How the Web Is Changing the Way We Trust

Dario TARABORELLI [1]

University College London
d.taraborelli@ucl.ac.uk

Abstract. Several studies have addressed the issue of what makes information on the *World Wide Web* credible. Understanding how we select reliable sources of information and how we estimate their credibility has been drawing an increasing interest in the literature on the Web. In this paper I argue that the study of information search behavior can provide to social and cognitive scientists an extraordinary insight into the processes mediating knowledge acquisition by *epistemic deference*.

I review some of the major methodological proposals to study how users judge the reliability of a source of information in the *World Wide Web* and I propose an alternative framework inspired by the idea that–as cognitively evolved organisms– we adopt to this aim strategies that are as effortless as possible. I argue in particular that Web users engaging in information search are likely to develop simple heuristics to select in a computationally viable way trustworthy sources of information and I discuss the consequences of this hypothesis and related research directions.

Keywords. credibility, authority, trust, heuristics, information search

1. Judging epistemic reliability in the *World Wide Web*

Possessing reliable knowledge and being able to select reliable sources of information are skills essential to our capacity to cope in an efficient way with the problems raised by our physical and social environment. As evolved cognitive organisms, we negotiate demanding cognitive problems by selecting parsimonious strategies that provide us with sufficiently accurate solutions. *Epistemic deference*–or the ability to trust external sources of information to form new beliefs–can be regarded as one of such strategies. In this paper I suggest that epistemic deference is a common aspect of information search in the *World Wide Web* and I argue that in order to be cognitively efficient it has to rely on simple and relatively effortless heuristic strategies.

1.1. Epistemic deference and the problem of selecting reliable sources

Social epistemology has introduced the concept of "epistemic deference" to refer to those processes of belief formation in which a subject (the *deferrer*) relies on an external source

---

(the *deferee*) in order to extend her knowledge to facts with which she has no direct acquaintance, or more generally, to use information from this source as "a model for what to believe".[18] Relying on experts in order to make a decision is a typical example of a process in virtue of which we adopt a *deferential stance* towards an other individual's opinions (an *epistemic authority*) to extend our system of beliefs beyond its individual boundaries. Epistemic deference is a constitutive trait of language competence, as the capacity by which we can "entertain thoughts through the language that would not otherwise be accessible to us" [31] as when – for example – we use the term "arthritis" in a conversation without exactly knowing the precise reference of this term. But the scope of deference is arguably broader than language use. Deference to an external source of information, insofar as we trust that source for epistemic matters, allows us to extend our beliefs and our ability to reason about facts which we do not thoroughly master.[10] As such, deference is a principle found everywhere in human cognition and possibly one of the common strategies used to bootstrap knowledge and language acquisition in young children.[see for instance 9]

In the general case we defer to external sources of information: (i) whenever we lack reliable knowledge on a given subject matter to ground a decision (in which case deference is a *necessary* condition) or (ii) when deference provides a convenient, parsimonious *sufficient* solution to meet the requirements of a problem.

(i) If I am not a doctor, looking up for symptoms of a disease in a medical encyclopædia can be regarded as an example of the first situation: trusting medical information about a disease from an encyclopædia is a *necessary* form of epistemic deference since I am not a medical expert and I could not acquire knowledge about this disease if not by deferring to an expert source.

(ii) Now consider the case in which I cannot directly recall a friend's phone number and I decide to call another friend who may know this number by heart: in this case, I trust the person that I am calling (or rather, her memory) as a sufficiently reliable source to provide me the information that I need, even if I could directly check this information by myself by other means. For example I could go back home and find the number in the copy of the address book I keep near my landline phone.

The massive availability of information in the *World Wide Web* is making deferential practices as those exemplified by (i) and (ii) a constitutive part of our belief-formation and decision-making strategies. Searching the Web typically yields multiple sources for the piece of information we are looking for, so the critical question we face is which of these sources should be trusted.

How do we select *trustworthy deferees* when we engage in information search in the Web? Unsurprisingly, it has been shown that we systematically rely on background knowledge and previous experience as a main factor to decide whether a source of information in the *World Wide Web* is trustworthy or not. Familiarity and *experienced credibility* [14] are among the most common grounds for the selection of trustworthy sources to which we defer. In the general case, though, we have no prior information on the trustworthiness of an external source and we need to estimate it.

1.2. Evaluative judgments of epistemic reliability

The problem of credibility of electronic information as been the object of a growing body of literature in the last decades.[33] Studies of Web credibility [14, 15] or the perception of epistemic trustworthiness of unfamiliar sources in the Web have mostly focused on *evaluative judgments*, i.e. judgments people make in order to estimate the trustworthiness of a source of information on the basis of extensive inspection of the content and credentials provided by the source. Evaluative judgments should be distinguished by *predictive judgments* or judgments about the expected reliability of a source prior to its actual inspection. [16]

I will come back later on the *evaluative* vs. *predictive* distinction, but it is worth asking why mainstream research on Web credibility has been focusing on evaluative judgments. Arguably, the main reason why the study of evaluative judgments has been privileged in the literature is that *in ideal conditions*, whenever users are required to estimate the credibility of a source, they are not subject to particular constraints of time or cognitive effort to make this judgment.

Traditionally, the study of evaluative judgments of credibility of a source of information has addressed two central issues:

1. how easily information acquired by deferring to external sources can be integrated into one's system of beliefs;
2. how prone such information is to subsequent revision.

Mainstream theories of persuasion [11, 22, 24] suggest that among the factors affecting the likelihood of subsequent revision of an evaluative judgment, the *amount of processed information* and the *degree of involvement* plays a critical role. Judgments based on small amount of information or in conditions of low involvement are more likely to be subsequently revised [12, 25]. It would then seem natural to assume that in the case of knowledge acquisition mediated by information search, people are likely to invest a large amount of information processing effort with the goal of identifying trustable sources of information. I will try to show that this assumption cannot be taken for granted.

Web credibility studies have indeed collected large datasets on evaluative judgments of source reliability by asking users to rate the quality of visited websites. [13] The results allowed to single down several factors affecting:

1. the likelihood that a specific features of the source be noticed by the subject;
2. the attribution of positive or negative values to features that are noticed.

Verbal reports and qualitative, *questionnaire*-based methods have been the most popular approach to studying judgments of credibility and source reliability at least since two decades [38] and are still the dominant approach adopted to investigate credibility in the *World Wide Web*. It should be noted that this approach is not limited to judgments of credibility of Web sources decontextualized from the specific task in which such judgments are required: even when researchers have looked at judgments of epistemic reliability in real human-computer interaction tasks, they tended to privilege verbal reports (or "think-aloud" protocols) as the main source of empirical evidence over other possible kinds of behavioral data [32].

1.3. Beyond qualitative analyses of epistemic reliability judgments

The use of verbal reports to understand Web credibility relies on the assumption that introspection is the best way to determine factors affecting judgments and decisions on the epistemic reliability of a source. As, for example, [32] observes,

> [t]he method used in this study is premised on the assumption that the users can identify and discuss the characteristics and features of information objects that influence their judgments of information quality and cognitive authority. (p.150)

However useful verbal reports may prove to study evaluative judgments in decontextualized conditions, they face a number of major limitations:

1. verbal reports assume that subjects are aware of the factors affecting the selection of a specific source as credible or epistemically reliable, but there is no need to assume that processes involved in such judgments should be explicit;
2. reports relying on extensive inspection of a Web source can hardly account for the kind of processes in which users engage when they are involved in real-world information search tasks, which are usually constrained by time and by limits on the cognitive effort the user is willing or capable to invest in the task;
3. qualitative studies based on in-depth, decontextualized source evaluation implicitly take for granted that *a posteriori*, evaluative judgments are immune to task-dependence effects (which may affect a user's perceived utility of the different features and credentials of a source);
4. evaluating credibility judgments against a list of predefined dimensions on the basis of verbal reports typically skews the results in favor of an (often arbitrary) class of credibility variables chosen by the experimenter.

These methodological issues affecting qualitative studies of Web credibility call for an alternative framework to study user behavior in ecologically valid conditions and under the typical constraints of real information search tasks.

1.4. Predictive judgments of epistemic reliability

Information foraging studies [26–29] marked a turning point in the literature on information search, by drawing the attention of the research community on the importance of studying *predictive judgments* of the value of a source of information in the context of an information search task as opposed to *a posteriori* evaluations. Predictive judgments are those that users make when they evaluate a source on the basis of information describing a source, like in the case of link descriptions. This "proximal information" can allow the user to estimate which source to visit prior to its in-depth evaluation.

Why are predictive judgments underrepresented in current research on *Web* credibility? One of the possible reasons is that low effort judgments may have been ruled out as irrelevant to the understanding of trust and epistemic deference. Since beliefs formed in low effort condition (i.e. those processes that tend to be classified as "peripheral routes to persuasion" [22]) have been shown to be more volatile and less predictive of behavior, researchers may have erroneously assumed that these are not as representative of deferential behavior as beliefs acquired on the basis of in-depth, more cognitively demanding evaluations.

The question boils down to understanding what is the average level of engagement typical of information search behavior on the Web. As Fogg observes:

> Web users typically spend small amounts of time at any given site or individual page, and are thus likely to develop strategies for assessing credibility quickly. One could argue that people typically process Web information in superficial ways, that using peripheral cues is the rule of Web use, not the exception. From a user perspective, there are too many competitors on the Web for deep credibility evaluation. Even the words people use to describe Web use–"visiting sites" and "surfing the Web"–suggest lightweight engagement, not deep content processing. Research has yet to examine the relationship between engagement level and credibility assessments online. [15, p.15]

It is reasonable to assume that, depending on the goals of the information search task, there may be different possible degrees of deference to a source [18, p.189], and–consequently–different degrees of cognitive engagement required to evaluate it. Seeking a reliable source of medical information to ground a critical decision regarding someone's health need not require the same level of cognitive engagement as searching for a reliable bibliographic source with the goal of writing an essay. However, I submit that the problem is not merely a matter of understanding the degree of engagement required by the domain of the query. Studying epistemic deference in real-word conditions must take into account general constraints that apply to judgments of epistemic evaluation in the context of information search tasks. Fogg's quotation evokes a number of such constraints.

First, *time constraints*. In a world in which online content is becoming massively and constantly available, user interactions with the Web naturally tend to become shorter, more frequent and increasingly mediated by search engines. *Information snacking* [20] can be seen as the application to the Web of a known principle of situated cognition that states that organisms tend to externalize the solution of demanding cognitive problems to the environment and use the environment as an external scaffolding to decrease cognitive effort. [8]

The second major class of constraints comes from *epistemic pollution* [35]. The larger the volume of potentially relevant but weakly authoritative information, the more urgent is the need of efficient and cognitively viable skllls for source selection. In conditions in which the number of items to evaluate increases beyond control, *a posteriori* evaluative judgments simply become intractable and predictive judgments seem to be the only viable solution. It should be noted, incidentally, that information pollution thrives precisely because of these lightweight, heuristic strategies in which Web users systematically engage. It requires little effort to forge an attractive link luring the user into believing that it will lead to the target source. Fighting epistemic pollution is mainly a matter of detecting cheaters (e.g. sources of unreliable information) on the basis of proximal cues, and this may be an even more effort-consuming task than individually evaluating each source. As Nielsen observes, "information pollution is, [for hungry wolves], like packing the forest with cardboard rabbits" [19]. Good heuristics are those that allow to tell a fake rabbit from a real one before even starting to hunt it.

The sum of time constraints and constraints imposed by epistemic pollution is the main rationale in support of the hypothesis on the nature of epistemic reliability judgments that I defend in this paper. I submit that due to these constraints, making a judg-

ment of source reliability in the *World Wide Web* is more likely to be the result of selecting appropriate heuristics, i.e. sufficiently reliable predictive strategies based on link evaluation, than time consuming and cognitively demanding *a posteriori* evaluation processes.

2. Heuristics for epistemic reliability

I presented in the previous section the main rationale to argue that the study of judgments of epistemic reliability in the context of real information search tasks in the *World Wide Web* should focus on predictive, heuristic strategies rather than in-depth source evaluation processes. Heuristics for the evaluation of credibility of a source as a precondition to epistemic deference can be seen as a subset of a broader class of cognitive heuristics that people assumedly adopt in assessing credibility of electronic information.[36] In this section I will focus on some broad theoretical implications of this hypothesis.

2.1. Proximality

One of the main limitations of traditional studies of Web credibility is the fact that they largely neglected the role of predictive judgments of reliability based on proximal cues about sources of information. The *World Wide Web* is rich of cues that represent (in a more or less reliable way) sources of information. These cues have been referred to in the information foraging literature as the constituents of *information scent*, i.e. a measure of the perceived profitability of a distal source prior to its selection [29]. The hypothesis endorsed by information foraging studies is that information seekers base the choice of optimal navigation patterns on the perceived strength of *information scent* and on the maximization of scent over effort (e.g. time and length of navigation patterns).

If we accept this assumption, information *about* a source (e.g. how a source of information is represented in search engine results) becomes much more critical for its evaluation than the actual content it delivers. It is then plausible to assume that the problem facing Web users seeking reliable information is a matter of understanding whether proximal cues (as those afforded by search engines, for instance) are good predictors of target sources. This, I submit, is possible only under the condition that this ecology is stable and sufficiently constrained.

2.2. Structured environments

The *ecology of the Web* has been the object of extensive studies in the information science literature [6, 17, 30]. The existence of strong ecological regularities constrains the way in which users learn the structure of the *Web* and determines to a large extent their preferential strategies in information search behavior. It is plausible to assume that information seekers are *situated* in this environment and rely on ecological regularities they have learnt in order to select effective solutions for negotiating source selection problems.

Cognitive technologies such as *search engines* aim at improving our information retrieval skills by reducing the cognitive effort required to solve particularly demanding tasks and by increasing the amount of information scent available to the user. In this sense, they tend to favor the selection of simple, effortless and automatic strategies over more costly processes. By enriching the user's ecology with highly informative cues and

making this ecology stable, technology aims at reducing information processing requirements on the user.

As Ecological Rationality theories suggest [3, 37], stable environments offer ideal conditions to favor the selection of shallow, effortless and relatively rigid computational strategies. These are typical features of *modular solutions* to the problem of negotiating cognitively demanding problems.

2.3. Modularity

Defendants of the modularity hypothesis insist that modularity arises as a viable solution in stable environments whenever an organism faces a problem of computational tractability of information [4, 5, 34]. As Carruthers observes,

> computational processes need to be *local*–in the sense of having a restricted access to background knowledge in executing their algorithms–if they are to be tractable, avoiding a "computational explosion". And the only known way of realizing this, is to make such processes modular in nature–[4].

If epistemic deference has to be cognitively profitable, then solutions to the problem of estimating the reliability of a source must be computationally tractable. Source evaluation processes whose cost outweighs the benefits of deferring to a source are unlikely to be selected as viable. I will call this a *cognitive affordability constraint* on deferential strategies. The selection of deferees is a paradigmatic case of problem that has to be solved in a cognitively tractable way by setting limits to background knowledge, in order to avoid computational explosion. In the case of reliability judgments, this means finding sufficiently local criteria for estimating the reliability of a source, that do not draw in turn on further reliability judgments and so on.

If local inferential strategies can be identified that accurately yield a representation of the trustworthiness of a source, then we can say that the basic conditions are met for the selection of a modular solution to the problem of source evaluation.

3. Empirical research directions

I have reviewed some of the broad implications of the hypothesis according to which reliability judgments in the Web can in principles be underpinned by highly specialized heuristics based on cues that allow accurate predictions of the reliability of a source. The question that needs to be answered on empirical grounds is then whether – given the ecology of the *World Wide Web* – there are specific heuristics based on information scent that users can adopt to predict the trustworthiness of a source. In this section, I sketch a programme that future research should aim to implement in order to empirically test this hypothesis.

3.1. Non-reputational cues in source evaluation

The first empirical research direction consists in studying how people use *non-reputational* proximal information to decide which sources are worth being selected. I refer to "non-reputational cues" as the class of properties of the proximal representation of a source (e.g. a search engine result) that do not contain explicit information about the

credentials of the source. In the case of common search engines, such cues include properties of the title, snippet and URL of an item in a search result page. It is a promising avenue for experimental research to study if we implicitly use properties such as *URL length*, *processing fluency of the snippet* or *density of keywords matching the query* in order to predict whether a specific item in a search engine result page is trustworthy (and hence worth being selected). Experimental designs will have to control in particular for effects depending on task assignment, as it is likely to expect that users trying to maximize semantic relevance in source selection may not use the same cues as users trying to maximize perceived trustworthiness.

3.2. Explicit reputational cues in source evaluation

A second research direction should focus on the study of the impact of *explicit reputational cues* on judgments of epistemic reliability. The Web offers a plethora of indicators of source "popularity" or "endorsement" that can be regarded as explicit reputational cues. Social software and Web 2.0 services have already made these indicators a promising avenue for future generations of search engines. [1, 39] These cues can be broadly grouped in six different categories:

1. *implicit indicators of individual endorsement* (such as indicators that a specific user selected/visited/purchased an item);
2. *explicit indicators of individual endorsement* (such as explicit ratings produced by specific users);
3. *implicit indicators of socially aggregated endorsement* (such as density of bookmarks or comments per item in social bookmarking systems like *del.icio.us*, *Digg*, *Reddit* etc.);
4. *explicit indicators of socially aggregated endorsement* (such as average ratings extracted from a user community);
5. *algorithmic endorsement indicators* (such as PageRank and similar usage-independent ranking algorithms [2]);
6. *hybrid endorsement indicators* (such as *interestingness* indicators in *Flickr*, taking into account both explicit user endorsement and usage-independent metrics);

Whereas in the general case, subjects should have no reason to trust the validity of such reputational cues other than trusting the provider of these cues, it is reasonable to expect that these indicators strongly bias the processes through which we select reliable sources. Experimental research will have to understand in particular:

- to what extent judgments of epistemic reliability are affected by different types of explicit reputational cues;
- to what extent the overall trust of the subject in the system providing these cues modulates their judgments.[see for instance the seminal work by Keane and O'Brien, 21]
- to what extent explicit reputational cues override implicit, non-reputational cues.

3.3. From reliability heuristics to biases

Possibly the most interest question is to study how these heuristics may result in large-scale biases in deferential behavior, which can be exploited by manipulating the per-

ceived trustworthiness of a source. It has already been shown that the sheer ranking of items in search engine result pages (i.e. the fact that top results attract the vast majority of clicks) produces strong asymmetries in the number of sources that are selected and visited by the majority of users. [7] Similar large-scale asymmetries in the distribution of visits are likely to be found as a result of heuristics that users adopt to evaluate sources of information in a fast and effortless way, especially if the outcome of their selection is fed back to other users. By making the link between individual endorsement and reputational indicators more and more technologically mediated (and hence less transparent to the end user) the Web is already massively biasing the way in which we decide which sources are worth being trusted, selected and visited. Future research will have to clarify the ethical implications of the increasing impenetrability of reputational cues Web users rely on and understand if policies need to be introduced to control this phenomenon.

4. Conclusions

In this paper I fleshed out the main rationale, theoretical implications and some potential research directions in the study of processes underlying epistemic reliability judgment in the *World Wide Web*. I proposed that such processes should be regarded as a class of capabilities depending on highly specialized heuristics and that heuristics-based predictive judgments are likely to be more ecologically valid than the kind of evaluative judgments studied so far in the Web credibility literature.

I suggested in particular the conditions under which such heuristics are likely to emerge and stressed how by decreasing the overall cognitive effort involved in source evaluation, they probably are in a better position to describe what users do when engaging in real-world information search behavior.

The rationale for this research programme does not rule out the necessity of studying effortful, *a posteriori* evaluative strategies, but calls for a better understanding of the contexts in which these strategies are deployed. Heuristics to ascertain the credibility of sources of information are likely to be privileged only in those cases in which (1) cognitive engagement is low, (2) the ecology in which these strategies apply is sufficiently stable to allow learning, and (3) simple cues are sufficiently accurate to allow the user to cope with epistemic pollution.

I proposed that such conditions may be more common than the current literature on Web credibility has realized so far, as they may provide a more realistic account of how we select sources to engage in deferential behavior.

References

[1] S. Bao, G. Xue, X. Wu, Y. Yu, B. Fei, and Z. Su. Optimizing web search using social annotations. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, pages 501–510, New York, NY, USA, 2007. ACM Press. ISBN 9781595936547. doi: 10.1145/1242572.1242640. URL http://dx.doi.org/10.1145/1242572.1242640.

[2] S. Brin and L. Page. The anatomy of a search engine. In *Proceedings of the Seventh International World Wide Web Conference*, pages 107–117. Amsterdam: Elsevier, 1998.

[3] S. Bullock and P. M. Todd. Made to measure: Ecological rationality in structured environments. *Minds and Machines*, 9(4):497–541, 1999.

[4] P. Carruthers. Moderately massive modularity. In A. O'Hear, editor, *Mind and Persons*. Cambridge University Press, Cambridge, 2003. URL http://www.philosophy.umd.edu/Faculty/pcarruthers/Moderate-modularity.htm.

[5] P. Carruthers. Simple heuristics meet massive modularity. In S. S. P. Carruthers, S. Laurence, editor, *The Innate Mind*. Oxford University Press, Oxford, 2006.

[6] L. D. Catledge and J. E. Pitkow. Characterizing browsing strategies in the world-wide web. *Computer Networks and ISDN Systems*, 27(6):1065–1073, 1995.

[7] J. Cho and S. Roy. Impact of search engines on page popularity. In *WWW '04: Proceedings of the 13th international conference on World Wide Web*, pages 20–29, New York, NY, USA, 2004. ACM Press. ISBN 158113844X. doi: 10.1145/988672.988676. URL http://portal.acm.org/citation.cfm?id=988672.988676.

[8] A. Clark. *Natural born cyborgs: Minds, technologies and the future of human intelligence*. Oxford: Oxford University Press, 2003.

[9] E. V. Clark. Color, reference, and expertise in language acquisition. *Journal of Experimental Child Psychology*, 94(4):339–343, August 2006. doi: 10.1016/j.jecp.2006.03.002. URL http://dx.doi.org/10.1016/j.jecp.2006.03.002.

[10] P. de Brabanter, D. Nicolas, I. Stojanovic, and N. Villanueva Fernández. Deferential utterances. In *"Referring to Objects" Interdisciplines Virtual Workshop*, 2005. URL http://www.interdisciplines.org/objects/papers/1.

[11] A. H. Eagly and S. Chaiken. *The psychology of attitudes*. Harcourt Brace Jovanovich., New York, 1993.

[12] A. H. Eagly and S. Chaiken. Attitude strength, attitude structure, and resistance to change. In Petty and Krosnick [23].

[13] B. J. Fogg. Prominence-interpretation theory: explaining how people assess credibility online. In *CHI '03: CHI '03 extended abstracts on Human factors in computing systems*, pages 722–723, New York, NY, USA, 2003. ACM Press. ISBN 1581136374. doi: 10.1145/765891.765951. URL http://portal.acm.org/citation.cfm?id=765951.

[14] B. J. Fogg and H. Tseng. The elements of computer credibility. In *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 80–87, New York, NY, USA, 1999. ACM Press. ISBN 0201485591. doi: 10.1145/302979.303001. URL http://portal.acm.org/citation.cfm?id=303001.

[15] B. J. Fogg, C. Soohoo, D. R. Danielson, L. Marable, J. Stanford, and E. R. Tauber. How do users evaluate the credibility of Web sites? A study with over 2,500 participants. In *DUX '03: Proceedings of the 2003 conference on Designing for user experiences*, pages 1–15, New York, NY, USA, 2003. ACM Press. ISBN 1581137281. doi: 10.1145/997078.997097. URL http://portal.acm.org/citation.cfm?id=997078.997097.

[16] R. Hogarth. *Judgment and choice. The psychology of decision*. John Wiley & Sons, Inc., New York, 1987.

[17] B. A. Huberman. *The Laws of the Web : Patterns in the Ecology of Information*. The MIT Press, April 2003. ISBN 0262582252.

[18] J. M. Joyce. VIII–Epistemic Deference: The Case of Chance. *Proceedings of the Aristotelian Society*, 107(1pt2):187–206, 2007. doi: 10.1111/j.1467-9264.2007.00218.x. URL http://www.blackwell-synergy.com/doi/abs/10.1111/j.1467-9264.2007.00218.x.

[19] J. Nielsen. Alertbox: Information Pollution. Website, August 11 2003. URL http://www.useit.com/alertbox/20030811.html.

[20] J. Nielsen. Alertbox: When Search Engines Become Answer Engines. Website, August 16 2004. URL http://www.useit.com/alertbox/20040816.html.

[21] M. O'Brien and M. T. Keane. Modeling result-list searching in the world wide web: The role of relevance topologies and trust bias. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, 2006.

[22] R. E. Petty and J. Cacioppo. The elaboration likelihood model of persuasion. *Advances in Experimental Social Psychology*, 19:123–205, 1986.

[23] R. E. Petty and J. A. Krosnick, editors. *Attitude strength: Antecedents and consequences*. Erlbaum, Mahwah, NJ, 1995.

[24] R. E. Petty and D. T. Wegener. The elaboration likelihood model: Current status and controversies. In S. Chaiken and Y. Trope, editors, *Dual-Process Theories in Social Psychology*. Guilford Press, New York, 1999.

[25] R. E. Petty, C. Haugtvedt, and S. M. Smith. Elaboration as a determinant of attitude strength: Creating attitudes that are persistent, resistant, and predictive of behavior. In Petty and Krosnick [23].

[26] P. Pirolli. Rational analyses of information foraging on the web. *Cognitive Science*, 29:343–373, 2005.

[27] P. Pirolli. *Information Foraging Theory: Adaptive Interaction with Information (Oxford Series in Human-Technology Interaction)*. Oxford University Press, USA, March 2007. ISBN 0195173325.

[28] P. Pirolli and S. Card. The evolutionary ecology of information foraging, 1997. URL http://citeseer.ist.psu.edu/pirolli97evolutionary.html.

[29] P. Pirolli and S. Card. Information foraging. *Psychological Review*, 106(4): 643–675, 1999.

[30] J. E. Pitkow and M. Recker. Results from the First World-Wide Web User Survey. *Computer Networks and ISDN Systems*, 27(2):243–254, 1994.

[31] F. Recanati. *Oratio Obliqua, Oratio Recta*. MIT Press, Cambridge, MA, 2000.

[32] S. Y. Rieh. Judgment of information quality and cognitive authority in the Web. *Journal of the American Society for Information Science and Technology*, 53(2): 145–161, 2002. doi: 10.1002/asi.10017. URL http://dx.doi.org/10.1002/asi.10017.

[33] S. Y. Rieh and D. R. Danielson. Credibility: A multidisciplinary framework. In B. Cronin, editor, *Annual Review of Information Science and Technology*, volume 41, pages 307–364. Information Today, 2007.

[34] D. Sperber. In defense of massive modularity. In E. Dupoux, editor, *Language, Brain and Cognitive Development: Essays in Honor of Jacques Mehler*. MIT Press, Cambridge, MA, 2002. URL http://www.dan.sperber.com/modularity.htm.

[35] K. Sterelny. Cognitive load and human decision, or, three ways of rolling the rock up hill. In P. Carruthers, S. Laurence, and S. Stich, editors, *The Innate Mind: Culture and Cognition*. Cambridge University Press, Cambridge, 2006.

[36] S. S. Sundar. Technology and credibility: Cognitive heuristics cued by modality, agency, interactivity and navigability. In M. J. Metzger and A. J. Flanagin, editors, *Digital Media, Youth, and Credibility (John D. and Catherine T. MacArthur Foundation Series on Digital Media and Learning)*. The MIT Press, December 2007. ISBN 0262562324.

[37] P. M. Todd. Fast and frugal heuristics for environmentally bounded minds. In G. Gigerenzer and R. Selten, editors, *Bounded Rationality*. MIT Press, Cambridge, MA, 1999.

[38] N. C. Wathen and J. Burkell. Believe it or not: Factors influencing credibility on the web. *Journal of the American Society for Information Science and Technology*, 53(2):134–144, 2002. doi: 10.1002/asi.10016. URL http://dx.doi.org/10.1002/asi.10016.

[39] Y. Yanbe, A. Jatowt, S. Nakamura, and K. Tanaka. Can social bookmarking enhance search in the web? In *JCDL '07: Proceedings of the 2007 conference on Digital libraries*, pages 107–116, New York, NY, USA, 2007. ACM Press. ISBN 9781595936448. doi: 10.1145/1255175.1255198. URL http://dx.doi.org/10.1145/1255175.1255198.