combining information about who comments on an image, who adds it as favorite, who tags it and how, Flickr now provides a new feature for browsing images: *interestingness*, http://www.flickr.com/explore/interesting/ which is an example of preference-based tools of categorization. As a Flickr user, I can decide to generate new categories of contents on the basis of an interestingness scale. A new category of the most interesting images on Flickr today is thus generated by sorting others' preferences. The success of this "fluid" way of constructing concepts and categories may depend also on the fact that it matches our cognitive capacities: it has been shown by cognitive psychologists (cf. (Barsalou 1995)) that concepts and mental categories are flexibly constructed in context.

In this perspective, the EC project LiquidPublication, (http://project.liquidpub.org/) in which both authors are involved, aims at developing "liquid" architectures for

producing, accessing and gathering scholarly information on the Web. Take for example the very concept of an academic journal: it is a selection of content based on a series of criteria of categorization: ISBN number, date of issue, etc. What we are working at in this project is a model of "Liquid Journal" which easily allows people to create selections of papers, articles, blog entries as a "My-journal" and then share them on the Web. One can imagine that, with the diffusion of such a model, the very category of "academic journal", or "journal issue" will be re-created by this particular form of information sharing, in which a user X can "conceptualize" a journal issue as for example: "all the content that the user Y is selecting in her journal". Here again, preferences of a user can be used by other people to re-organize information in a creative way. Virological examples of information diffusion based on a Twitter-logic of followers and leaders may be another example worth mentioning of reputational tools that create new categories of information.

Although the information-dense environment provided by the Web is the obvious locus in which examples bloom, we do not think that our analysis should be restricted to the case of the Web: in many other domains where information about the items at stake is very costly or difficult to obtain, reputational cues become an unavoidable way of organizing knowledge. Different cultural domains such as wine labeling systems and academic citation systems are based on rating devices that classify the underlying information by evaluating it (see (Origgi 2007), (Origgi 2009)).

5: Epistemic Injustices: On the Dangers of Using Social Information for Epistemic Purposes

The model of rational consensus as well as the Web applications that we have introduced are clearly examples of how reputation can be used to decide on content information, resp. on how social and content information might be productively merged to achieve better epistemic results. However, where there is use, there also is potential misuse. And in the case of reputational cues, these dangers might be inherent in the very concept of reputation as the “recognition by other people of some characteristic or ability” ((Merriam-Webster-Online-Dictionary 2009)).

More precisely there are two threats. First of all, the use of reputation to assess content can be epistemically beneficial while being morally questionable. This problem already becomes obvious in the first example we chose to open this article: Hawthorne’s A-shaped scarlet piece of tissue. Although classifying someone as a sinner, hero or drunkard – or as an expert, layperson or lobbyist - based on some cues might prove epistemically useful in certain situations, we would have to decide whether we are willing to pay the moral price of possible discrimination that comes with such stereotypical evaluation. More generally, once social information is taken into account to rate the quality of content, the door is open for social biases, prejudices and discrimination, which are as prevalent on the Web as in the societies that have developed and maintained it. These problems are not new and have long been identified for science and other epistemic fields by feminist epistemologists. In addition to raising awareness about these problems, various scholars have also developed tools and strategies to counter these epistemic injustices ((Fricker 2007), (Scheman 2001), (Alcoff 2001)). Miranda Fricker for instance distinguishes between testimonial and hermeneutic injustices as two instances in which someone is wronged in his capacity as a knower based on his social position. According to her “testimonial injustice occurs when prejudice causes a hearer to give a deflated level of credibility to a speaker’s word,

whereas hermeneutic injustice "[...] occurs at a prior stage, when a gap in collective interpretative resources puts someone at an unfair disadvantage when it comes to making sense of their social experience" ((Fricker 2007) 1). Clearly, both forms of injustice are easily conceivable when reputational cues and their epistemic usage are not critically reflected upon and kept open for constant scrutiny and revision.

The second problem concerns the limits of the epistemic usefulness of this type of information itself. The first question is how you calculate the reputation of someone else in the first place, resp. which proxies you use. Do you use the person's academic development, his institutional background, some form of communal evaluation, such as ratings or recommendations that he has received from other people as a cue to assess someone's reputation? Do you rely on your own experience with her? On some indicator of the quality of his former research? On her track record of different academic achievements? Her H-index or impact factor? Which of these proxies are valid and which are not? The second crucial questions concerns the stability of reputation, resp. the way you deal with evidence that supports or contradicts your view on the reputation of others. When, under which conditions and up to which point of counter-evidence or you warranted in keeping your reputation value for someone or something? Clearly, these issues as crucial as they are cannot be answered given the brevity of this paper. However, if we want to explore the utility of reputation for epistemic purposes, we have to analyze the potentials and possible dangers very carefully. That reputation is used to assess information and epistemic claims goes without saying – and it comes with benefits as much as with problems. So the question should be less how to avoid using reputation as epistemic tools, but rather how to use them wisely.

6: Conclusion

Our preliminary analyses indicate that ratings and reputational tools in knowledge organization have epistemological, cognitive, practical as well as ethical implications. From an epistemological point of view, a priority of rating tools and reputational scales over classification leads to a re-conceptualization of the "facts/values" dichotomy. Another epistemologically pressing question concerns the validity of reputation mechanism as epistemological tools. How epistemically warranted is the use of these tools? Is it just based on blind and imperfect heuristics that have a serendipitous effect on our search of information, or is it possible to conceive second order epistemic criteria that allow us to pry apart "good" and "bad" practices of trust and reliance on these reputational metrics?

For cognition, this implies to take into account a pragmatically oriented way of creating concepts and categories (i.e. the most "valued", items preferred by "x"), as it has already been argued in some works in cognitive psychology (cf. (Barsalou 1995)). From a practical point of view, this perspective may help to rethink the bottom up/top down distinction in designing categories by suggesting ways in which rating systems can serve as middle-ground categorizations that are neither imposed from above, not completely generated from spontaneous tagging. Rather they are user-driven meta-categorizations that inform the users.

The ethical and political aspects become obvious when taking feminist critique concerning the danger of epistemic injustices into account. Miranda Fricker's emphasis of the danger of testimonial and hermeneutic injustices are particularly pressing when reputational cues are used uncritically. It is especially when reputation mechanisms

become automatized in algorithms, there is a clear danger that epistemic injustices are inscribed in and reinforced by technology. Such an entanglement between ethics and epistemology in information design has been shown for trust-aware recommender systems((Simon 2008; Simon 2009)). Different trust metrics not only yield to different search results, but that they also correspond to different views concerning the organization as well as even more fundamentally the very concept of knowledge. Moreover, different trust metrics value different people differently and depending on the algorithm, some users are automatically silenced and "sorted out" ((Bowker and Star 1999)), while others "count". Thus, when developing reputational tools, the possibility of injustices has to be accounted for.

This example suggest that a purely epistemological or cognitive analysis of using reputation for epistemic purposes will not suffice for knowledge organization: the goals and standards for knowledge organization and epistemic practices have to be discussed and decided upon taking political and ethical considerations into account. Reputational tools open up new possibilities for knowledge organization, but they also bring with them their own problems. Raising awareness for the values as well as the dangers of using reputational cues for epistemic assessment will be the major goal of our talk.

References

Alcoff, L. M. (2001). On Judging Epistemic Credibility: Is Social Identity Relevant? Engendering Rationalities. N. Tuana and S. Morgen. Albany, SUNY Press**:** 53-80.

Barsalou, L. (1995). Flexibility, structure and linguistic vagary in concepts. Theories of Memory. A. F. Collins, S. E. Gathercole, M. A. Conway and P. E. Morris, Psychology Press, Taylor and Francis.

Bowker, G. C. and S. L. Star (1999). Sorting Things Out: Classification and Its Consequences. Cambridge, MIT Press.

Fricker, M. (2007). Epistemic Injustice. Power and the Ethics of Knowing. Oxford, Oxford University Press.

Lehrer, K. (1990). Metamind. Oxford, Claredon Press.

Lehrer, K. and C. Wagner (1981). Rational Consensus in Science and Society. Dordrecht, Reidel.

Linestone, H. A. and M. Turoff (2002). The Delphi Method: Techniques and Applications, Addison-Wesley Publishing Company.

Merriam-Webster-Online-Dictionary (2009). Reputation. Merriam-Webster Online Dictionary.

Origgi, G. (2007). Wine epistemology: The role of reputational and rating systems in the world of wine. Questions of Taste. B. Smith. Oxford, Oxford University Press**:** 183-197.

Origgi, G. (2009). Designing wisdom through the web. The passion of ranking. Collective Wisdom. J. Elster and H. Landermore. Cambridge, Cambridge University Press.
Scheman, N. (2001). Epistemology Resuscitated: Objectivity as Trustworthiness. Engendering Rationalities. N. Tuana and S. Morgen. Albany, SUNY Press**:** 23-52.

Simon, J. (2008). Knowledge and Trust in Epistemology and Social Software/ Knowledge Technologies. Culture and identity in knowledge organization: Proceedings of the Tenth International ISKO Conference. C. Arsenault and J. T. Tennis. Montréal, Canada, Würzburg: Ergon**:** 216-221.

Simon, J. (2009). MyChoice & Traffic Lights of Trustworthiness: Where Epistemology Meets Ethics in Developing Tools for Empowerment and Reflexivity. Proceedings of the 8th International Conference of Computer Ethics, Corfu, Nomiki Bibliothiki.