Adam Wierzbicki

Web Content Credibility

Springer

Web Content Credibility

Adam Wierzbicki

Web Content Credibility

Springer

Adam Wierzbicki
Polish-Japanese Institute of Information Technology
Warsaw, Poland

ISBN 978-3-319-77793-1 ISBN 978-3-319-77794-8 (eBook)
https://doi.org/10.1007/978-3-319-77794-8

Library of Congress Control Number: 2018941473

*To all participants of the **Reconcile project**. Thank you for a rewarding and inspiring collaboration.*

***To my mother and father**, with love and respect.*

***To my wife Małgosia**, who is my best friend, partner, and my better half.*

***To my son Tomek** and his sisters Asia, Róża, and Ania. The future belongs to them.*

Preface

The lack of support for credibility evaluation is one of the major weaknesses of the Web today. Despite mature search services, as well as advanced content filtering and recommendation methods, only a few services attempt to support credibility evaluation on the Web. This situation allows the spread of fake news, rumors, and factually incorrect Web content. The increasing reliance of our society on the Web makes solving this problem an important goal for informatics.

This book has been written for an audience of researchers and developers who wish to design and implement credibility evaluation support systems for the Web. I hope that the book can serve as a basis to accelerate the research in this area.

The study of Web content credibility evaluation is an interdisciplinary area. The book attempts to bridge the gap between research on credibility evaluation in media science, social science, psychology, and informatics. The book contributes operational definitions, as well as models, of source and message credibility that base on theoretical work in disciplines other than informatics.

Supporting better Web content credibility evaluation is an important social goal. Because of this, the book falls into the broader discipline of social informatics: *a discipline of informatics that studies how information systems can realize social goals, use social concepts, or become sources of information about social phenomena*.

This book is based on research supported by the grant "Reconcile: Robust Online Credibility Evaluation of Web Content" from Switzerland through the Swiss Contribution to the enlarged European Union (http://reconcile.pja.edu.pl).

Warsaw, Poland
May 2018

Adam Wierzbicki

Contents

Chapter 1
Introduction

> Lying only works if there is first a mutual assumption of cooperation and trust: you only lie because you know that I will trust your information as truthful and act accordingly.
> Michael Tomasello

Credibility has recently become a hot topic in Web content research. Companies such as Google aim to discern the veracity of statements of fact contained in webpages.[1] Crowdsourced services striving to filter out non-credible information have been subject to research and are applied in practice. Among systems using that approach are the Article Feedback Tool on Wikipedia, the TweetCred[2] system for Twitter, or the WOT system for evaluating Web portal credibility.[3] Content evaluation can also be supported by machine classification approaches that attempt to learn quality ratings and predict the ratings of new content.

Diverse techniques have been applied to address the problem of supporting Web content credibility evaluation. They range from Crowdsourcing, machine classification, reputation systems, to Natural Language Processing. The domains of application also vary: credibility evaluation has been applied in the domain of health information on the Web (Health on the Net (HON)[4]) as well as in the domain of politics and media (PolitiFact[5]). The first goal of this book is to provide a sound theoretical foundation for research on Web content credibility and present an overview of the state of the art. This goal is achieved in part by presenting openly available datasets (such as the Web Content Credibility Corpus produced

[1] http://www.newscientist.com/article/mg22530102.600-google-wants-to-rank-websites-based-on-facts-not-links.html.

[2] http://twitdigest.iiitd.edu.in/TweetCred/.

[3] http://www.mywot.com/.

[4] http://www.hon.ch/.

[5] http://politifact.com/.

A. Wierzbicki, *Web Content Credibility*,
https://doi.org/10.1007/978-3-319-77794-8_1

as a result of the Reconcile project[6]) for reproducible evaluation of algorithms for Web content credibility evaluation, as well as methodologies for carrying out experimental evaluations of systems for supporting Web content credibility evaluation of users. The book contributes to theory by formulating a new definition of credibility (Chap. 2) and theoretical models of credibility based on game theory (Chap. 5). Another goal of the book is to provide reference models and designs of Web content credibility evaluation support methods, taking into account the diversity of these methods for various types of Web content. This goal is achieved in Chaps. 3 and 4.

1.1 Credibility and Relevance of Web Content

The enormous utility and ubiquity of the Web are due not merely to the Web's information content, but rather to the increasingly intelligent functions of the Web that enable the discovery of the most useful Web content. Among these functions, Web information retrieval (Web search) plays a central role. Web search is an unquestionable success of information and computer science, as well as one of the most prolific research areas in these disciplines. There are widely used and highly effective practical Web search systems, but even today there are ongoing projects to develop new search engines (such as Wikipedia's new search engine project, the Knowledge Engine,[7] or continuous commercial efforts).

This striking success of Web information retrieval is in contrast to the current state of the art of research on Web content credibility evaluation support. While research on credibility is very active, it has a long way to go before reaching the mature state of research on Web searching. In order to try to understand the reasons for this difference, one must first consider the reasons for success of Web information retrieval.

Research on Web information retrieval has a very long history when compared to other research areas concerning the Web (semantic Web, Web intelligence, or Web content credibility evaluation). First, large-scale information retrieval systems have been developed in the early 1970s. It can be said that the commercialization of the Web in the early 1990s merely gave a new application domain to an already mature research area. Later research that resulted in breakthroughs, such as the invention of PageRank in 1996, may have been inspired by earlier results on the use of eigenvalues in scientometrics for ranking scientific journals, published 20 years before.

Another important factor in the successful development of Web information retrieval is the research methodology. In 1992, the US Department of Defense along with the National Institute of Standards and Technology (NIST) organized

[6] http://reconcile.pjwstk.edu.pl/.

[7] https://en.wikipedia.org/wiki/Knowledge_Engine_(Wikimedia_Foundation).

the Text REtrieval Conference (TREC) that became a series of annual conferences, continuing until today. TREC allow for evaluation and comparison of competing Web search systems and algorithms on the same, public datasets. According to the organizers (NIST), in the first 6 years of the TREC organization, the effectiveness of Web information retrieval systems approximately doubled. TREC today present challenges for the research community in various areas of Web information retrieval, including the newer discipline of Human-Computer Information Retrieval (HCIR) which tackles issues such as automatic query reformulation, faceted search, lookahead, or relevance feedback.

Time and a good methodology brought about the development of sophisticated algorithms and some of the most advanced information systems practically used today (such as the Google search engine). Web search algorithms use diverse mathematical models, such as set-theoretic models, algebraic models, and probabilistic models. These models can take into account term interdependencies, as well as contextual information.

A final, and perhaps most important, reason for the success of research on Web information retrieval is the good understanding of basic concepts involved in an evaluation of search results by computer scientists. In Web information retrieval, the notion of quality of retrieved documents is called relevance. Relevance is understood by computer scientists as a measure of how well retrieved Web content matches the information need of the user. While it is understood that ultimately relevance judgments are made by humans, the success of currently used algorithms suggests that there is little inherent difficulty in predicting human relevance judgments based on available data.

The state of research on Web content credibility evaluation is far different. From the beginning, the concept of credibility has been considered by computer scientists as difficult to understand, fuzzy, and uncertain. One of the reasons for this difference in understanding of the two concepts of relevance and credibility may lie in the origin of early credibility research: media science and social science. The goals of these two fields of science are different from the goals of computer science or information science, which are applied sciences. Media science and social science aim to achieve an in-depth understanding of studied social phenomena or concepts, taking into account all possible aspects. On the other hand, informatics requires an operational understanding of the concept of credibility, because this is a prerequisite for developing information systems and algorithms that may be used in practice.

The difficulty of understanding credibility of information, as well as the late start of research on Web content credibility in computer and information science—which dates back to the 1990s—results in the current state of the Web that can successfully support search, but has little functions or tools for supporting Web content credibility evaluation. This does not mean that credibility is less important than relevance on the Web. In fact, research on the nature of human communication and human thought suggests the opposite, as explained in the next section.

1.1.1 Credibility and Relevance in Human Communication

The human condition has been subject to inquiry since the ancient times, as the entire discipline of philosophy has aimed to understand what makes us human. Recently, however, this question has become a subject of empirical sciences such as anthropology, evolutionary psychology, or evolutionary linguistics, and significant progress has been made. Scientists seek to understand what makes us human by experimentally investigating primate behavior and comparing the results to human behavior (especially child behavior). Seminal research by Michael Tomasello [192], as well as by Robin Dunbar, has shed new insight on the human condition and highlighted the importance of two specifically human concepts: credibility and relevance.

Tomasello argues that human cognition is special due to its focus on collaborative communication and emergent states of the mind called joint attention and joint intentionality. In short, humans are able to understand not just the received communication but also the intentions of the communicating partner and to act accordingly based on this understanding.[8] This process is, in fact, recursive, as a communicating human is able to anticipate the understanding of their intentions by the receiver, and can take this knowledge into account. However, human communication is essentially collaborative: its goal is a joint action that should reach a joint goal. This is why it is so important that this communication should be true. To quote Tomasello,

> The notion of truth entered the human psyche not with the advent of individual rationality (...), but, rather, with the advent of joint intentionality and its focus on communicating cooperatively with others.
> *A Natural History of Human Thinking*. Michael Tomasello

While Tomasello does not explicitly use the term, evaluating credibility (or believability) of communicated information is an essential part of his theory of what makes us human. Humans do not merely communicate informative facts; rather, they commit to the truth of a statement and back up this commitment with reasons and justifications as necessary. This kind of cooperative argumentation is another important human characteristic, requiring a special way of thinking, which Tomasello refers to as the "Web of beliefs":

> The ability to connect thoughts to other thoughts (both those of others and one's own) by various inferential relations (prototypically by providing reasons and justifications) is key to human reason in general, and it leads to a kind of interconnection among all of an individual's potential thoughts in a kind of holistic "web of beliefs."
> *A Natural History of Human Thinking*. Michael Tomasello

[8] Other neuroscientists, such as Graziano, go even further and hypothesize that perceiving and understanding the intentions, emotions, and minds of others is the basis of perception of our own consciousness, as well as the basis of spiritual and religious beliefs [53].

Tomasello's idea of recursive credibility evaluations points out the fact that not just the receiver but also the sender of information evaluates the information's credibility. Furthermore, the goal of the sender's evaluation is different from the receiver's. The receiver wishes to know whether the received information is true (or sufficiently reliable). The sender's credibility evaluation of his own information, however, has the goal of predicting whether the information will be judged as credible by the receiver. These two (often conflicting) goals point out the essential difference between truth and credibility: the receiver wishes for information that is true, while the sender wishes to produce information that is credible. Further, it becomes clear that the sender has many means to modify his information in order to persuade the receiver that it is credible. In fact, humans' invention of truth also resulted in the invention of lying, as succinctly expressed by Tomasello in the motto of this chapter. The differences between truth and credibility leads to a new definition of credibility, discussed in Sect. 2.3.

Another basic characteristic of human thinking is an evaluation of relevance. However, according to Tomasello, relevance evaluations are actually secondary to credibility evaluations: they can only be made once the receiver is assured (or assumes) that the received information is credible. Then, it becomes important to understand why the information is relevant for the receiver:

> The second important consequence of this new cooperative way of communicating was that it created a new kind of inference, namely, a relevance inference. The recipient of a cooperative communicative act asks him or herself: given that we know together that he is trying to help me, why does he think that I will find the situation that he is pointing out to me relevant to my concerns.
> *A Natural History of Human Thinking*. Michael Tomasello

The relevance judgment is another human ability that requires the establishment of a common (joint) basis of facts and attention between the receiver and the communicating partner: the Web of beliefs is therefore also fundamentally important for relevance evaluations. Tomasello's view on the fundamental role of credibility also points out an underlying assumption on Web information retrieval that is so obvious that one usually forgets about it. Obviously, retrieving the most relevant document will result in the best satisfactions of a user's information needs, *assuming that the retrieved document is credible*.

Tomasello's theory on what makes human thinking special shows that the two concepts of credibility and relevance are inherently connected. They are also two important, basic human concepts, similar to trust and fairness. Since much, if not the majority, of human communication occurs on the Web today, a solid understanding of credibility and relevance should therefore be the basis of any technological solutions that support communication or information retrieval on the Web. The concept of relevance is the basis of Web information retrieval, and has been studied in computer and information sciences for several decades, resulting in enormously successful technologies such as contemporary search engines. On the other hand, the equally important concept of credibility has not yet reached a similar understanding. The aim of this book is to make a first step towards changing this situation.

1.1.2 Epistemic Similarities of Credibility and Relevance Judgments of Web Content

The predominant view on the main reasons (besides lack of time) for the less advanced state of the art of research on Web content credibility evaluation support is that it's a result of lacking understanding of credibility in computer science. What follows is the comparison of the concepts of credibility and relevance from a methodological point of view. The goal is to understand whether credibility is inherently more difficult to investigate than relevance.

To compare the difficulty of researching credibility and relevance, it is possible to start by comparing the methodologies currently used to further this research. Credibility is the first issue to be addressed. Here, the main methodology is empirical: researchers conduct experiments during which human users evaluate Web content from specially prepared corpora. These experiments are used to gather datasets of human credibility evaluations, which are then used to study factors affecting human evaluations, or to train and evaluate machine learning algorithms that aim to predict credibility evaluations. Several such experiments will be described in detail in the second chapter of this book.

The research methodology used in Web information retrieval is best described by the methodology of the TREC. First, large corpora of Web content are prepared by the conference organizers. These corpora depend on the track; however, for simple Web search, the corpora are usually real webpages or documents, chosen to reproduce topical diversity. Next, the organizers prepare search topics that represent a user's information needs. The description of the topics is quite detailed, for example [60]:

> A relevant document will either report an impending rail strike, describing the conditions which may lead to a strike, or will provide an update on an ongoing strike. To be relevant, the document will identify the location of the strike or potential strike. For an impending strike, the document will report the status of negotiations, contract talks, etc. to enable an assessment of the probability of a strike. For an ongoing strike, the document will report the length of the strike to the current date and the status of negotiations or mediation.

Finally, TREC participants devised algorithms for searching that took as input the document corpora and the search topics. The result was a ranking of documents. TREC organizers used a pooling method that took the documents from the top of the ranking of all algorithms for a given search topic and merged them into one set (removing repetitions). The relevance of these pooled documents to the search topic was then evaluated by a single expert [60, 61].

To summarize, relevance judgments used to evaluate and improve Web information retrieval systems and algorithms at the TREC were made by human evaluators. This is a similar approach to the one used in research on Web content credibility today.

In order to further the comparison of epistemic difficulty involved in research of credibility and relevance on the Web, one may wonder whether relevance judgments are more or less uncertain or subjective than credibility judgments.

This question could only be answered with certainty by a comparison of datasets that contain multiple human evaluations of credibility or relevance per document. Unfortunately, such datasets are not available today. It is, however, possible to carry out a thought experiment that would attempt to apply one of the theories developed to understand Web content credibility judgments to relevance judgments. The reader is left to judge whether the theory proposed to understand credibility applies well to relevance judgments. If it does, then relevance judgments should not be any easier to understand and study than credibility judgments.

The theory proposed to apply to relevance judgments is called Prominence-Interpretation theory. It has been introduced by Fogg based on research by the Stanford team on Web content credibility [40] during 1999–2003. According to the Prominence-Interpretation theory, there are two stages in Web users' credibility evaluations. In the first stage, users investigate the Web content and notice some credibility cues. Some other (possibly important) cues may go unnoticed at this stage. Prominence is the likelihood that an information contained in a webpage will be noticed or perceived by the user. Fogg claims that at least five factors influence prominence: the motivation ("involvement") of the user, the ability ("experience") of the user, the task of the user, user's individual characteristics, and the topic of the webpage.

During the second stage, Interpretation, users make judgments about the cues they have noticed in the first stage. According to Fogg, many factors can influence interpretation: a user's level of competency (expertise) regarding the webpages topic, the context of the evaluation (consisting of the user's environment, expectations, and situational norms), as well as a user's assumptions consisting of heuristics, past experience, and culture.

In this thought experiment, the example search topic from the TREC, quoted above, is considered. A situation is given when Web users, or evaluating experts, need to judge the relevance of a retrieved webpage to the search topic. The judgment would not necessarily be binary, but could be done on some evaluation scale (such scales will be discussed in more detail in the next chapter). The search topic description asks the user to look for several kinds of information. The first question is whether the retrieved webpage is related to strikes. However, it should not be too difficult for a good retrieval algorithm to disregard pages that do not include strikes as the keyword (the only problem could be multiple meanings of a search term). Next, the user has to check whether the webpage contains information on the strike's location, ongoing negotiations, date and length of the strike, and current status of the strike or negotiations. The presence or absence of any of these kinds of information should affect the relevance judgment. However, it is clear that in a larger webpage, depending on user's attention, concentration, and time used for evaluation, some of these kinds of information may go unnoticed by the user (even if they are present in the page). These factors are influenced by the user's involvement, experience, and individual characteristics, as foreseen by the first stage, Prominence, of the Prominence-Interpretation theory.

Once the user has searched for the important content in the webpage, he or she still needs to make a relevance judgment. The search topic does not prioritize

the kinds of information to look for (with the exception of information about a strike's location, which is explicitly mentioned as being a necessary condition for relevance). Therefore, the user is left to decide how relevant a webpage is that, for example, contains information about a strike's location and time, but no information about negotiation or current status. This stage requires an interpretation of the kinds of information in the search topic. Such decisions by the Web user clearly resemble the Interpretation stage of the Prominence-Interpretation theory. Clearly, the decision would also depend on a user's experience and understanding of the topic. For example, expert users who have experience in evaluating relevance may find it easier than ordinary users to interpret the information contained in the webpage. Still, even for an expert, determining the relevance ranking of two webpages about strikes may be difficult, if—for example—one webpage contains only information about the time of the strike, while the other only about the negotiation status.

The reader is left to judge for himself whether the proposed application of the Prominence-Interpretation theory is valid. There is, at the very least, a strong analogy between the decision process of evaluating the relevance and the credibility of a webpage. The conclusion is that relevance judgments are not any easier to make, nor less dependent on a user's attention, concentration, and experience than credibility judgments. Therefore, the creation of information systems and algorithms that support the evaluation of Web content credibility should in principle be not more difficult than the creation of such tools for supporting Web search. This is an optimistic conclusion for researchers working on Web content credibility evaluation support. However, an improved understanding of the concept of credibility by computer science is a prerequisite for advancing the state of the art of research on Web credibility. It is also one of the main goals of this book.

1.2 Why Does Credibility Evaluation Support Matter on the Web?

The current World Wide Web is characterized by two factors: a low price of producing Web content and very high incentives for producing Web content that can influence the beliefs of Web users. The incentives are, firstly, of commercial nature, which accounts for the rapid development of Web-based marketing and advertising, as well as the success of Google. However, Web content affects not just our purchasing decisions but also several other aspects of human behavior, ranging from lifestyle to political decisions.

Credibility evaluations of Web content are regularly made by ordinary users. The goal of research and development on Web content credibility evaluation should be to support and improve these evaluations, so that Web users can distinguish between truthful and untrue Web content (whenever this distinction is possible). The following section contains examples of Web content that is untrue, and can potentially have a high social impact. It also presents examples of existing services aiming to support Web content credibility evaluation.

1.2.1 Examples of Non-credible Medical Web Content

The Web contains information on almost all aspects of human behavior, and human knowledge. In order to focus our analysis on the credibility of current Web content, the discussion will be limited to a few areas of Web content. The first is Web content related to health and medicine. This is a large and important category of Web content, to the extent that Web users have coined the term "Dr Google." The medical domain is important because it is an instance of a larger category of Web content based on factual, specialized knowledge.

This section describes just three examples of non-credible medical Web content, which should be enough for the reader to realize the potential impact of this content, and therefore the significance of research on Web content credibility evaluation support in the medical domain. Still, it is important to remember that these examples are just a tip of the iceberg of non-credible medical content on the Web.

1.2.1.1 Vaccines and Autism

A Web search for the query "vaccines and autism" will have a plethora of results. Many of the found webpages will still support the view that a connection between vaccines and autism exists, although as of now a majority of webpages will refute the claim. The supposed connection between vaccines and autism dates back to the thiomersal controversy. The causal link between the mercury-based vaccine preservative, thiomersal, and autism has been proposed in a scientific article published in 1998 by *The Lancet*, a highly reputed medical journal. The author of the study, Andrew Wakefield, has been found guilty of dishonesty by the British General Medical Council in 2010[9] and subsequently barred from practicing medicine in the UK. *The Lancet* has retracted Wakefield's article.[10] However, for 12 years, the controversy remained unresolved in the scientific community, and even today it lingers in the opinion of the public.

The anti-vaccination movement is a global community, largely present on the Web. One of the strongest supporters of the link between vaccines and autism was Jenny McCarthy, an American model and celebrity. Her son was diagnosed with autism at the age of 2.5, following a series of vaccinations. McCarthy had launched a foundation called "Generation Rescue" that promotes alternative therapies and alternative theories debating environmental causes of autism, including vaccination. The "Generation Rescue" website is an excellent example of the difficulty involved in evaluating Web content credibility. It is professionally designed, informative, and supported by a large community of users. Nevertheless, it is also rated as non-credible by existing credibility evaluation support tools, such as WOT.

[9] http://briandeer.com/solved/gmc-charge-sheet.pdf.

[10] http://www.thelancet.com/journals/lancet/article/PIIS0140-6736(97)11096-0/abstract.

The proposed example is one of many, and a good example of a specific type of medical Web content: content based on scientific facts (frequently, however, subject to simplification or incorrect interpretation). Only the scientific community could have refuted the claim underlying the belief that vaccines cause autism. The majority of this community has agreed that this belief is unsubstantiated. However, the process was time-consuming and obscure to the lay Web users and thus could not avert the general surmise that vaccines may indeed cause autism. This resulted in a decrease in the number of vaccinations in the USA and worldwide. The seriousness of the situation manifests in increasingly large measles outbreaks, such as the recent outbreak in 2015 in Germany. It involved over 570 reported measles cases and caused infant deaths.

1.2.1.2 Consuming Placenta

Examples of non-credible medical Web content can be much more humorous than the case of the anti-vaccination movement. An increasingly popular, and probably harmless, trend is for women to consume their own placenta post-partum. Ingesting one's own placenta is supposed to help avoid after-birth mood swings or depression (the so-called baby blues). The trend has supporters among celebrities, including Kim Kardashian.[11] This example is food for thought regarding the link between non-credible, controversial health and medical decisions, and the notoriety (and public exposure) of celebrities. Whatever the reason for her decision, Kardashian's impact on other women may be similar to her influencing women by declaring that she is using a new cosmetic.

Similarly to the previous example, a critical Web user should be able to find reputable Web content that refutes the claim that eating a placenta has health benefits.[12] However, the exposure of such content on the Web is much lower, proving that medical professionals are no match for celebrities online.

1.2.1.3 Colloidal Silver

The Web is a medium that is almost ideally suited for the proliferation of alternative medicine. One of the reasons is that many of the commodities required for alternative therapies can be readily purchased on the Web. One such example is colloidal silver. A search query for this term will result in many pages that recommend supposed benefits of using colloidal silver for the treatment of infections, ranging from ear infections to pneumonia, as well as on wounds or lesions.[13] This view is

[11] http://www.mirror.co.uk/3am/celebrity-news/kim-kardashian-eating-placenta-again-7013283.

[12] https://www.nichd.nih.gov/news/releases/Pages/062615-podcast-placenta-consumption.aspx.

[13] https://draxe.com/colloidal-silver-benefits/.

not supported by established medicine.[14] Rather than healing sinus trouble, colloidal silver (if ingested) can cause argyria, a permanent bluish-gray discoloration of the skin. Similarly to placenta ingestion, the use of colloidal silver is promoted by celebrities: this time, it is the well-known actress, Gwyneth Paltrow, who promoted the use of a spray with colloidal silver in the American health show, *Dr Oz*.[15] The price of colloidal silver health supplements ranges from 10 to 60 US dollars on Amazon (where it is one of the most popular health products).

The above three examples have one thing in common: Web content from reputable sources (such as the National Institutes of Health[16]) exists that contradicts non-credible claims present in other places on the Web. Critical Web users, who are willing to spend time and effort on verifying medical claims made on the Web, should be able to correctly evaluate the credibility of content from these three examples. This might not be the case for medical content that concerns new or emerging alternative medicine therapies or unconfirmed medical claims.

1.2.2 Fake News in Web-Based Social Media

"Fake news" is a term that has "gone viral" after the 2016 American presidential election. While the existence and potential political consequence of fake news is probably as old as civilization itself, the situation in 2016 has been significantly different: in the USA, over 60% of adult citizens get their news on social media [4]. Popular fake news can be more widely shared on Facebook than mainstream news. Last but not least, fake news stories circulated on social media during the 2016 election in the USA has favored one candidate at the expense of another. The ratio of the amount of fake news that were pro-Trump (or anti-Clinton) with respect to those that were pro-Clinton (or anti-Trump) was almost 3:1, while the ratio of Facebook shares of these fake news is over 4:1. This strong partisanship has most likely contributed to the results of the American presidential election in 2016 [4].

Here, let us give one striking example of political fake news from the American election in 2016. While many other examples have been identified, the case of the infamous "Pizzagate" is sufficient to describe the problem.

In early November 2016, Internet users received access to the e-mails of John Podesta, Hillary Clinton's campaign manager. The e-mails have been released on WikiLeaks after a successful phishing attack. At the same time, another anonymous source released the claim that the New York City Police Department (NYPD) was investigating a pedophile ring linked to members of the Democratic Party. The combination of these two news has led to claims that food-related words in Podesta's e-mails were code words for pedophilia, sex trafficking, and even satanist ritual

[14] https://nccih.nih.gov/health/silver.

[15] http://colloidalsilversecrets.blogspot.com/2013_08_01_archive.html.

[16] https://nccih.nih.gov.

practices. For example, the words "cheese pizza" were claimed to be a code word for "child pornography," since they have the same initials.

Altogether, claims related to "Pizzagate" formed a full-fledged conspiracy theory. Multiple fake news emerged from that theory, such as the alleged raid of the NYPD on Hillary Clinton's property, or a claim that the FBI had confirmed the existence of an underground sex network. The spread of these fake news on Twitter and several news websites was viral. It is estimated that over 1 million tweets related to "Pizzagate" have been published on Twitter by the end of the 2016 US election.

Apart from political consequences, the fake news' spread on social media resulted in real consequences that affected owners of restaurants claimed to be involved in sex-ring meetings. These businesses received large amounts of threatening phone calls, including death threats, and also experienced online harassment. The public reaction to fake news culminated in a shooting in one of the restaurants, Comet Ping Pong.[17] A self-proclaimed "investigator" fired three rifle shots in the restaurant, but let himself be arrested without resistance afterwards. Luckily, no one was hurt.

The "Pizzagate" conspiracy theory has been finally discredited on the Internet. A significant role in the process has been played by one of the credibility evaluation support systems described in the next section: Snopes.com. Mainstream news organizations, such as *The New York Times*, *The Washington Post*, and others, have also played a part in discrediting and denouncing "Pizzagate." Debunking of this conspiracy theory involved disproval of several fake news that used images of children from Instagram and falsely claimed that these were sex-ring victims, kept in a non-existent basement of the Comet Ping Pong restaurant. Despite early warnings that the news related to "Pizzagate" were fake, the "Pizzagate" theory continued to spread on social media and alternative Web-based news sites. It is important to note that an author of a fake news article about "Pizzagate" has been quoted to express her satisfaction by saying: "It's honestly really grown our audience."[18] This honest admission clearly shows the economic motivation of Web-based content producers in the production and dissemination of fake news, as well as other non-credible Web content.

1.2.3 Examples of Credibility Evaluation Support Systems

As it was already pointed out on the example of vaccines and autism, there are services aiming to support Web users in the evaluation of Web content credibility. In this book, such services will be referred to as credibility evaluation support (CS). While their current impact and popularity is still too low, they point out various

[17] https://en.wikipedia.org/wiki/Pizzagate_conspiracy_theory.

[18] https://www.thestar.com/news/canada/2016/12/07/belleville-woman-helped-cook-up-pizzagate.html.

possibilities for future research and development. A common feature of all the discussed services is that they are based on human evaluations, sometimes using a Crowdsourcing approach, and sometimes relying on experts.

1.2.3.1 Health on the Net

The first example comes from the medical domain and provides a follow-up for the discussion of non-credible medical Web content. Although activities of official health institutions (such as the National Institutes of Health) providing access to credible medical content are not a service supporting medical Web content credibility evaluation, such services exist. They are usually based on the principle of certification. One of the services providing certificates to credible medical websites is Health on the Net (HON).[19]

Established in 1996, Health on the Net is a foundation based in Switzerland. HON has been founded by experts on telematics and e-healthcare. HON collaborates with several medical institutions, as well as with the Economic and Social Council of the United Nations, the WHO, and the International Organization for Standardization's (ISO) technical committee for Health Informatics. To date, HON has certified over 8000 websites. While this number may seem small when compared to the number of websites that deliver medical content, it can be said that HON covers websites with highest source credibility.

Web content producers can apply for HON certification. They have to submit their websites for reviewing by a team of HON experts. The review is based on a publicly available set of criteria, summarized in the "HON code." According to the code, medical Web content needs to meet the following requirements:[20]

1. Authoritative. Any medical or health advice provided and hosted on this site will only be given by medically trained and qualified professionals unless a clear statement is made that a piece of advice offered is from a non-medically qualified individual or organisation.
2. Complementarity. The information provided on this site is designed to support, not replace, the relationship that exists between a patient/site visitor and his/her existing physician.
3. Privacy. Confidentiality of data relating to individual patients and visitors to a medical/health website, including their identity, is respected by this website. The website owners undertake to honour or exceed the legal requirements of medical/health information privacy that apply in the country and state where the website and mirror sites are located.
4. Attribution. Where appropriate, information contained on this site will be supported by clear references to source data and, where possible, have specific

[19] www.hon.ch.

[20] http://www.hon.ch/HONcode/Webmasters/Conduct.html.