

TruthSift: A Platform for Collective Rationality

Eric B. Baum

eric@truthsift.com

TruthSift, Princeton, NJ 08540, U.S.A.

TruthSift is a cloud-based platform that logically combines members' contributions into a collective intelligence. Members add statements and directed connectors to diagrams. TruthSift monitors which statements have been logically established by demonstrations for which every challenge raised has been refuted by an established refutation. When members run out of rational objections, the result is a converged diagram succinctly representing the state of knowledge about a topic, including plausible challenges and how they were refuted. Previous computer systems for collaborative intelligence did not have a qualitatively better solution for combining contributions than voting and are subject to groupthink, interest group capture, and inability to follow a multistep logical argument. They did not settle issues automatically point by point and logically propagate the consequences. I review indications that many practically important statements most people believe to be firmly established will be revealed to be firmly refuted upon computer-assisted scrutiny. TruthSift also supports construction of powerful probabilistic models over networks of causes, implications, tests, and necessary factors.

1 Introduction ---

What is a proof? According to the first definition at Dictionary.com, a proof is "evidence sufficient to establish a thing as true, or to produce belief in its truth." In mathematics, a proof is equivalent to a proof tree that starts at axioms, which the participants agree to stipulate, and proceeds by a series of steps that are individually unchallengeable. Each such step logically combines several conclusions previously established or axioms, or both. The proof tree proceeds in this way until it establishes the stated proved conclusion. Mathematicians often raise objections to steps of the proof, but

Readers who are interested in exploring applications of the technology explored in this letter to their own problems or organization, please contact me. I also suggest going to <https://tssciencecollaboration.com/graph/You%20should%20use%20TruthSift%20for%20your%20or%20your%20organizations%20plans%2C%20projects%2C%20documents%2C%20and%20decisions/289/0/-1/-1/0/0#InkNameGraph>

if it is subsequently established that all such objections are invalid or if a workaround is found around the problem, the proof is accepted.

The scientific literature works in a similar way. Each paper adds some novel argument or evidence that previous work is true or is not true or extends it to establish new results. When people run out of valid, novel reasons why something is proved or not proved, what remains is an established theory or a refutation of it or of all its offered proofs.

TruthSift is a platform for diagramming this process and applying it to any statements members care to propose to establish or refute. One may state a topic and add a proof tree for it, which is drawn as a diagram with every step and connection explicit. Statements are simply added as replies to previously added statements, either a supporting reply or a refuting reply. If somebody thinks they have found a hole in a proof at any step or thinks one of the original assumptions needs further proof, they can challenge it, explaining the problem they see. Then the writer of the proof (or others if it is in collaboration mode) may edit the proof to fix the problem or make clearer the explanation if they feel the challenger was simply mistaken and then may counterchallenge the challenge explaining that it had been resolved or mistaken. This can go on recursively, with someone pointing out a hole in the proof used by the counter-challenger that the challenge was invalid. With TruthSift, the whole argument is laid out graphically and essentially block-chained, which should prevent the kind of edit wars that happen for controversial topics on Wikipedia (Attkisson, 2015; Wilson & Likens, 2015). Each challenge or post should state a novel reason, and when the rational arguments are exhausted, as in mathematics, what remains is either a proof of the conclusion or a refutation of it or all of its proofs.

As statements are added to a diagram (see Figures 1, 2, and 3), TruthSift keeps track of what is established and what is refuted, writing "TE" on tentatively established statements and "TR" on tentatively refuted statements ("tentative" meaning based on all the statements that have been added to the graph so far though more replies could be added and change the result). You can instantly tell what is established and what refuted. TruthSift computes this by a simple algorithm that starts at statements with no incoming edges (no incoming challenges or proofs), which are thus unchallenged as assertions that prove themselves, are self-evident, or appeal to an authority everybody trusts. These are considered established. (If you don't trust the authority, that is a valid reason to challenge.) Then it walks up the diagram rating each statement in turn after all its parents have been rated. A statement will be established if none of its challenges are, and if it has proofs, at least one is established. A challenge may request that a proof be added if a statement does not have one already or adequately prove itself, or if it states a reason that the existing established proofs do not, taken together, establish the statement. If a statement has an established challenge or if all of its proofs are refuted, it is refuted.

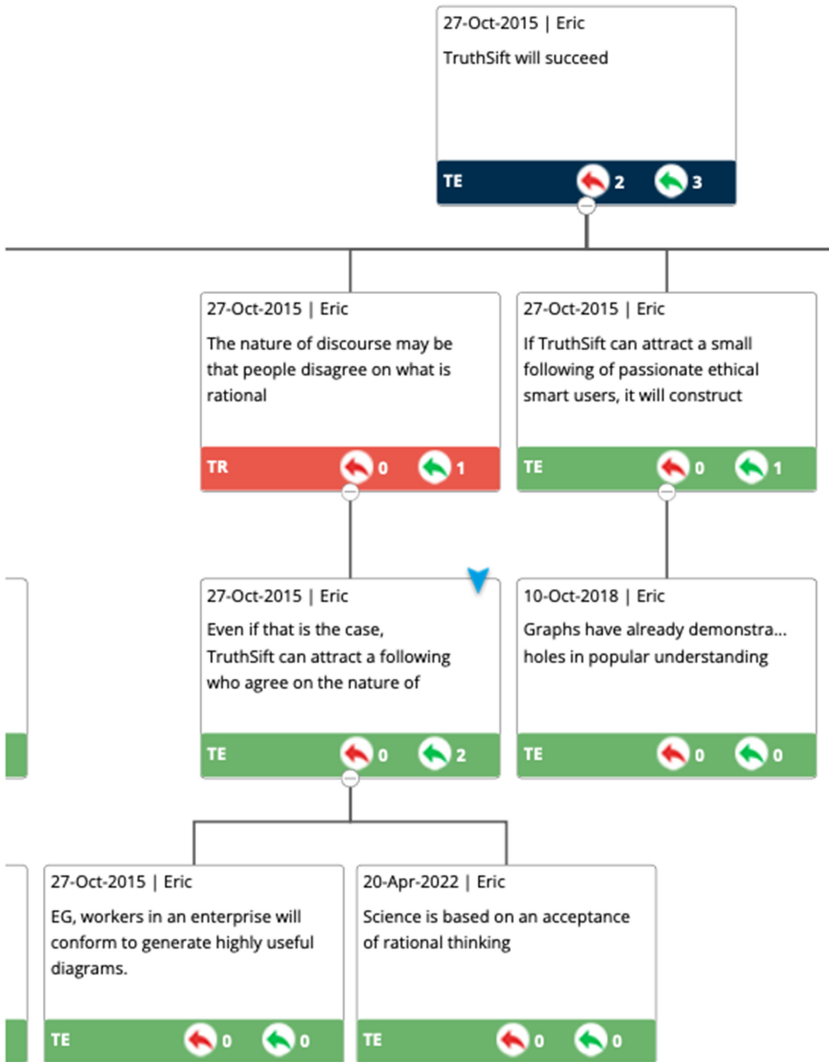


Figure 1: Screenshot of a portion of a graph. The conclusion (a.k.a. the “topic statement”) of the graph is at the top. Replies are rectangles showing the title of the reply, whether it is TE or TR, and with a green bar at the bottom if the reply supports the topic statement (PRO) and a red bar if it opposes it (CON). A supporting reply of the topic statement or of a PRO statement is PRO, as is a rebuttal to a CON statement. A supporting reply of a CON statement is CON, as is a rebuttal to a PRO statement.

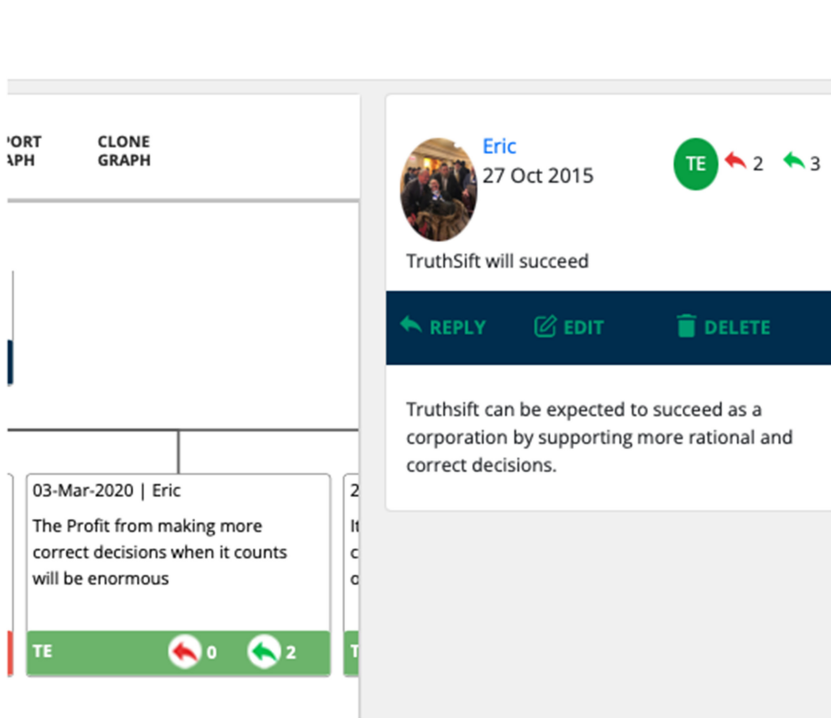


Figure 2: Detailed screenshot of the upper right corner of Figure 1. Whichever node of a graph is selected, that statement is shown on the right side. Selecting “REPLY” pops a window in which you may choose if the reply is SUPPORT or REFUTE and enter content. Thus, the graph can be built up. After each statement is added, TruthSift updates the ratings. The body of the statement defines it.

If it is not refuted, then it follows that all of its challenges are refuted, and it has an established proof or is accepted as providing one, and so it is established. Work your way backward on a diagram, centering on each statement in turn, and examine the reasons why it is established or refuted.

It is important to understand that TruthSift directly maps mathematical practice. If, based on published proofs and refutations and proof fixes and counter-refutations, a rational human would consider a proof to be established by the TruthSift diagram, then so will TruthSift’s simple rating algorithm. In mathematical practice, what it means to be established is whether there is a step-by-step proof offered with all the proposed refutations refuted.

TruthSift implements a mapping of actual human mathematical practice into a machine-human collaboration that enables it to extend to all fields

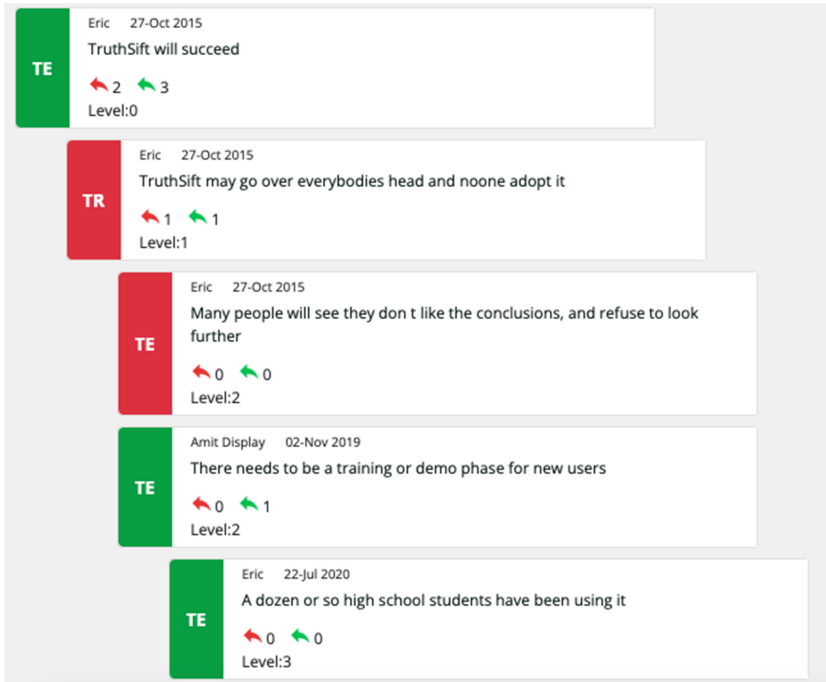


Figure 3: TruthSift also supports an alternative tree view of graphs. In the tree view, the structure of connections is shown by indents, so that just as in making an outline, the user indents one tab each time he or she goes one deeper in the hierarchy and shifts left one tab each time he or she goes up a level in the hierarchy. I find it easier to navigate large graphs in the Tree view and far better for viewing graphs on cell phones. Only the top of this graph is shown.

of human endeavor, not just mathematics. The algorithm decides the same thing a rational human would about what has been established given the contributions to date, and thus serves as an unemotional arbiter, enabling human collaboration, even if it is in some cases adversarial, to achieve superhuman feats of reasoning.

It also forces the discussion to be more precise. If you believe it has been demonstrated that “vaccines are safe,” how much residual damage was allowed in the determination? What precisely is the definition of *safe*? With TruthSift, there is a spelled-out topic statement. Careful phrasing of statements is essential (just as in mathematical practice) to being able to actually establish them and enables more to be established than one might have predicted.

Without TruthSift, people are often unable to agree on what the truth is. With TruthSift, they have a block-chained, semantically organized diagram of the course of their argument, and when they've reached a conclusion, all objections raised have been addressed. TruthSift doesn't always succeed in establishing truth, but it does show when something is provable.

The process divides a problem into natural subtopics that get settled point by point, although not sequentially in time: an old conclusion about a subtopic can always be challenged with a new rational argument or evidence. These different subtopics may involve different collaborators or opponents. The process guides people in the essence of critical thinking and fruitful collaboration.

The process leaves a careful record of how decisions were made and how all the arguments against every subpoint were rebutted. This is like the scientific literature diagrammed and automatically updated as new conclusions are added. This documentation is useful for businesses quite aside from the advantages of getting the right answer and facilitating collaboration. The transparency exposes if some employee has an agenda other than the bottom line or truth. The process is also useful for teaching the essence of critical thinking and would make for far better refereeing and author response process than currently in use in the sciences (TruthSift, n.d.c). The documented decision processes provide data to refine decision making over time.

For TruthSift to work properly, posters should respect the guidelines and make a good-faith effort to post only proof or challenge statements that they believe rationally support or rebut their target. Posts do not have to be correct (that's what challenges are for), but they have to be honest attempts, not spam or ad hominem attacks.

Frequently a user wants to assemble arguments for a proposition stating something like, "The preponderance of the evidence indicates X," and these arguments are not individually proofs of X. It is safe to simply add them as proofs of the proposition. If not enough of them are established, the target may be challenged on that basis. The goal is a diagram that transparently explains a proof and what is wrong with all the objections people have found plausible. Edits that move in that direction are useful and desired.

TruthSift ran a public site where members could debate issues of their choosing. There were a few dozen dedicated users, but we believe that the concepts of establishing proofs are difficult for a consumer audience. Another problem is that many things people believe are not in fact true (Le Bon, 1895; Bernays, 1928; TruthSift, n.d.a), and they are disturbed to see this and turn away from a technology they do not fully understand, thinking it must be crackpot. Also, I designed the user interface, and it was too sophisticated and not user-friendly enough. I thought it was great for a power user, but opinion is unanimous that it was hard to train novices in.

We have now pivoted and are trying to sell decision-making services to corporations. People are more motivated to learn to use TruthSift for their

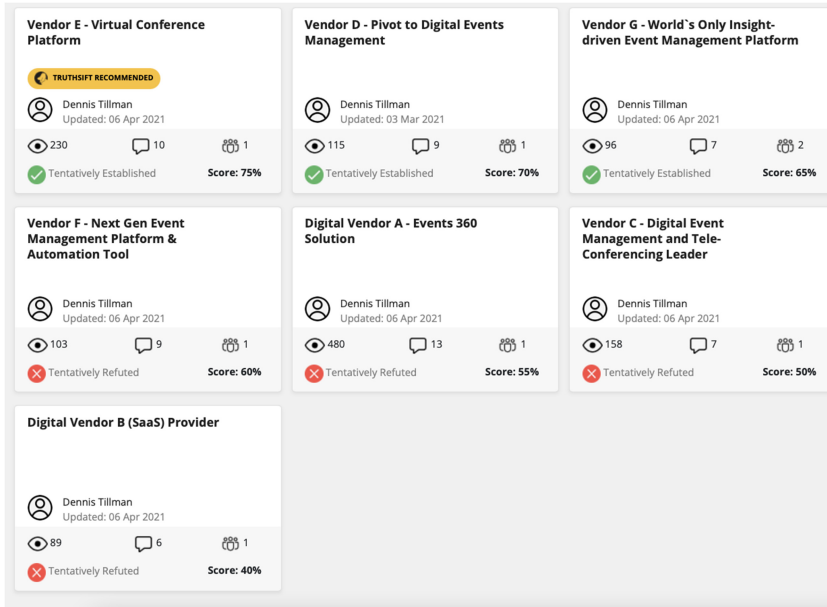


Figure 4: Screenshot of a workspace containing seven graphs, each making a case for a different vendor.

job than as a social network, we can provide guidance, and they are less likely to engage in ad hominem attacks and the like. Corporations need documentation on their decisions, and they need to make better decisions, so there is a lot more motivation. Basically, we are bringing the scientific method to business decision making, so we expect there will be a lot of room for profit in it.

The user interface as presented here has been greatly simplified. For example, we cut back from multiple reply types to just two: support and refute. (We are currently supporting a site where users may freely debate scientific topics at <https://tssciencecollaboration.com>. Users can sign up for a free account.)

For business users, we have added features such as a dashboard with multiple work spaces. One business decision for which TruthSift has been used repeatedly is deciding among alternatives such as applicants for a position, vendors for a technology, or marketing plans (see Figure 4). We support multiple shared workspaces, each of which may be a graph debating, say, whether a particular applicant is qualified for the job. The author of a workspace that contains graphs for each applicant also may designate several features on which to weigh the different candidates. The contributors to the individual graphs provide numerical assessments of the candidate

according to these criteria, in addition to providing pro and con statements to the graphs. TruthSift keeps track of a weighted score for each of the candidates, as well as whether his or her qualifications are tentatively established. This aids the comparison of the qualified candidates, and one may also create a graph establishing that the selected candidate is in fact the right one. Templates are available for various types of business decisions. As a business refines its decision process with experience, these will likely be improved.

2 This Technology Is Important Because We Have Suffered from Its Lack

Peer-reviewed surveys agree: a landslide majority of medical practice is not supported by science (Ezzo et al., 2001; Garrow, 2007; Greenberg, 2009; Office of Technology Assessment, 1978). Within fields like climate science and vaccines that badly desire consensus, no true consensus can be reached because skeptics raise issues that they feel are not adequately addressed by the majority (exactly what Le Bon, 1895, warned of more than 100 years ago). Widely consulted sources like Wikipedia are reported to be largely paid propaganda on many important subjects (Attkisson, 2015) or, at best, the most popular answer rather than an established one (Wilson & Likens, 2015). Quora shows the most popular answer, not necessarily the correct one, and the answers are individual contributions, with little or no collaboration, and often there is little documentation of why anyone should believe them. Existing systems for crowd-wisdom thus largely compound group think rather than addressing it.

Corporate or government planning is no better. Within large organizations, where there is inevitably systemic motivation to not pass bad news up, leadership needs active measures to avoid being kept in the dark as to the real problems (Siitari et al., 2014). Corporate or government plans are subject to groupthink or takeover by employee or other interests competing with the mission. Individuals who perceive mistakes have no recourse capable of rationally persuading the majority and may anyway be discouraged from speaking up by various consequences (LeBon, 1895).

Feynman (1974) famously described the phenomenon of cargo cult science. In the South Seas, the natives saw planes landing on airstrips and delivering cargo, so in an effort to get cargo, they built detailed replicas of landing strips, including wood radios. Feynman observed that the science that results when practitioners do not rigorously discuss and address all the opposing arguments, but instead ignore them or replace them with strawmen, is cargo cult science. Like the islanders' radios, it is missing the key ingredient—the integrity to examine publicly every objection raised and rationally refute it—and so will not function.

We see such examples when the medical literature is assessed. Apparently many researchers have been using the wrong cell lines and know it,

but haven't done anything about it. When a doctor prescribes medicine, it may have been mistakenly tested on, say, a different cancer (Fung, 2014).

In the vaccine literature, apparent systematic biases corrupting almost all their measurements are never discussed (Baum, 2016). The same is true of areas of proven danger such as contaminants, aluminum, timing during development, or multiple vaccine interaction. (Stratton et al., 2002, 2012, provide Institute of Medicine safety surveys). For a graph proving the majority is often wrong, even when they are credentialed and widely admired, see TruthSift (n.d.a.).

3 Private Planning Anecdotes

Users also report that private diagrams have unique potential for intuitive yet rigorous personal and business planning. TruthSift breaks debates down and settles them point by point, propagating the rational consequences of one conclusion on to other decisions. To decide *A*, a user can create a statement, "Evidence for *A*," with all the arguments he or she can think of as proof statements and "evidence for Not *A*" with proof statements of that. Then this person should consider the counterarguments on the specific arguments, followed by what actual proofs and refutations can be offered to the overall conclusions. Everybody can contribute their point of view. Experience indicates that users find the process intuitive and helpful.

4 Previous Related Work

Mathematicians usually write proofs to appeal to other mathematicians, leaving many of the steps somewhat intuitive and incompletely described. However, in principle, every correct mathematical proof should be reducible to a series of syntactic checks. Languages like HOL Light, Coq, PVS, Isabelle, and Mizar have enabled mathematicians to write computer programs checking many major results (Hales, 2008), and there is an ongoing project to verify all of mathematics (Formalizing 100 Theorems, 2016). TruthSift, by contrast, aims to provide formal verification for informal rational discourse such as occurs in the social and physical sciences.

There has been research going back to the Middle Ages on formal systems for persuasion dialog, reviewed by Prakken (2006). These to some extent mirror (but formalize) the intuitive discussion of the mathematical process, as does TruthSift. One way I differ philosophically from some of this work is in hypothesizing that there is an objective reality and maybe also a Platonic reality that can help guide members to discovery and agreement. Another example of modeling mathematical argument was Lakatos games (Pease et al., 2014). TruthSift appears to be the first to implement a representation of such a system on a platform for collaboration and point-by-point proof and refutation, with machine verification of what is established, and to test it on general topics.

Feynman (1968) wrote, "That is what science is: the result of the discovery that it is worthwhile rechecking by new direct experience, and not necessarily trusting the [human] race[']s experience from the past. I see it that way. That is my best definition. . . . Science is the belief in the ignorance of experts." TruthSift directly implements this vision of science, requiring that all objections raised be rationally refuted before a conclusion can be established. What is constant in cargo cult science and other delusions is that the objections raised are simply ignored and usually censored.

There have been a number of previous efforts to achieve crowd intelligence or to map arguments, or both. (For recent reviews, see Michelucci & Dickinson, 2016; Klein, 2015; and Quinn & Bederson, 2011. For a survey of argument maps, see Dwyer et al., 2012.) In Klein's (2015), classification, TruthSift may be considered to have a novel means of aggregation of contributions compared to any surveyed methods, combining contributions according to logic, and a novel means of quality control, using the aggregation to compute what is logically established. Another characteristic used to classify existing systems is the motivation of users to contribute to public collaboration, for which they list the existing alternatives: pay, altruism, enjoyment, reputation, and implicit work (Klein, 2015). In addition to invoking altruism, enjoyment, and reputation, TruthSift also offers the motivation for individuals or organizations (e.g., advertisers or advocates) that want to publish a verified proof of some proposition. It also supports private collaborations within corporations to make better decisions.

In the classification of Michelucci and Dickinson (2016), "The human computation ecosystems of the future" are type C systems that "combine the cognitive processing of many human contributors with machine based computing to build faithful models of the complex, interdependent systems that underlie the world's most challenging problems." However, the existing prototype C systems seem to be for specialized applications and also not entirely automated, like the polymath project (Gowers & Nielsen, 2009).

Notable efforts to provide crowd-sourced question answering or information include Wikipedia and Quora. Wikis suffer on controversial topics, and Wikipedia has been reported to represent paid propaganda on controversial scientific topics (Attkisson, 2015; Wilson & Likens, 2015). Wikipedia reports only the last edit, and this is often apparently controlled by influenced parties. On TruthSift, such edit wars would be transparent, flagged as contrary to the guidelines, and possibly available as specific alternative stipulations. Quora reports the most popular answer to a question rather than any more powerful collaboration, together with a list of alternative answers. Other fact-checking or question-answering websites present some "expert opinion," justified if at all by argumentation that demonstrably often comprises logical fallacies such as straw men and ad hominem attacks.

Klein's (2015) review of online deliberation technologies classifies systems into time-centric, question-centric, topic-centric, debate-centric, and

argument-centric. Time-centric sites like Twitter or comment threads organize content by when it is contributed. Question-centric sites like Quora organize answers by questions. As Klein notes, both of these types of systems produce voluminous output, most of which is of low quality thought there might be some concealed. There is little collaboration beyond voting on one most popular individual contribution.

Topic-centric sites like wikis organize content by topic, producing a more concise output, but they suffer on controversial topics, becoming battlegrounds that are not won on the basis of logic. Debate-centric sites like *whysaurus.com*, *debatepedia.com*, *debatewise.org*, and *debate.org* allow users to contribute pros and cons. However, they have no method of breaking down issues into subpoints and establishing point by point what can be established. They do not allow linking of arguments to arguments, much less automatic update of consequences, and they do not support reasoning forward on open-ended problems, as is possible with TruthSift.

Argument-centric systems like Klein's MIT Deliberatorium (Klein, 2011) (and TruthSift) have advantages. Members prepare a concise and informative diagram summarizing an argument. Solutions are genuinely collaborative, and thus potentially powerful. The idea of determining statements to establish and work forward to establish other statements is integral to science but missing from previous argument map systems like the Deliberatorium. Without that, as with Quora and Wikipedia, there is no good way of composing the right answer, as opposed to merely selecting the most popular answer at each choice. Crowd-based systems mostly compound groupthink, not correct it.

The best existing system for determining truth using natural language argumentation and evidence by a large collaboration has been the scientific literature. It too, however, has suffered under questionable refereeing, the lack of a good mechanism to agree on what is actually established at any given time, and the ignorance or disregard of contrary arguments by authors and referees or editors. It would be much improved if it adopted TruthSift for refereeing (TruthSift, n.d.c).

5 Probability Mode

TruthSift supports probabilistic ratings if switched to probability mode. (See Figure 5.) This allows users to easily collaborate on constructing a Bayes net modeling any question on all supplied evidence and automatically computes the probabilities predicted using a fast Monte Carlo algorithm. You may make a number of observations about the world that are pertinent to whether a given hypothesis is true, such as, "It was reported in the newspaper" or "The newspaper isn't always accurate," or "If that hypothesis was true, some other thing I observed would be much more likely to hold than if it were false." TruthSift offers combining assessments and intuitions

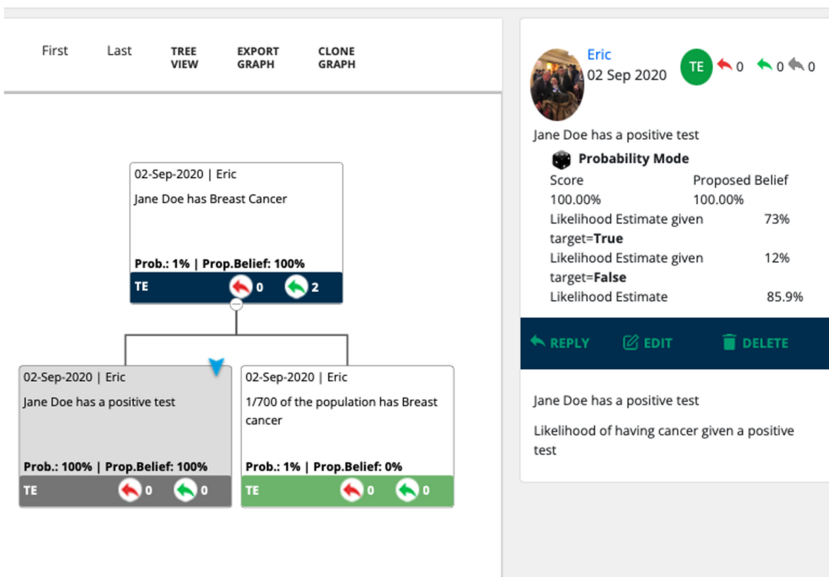


Figure 5: A graph drawn in probability mode. A Bayesian graph may be composed of proof statements whose truth makes their parent more likely, challenges whose truth makes the parent less likely, and test statements for which the author supplies two parameters: the likelihoods given target true and false. TruthSift runs a fast Monte Carlo estimation of the probability of each node. The selected test node is shown in the upper right. The likelihood that Jane actually has breast cancer is reported as the probability of the top statement as 1% after rounding or more precisely if the top node is selected.

about different pieces of evidence into a rigorously calculated probability estimate.

The author of a topic can set it to probability mode by checking the “make this a probability graph” box right below the title on the “add new topic” window. In probability mode, contributors are asked to assign a “proposed belief” to the statements they add, between zero and one, to reflect how much confidence they have that their statement either proves or refutes its target.

In probability mode, in addition to pro and con statements, users can add test statements, which are like epidemiological tests that provide evidence favoring the truth or falsehood of their target. Given some observation claimed in the body of the test statement, the author of the test statement supplies a likelihood of the observation given that the target statement is true and the likelihood of the observation given that the

target statement is false. According to Bayes's law, this evidence multiplies the likelihood of the target statement by the ratio of the first number to the second.

For example, consider the medical test for breast cancer (see Figure 5). It has a false-positive probability of .12, meaning for a woman without the condition. There is a .12 probability the test will show positive anyway, and it has a false-negative probability of .27, meaning that if she does have the condition, there is a .27 the test says she does not and thus a .73 chance it says she does. Thus, a woman who has breast cancer is $.73/.12 = 6$ times as likely to get a positive test as one who doesn't. But since only 1 in 700 women has breast cancer, the likelihood of a woman with a positive test having breast cancer is $6/700 = .00857$.

TruthSift probability mode allows the woman to represent this and calculate the likelihood of having breast cancer, using a test node with a likelihood given true target = .73 and a likelihood given false target of .12, and a prior node giving proof with expected belief $1/700$. (See for this probability graph, TruthSift, n.d.b. This simple example is taken from the discussion in Pearl & Mackenzie, 2018, pp. 104–105.)

TruthSift then calculates the probability that the topic node is true (the woman actually has cancer) by drawing instances from the population conditioned on having a positive test. It does this by sampling instances with a 1 in 700 chance of having cancer and then weighting the women who actually have cancer by the factor .73 (the likelihood estimate given target true) and then weighting ones who do not by the factor .12 (likelihood estimate given target false). Then the likelihood that the topic node is true is given as the ratio of the total weight of positive instances divided by the total weight of examples. This is a Monte Carlo calculation that gives as the expected value the same result as the Bayesian calculation.

TruthSift thus allows building a big causal model with epidemiological tests bringing to bear all of the evidence on a given subject and rapidly estimates the probability of the statements. Those who are unskillful in adding to the graph multiple times different nodes that are highly correlated, basically the reflection of a single measurement and not adding independent information, cannot expect the results to be accurate. However, my experience is that in the hands of a practiced analyst, it is invaluable for intelligence modeling.

6 Discussion and Future Work

TruthSift is attempting to create a collective rationality in a way that has not been tried before. We have had in the past a number of public users on our consumer websites and now have corporations testing for private decision making, but it remains to be seen whether it will recruit a huge base of users or what would result from large-scale public or enterprise use. The expectation is that members will explore rational diagrams as is somewhat

achieved in the scientific literature, but it is possible that too many members will make irrational posts or that it will be hard to agree on what is rational or that people won't get over the learning curve to use it.

TruthSift diagrams are currently demonstrating that there are huge logic holes in things many people believe, so the potential impact of Truth Sift on society is substantial. Moreover, it may demonstrate that there are many similar costly misconceptions and groupthink delusions throughout the corporate, intellectual, and government worlds. A noncontroversial example of a widely held confusion in the business world is given by *Moneyball* (Lewis, 2004). Baseball managers were systematically misvaluing assets worth hundreds of millions of dollars, although Bill James (1982) had been trying to point this out for decades. Without TruthSift, people mostly disregard any evidence or arguments against their cherished positions. TruthSift flags their conclusions until they have come up with a rational argument addressing the issues. After a TruthSift diagram is built, assessing the state of the argument is quick.

A critical reason for the propagation of groupthink delusions is that most people do not take the time to understand the issue but assume others are. But with TruthSift, they do not have to invest much time to verify so can come to rely on TruthSift as a better authority than the group. Thus, a relatively small group of contributors can have a large impact on society.

TruthSift expects to add more features, including n -choice statements (one and only one of the n statements may be established, such as negations that are 2-choice); connectors between diagrams, allowing a big web of verified concepts to be built; and a system allowing rewards to be posted for contributions to topics that are significant because they affect the status of the topic and remain established. This should allow members to motivate researchers to a question and also to make claims that they have demonstrated something in spite of offering a reward for disproof.

TruthSift could be used by AIs rather than or in addition to humans, provided they could sometimes suggest proofs or refutations or tests of statements. TruthSift is currently experimenting with AI-human collaborative posts. The criticism of TruthSift may enable a collection of AIs, perhaps with some humans, to produce a higher-level intelligence than any of them individually could achieve. Even if they sometimes interject challenges or proofs that are off-base, if they more often can identify the problem and challenge it, they may be able to bootstrap to a more rational intelligence. In the meantime, they can help create a wider base of interesting and explored topics.

Readers who are interested in exploring applications of the technology explored in this letter to their own problems or organization, please contact me. I also suggest going to <https://tssciencecollaboration.com/graph/You%20should%20use%20TruthSift%20for%20your%20or%20your%20organizations%20plans%2C%20projects%2C%20documents%2C%20and%20decisions/289/0/-1/-1/0/0#InkNameGraph>.

Conflicts of Interest

I founded/own TruthSift Inc and have patents on the technology.

Coda: Many Worlds Theory: An Empirical Observation and Memoir

When in graduate school at Princeton, I think 1978, I conceived the notion that the many worlds theory of quantum mechanics predicted that I would eventually be the oldest person in the world because there would always be a branch in which I was still alive, and in that branch I would remember making this prediction, and since a 1 in 8 billion event had successfully predicted, I would eventually know in finite time that the many worlds theory was correct (although if it was false I might never gain such confidence).

On October 7, 2016 in the evening I was at a bonfire with my wife and small daughters on Lake Tahoe Beach thrown by the Princeton Club of Northern California. My daughters' fingers got sticky on smores, so my wife took them to the lake to wash them, leaving me talking to somebody who had introduced himself as also a Baum. When they came back a few minutes later, I (and he) were gone. Shortly after she came upon me where I had been found by a passing doctor lying in soft sand with my neck broken and my heart stopped. I remember none of this; my first recollection is from waking up later in an excellent trauma center in Reno.

Since then I have been a vent-dependent quadriplegic. I have a good life still. I have a wife and daughters who love me and I love them. In spite of the fact my heart stopped for some indeterminate period in the dark before I was found, my mind is still sharp. I play far better chess than I ever did (2350 on worldchess.net). I get Queen Bee level on the *New York Times* website "spelling bee" in under 30 minutes almost every morning. I read widely and have had about the best understanding of the bizarre evolution of the world over the last decade of anyone I know. I am fortunate for a fracture of my type in being able to eat anything normal people do and speak well. I have excellent technology to communicate with my computer (Glassouse Bluetooth device I wear like glasses that moves the mouse with my head, bite switch to click, speech recognition), have a mobility van, a great 24/7 nursing staff and a power wheelchair I operate with a mouth joystick.

A few years before my injury I invented TruthSift. It supports the construction and verification of proof refutation trees by individuals or collaborators for any statements you care to debate. This settles issues point by point and propagates the conclusions up. People always have a strong tendency to ignore the arguments against their position and TruthSift prevents this. It makes clear what you actually know, how you know it, and how all arguments against have been refuted. Basically it implements the Platonic ideal of the scientific literature as a truth verification and discovery system and applies it to any situation in life. We are currently marketing it as a

scientific and far superior approach to business decision making, including specializations for various subfields such as hiring.

Living as I do on mechanical ventilation, my life is precarious. I began by lying for an indeterminate time with a broken neck and my heart stopped before presumably somebody walking through the dark on the beach found me, and he happened to be a doctor and revived me. On at least a dozen occasions since then, my airway has become obstructed or disconnected to the point that I passed out. This mostly happens due to the fact that my body regards the trach as a foreign object and produces mucus in an effort to get rid of it. This mucus has to be suctioned out periodically so that I can breathe, it can get dried into plugs that can clog the airway. For example, as I write, this happened this morning. Then the oxygen in my blood drops while a nurse comes running to try and suction the plug out of the way. But it's not always that easy to get the plug out, and a nurse doesn't always respond immediately, so on a dozen occasions, I have passed out before they removed the plug, and on many more come close.

I haven't quite become the oldest man in the world yet, but I have survived a series of threats that seem to be very long odds, so I am basically ready to conclude that the many worlds theory is correct. It's unfortunate my loved ones are mourning me in all the other universes.

Acknowledgments

An earlier version of this letter was presented as "TruthSift: A Platform for Collective Rationality" at the IEEE 2016 Future Technologies Conference. Thanks to Chaitra Keshav-Baum, Amit Sengupta, Vineet Velso, and other team members and contributors to TruthSift for contributions to website design, software, and diagrams, and to Stefi Baum for a critical reading of the manuscript.

References

- Attkisson, S. (2015). Astroturf and manipulation of media messages. TEDx. University of Nevada, <https://www.youtube.com/watch?v=-bYAQ-ZZtEU>
- Baum, E. (June 28, 2016). The top ten reasons I believe vaccine safety is an epic mass delusion. *Atlantic* (April 23).
- Bernays, E. (1928). *Propaganda*. Ig Publishing.
- Dwyer, C. P., Hogan, M. J. & Stewart, I. (2012). An evaluation of argument mapping as a method of enhancing critical thinking performance in e-learning environments. *Metacognition Learning*, 7, 219–244. 10.1007/s11409-012-9092-1
- Ezzo, J., Bausell, B., Moerman, D. E., Berman, D., & Hadhazy, V. (2001). Reviewing the reviews. *International Journal of Technology Assessment in Health Care*, 17, 457–466. <http://journals.cambridge.org/action/displayAbstract?fromPage=online&aid=101041>. 10.1017/S0266462301107014
- Feynman, R. P. (1968). What is science? *Physics Teacher*, 7(6), 31–320. http://www-oc.chemie.uni-regensburg.de/diaz/img_diaz/feynman_what-is-science_68.pdf

- Feynman, R. P. (1974). *Cargo cult science*. Commencement address at CalTech. <https://sites.cs.ucsb.edu/~ravenben/cargocult.html> <http://discovermagazine.com/2014/nov/20-trial-and-error> <https://lifeboat.com/blog/2016/06/the-top-ten-reasons-i-believe-vaccine-safety-is-an-epic-mass-delusion>
- Formalizing 100 theorems. (2016). <http://www.cs.ru.nl/~freek/100/>
- Fung, B. (2014). "Oops, wrong cancer": How contaminated cell lines produce bad research. *Atlantic* (April 23).
- Garrow, J. S. (2007). What to do about CAM? How much of orthodox medicine is evidence based? *BMJ*, 335(7627), 951. <http://www.dcsociety.net/garrow-evidence-bmj.pdf>. 10.1136/bmj.39388.393970.1F
- Gowers, T., & Nielsen, M. (2009). Massively collaborative mathematics. *Nature*, 461(879), 879–881. 10.1038/461879a
- Greenberg, S. A. (2009). How citation distortions create unfounded authority: Analysis of a citation network. *BMJ*, 339, b2680. 10.1136/bmj.b2680
- Hales, T. P. (2008). Formal proof. *Notices of the A.M.S.* 55 (11), 1355–1380.
- James, B. (1982). *The Bill James baseball abstract*. Ballantine Books.
- Klein, M. (2011). The MIT Deliberatorium: Enabling large-scale deliberation about complex systemic problems. *International Conference on Agents and Artificial Intelligence*. Rome, Italy.
- Klein, M. (2015). *A critical review of crowd-scale online deliberation technologies*. SSRN. <http://ssrn.com/abstract=2652888>
- Le Bon, G. ([1895] 1995). *The crowd*. Transaction.
- Lewis, M. (2004). *Moneyball: The art of winning an unfair game*. Norton.
- Michelucci, J. L., & Dickinson, J. L. (2016). Human computation: The power of crowds. *Science* 351 (6268), 32–33. 10.1126/science.aad6499
- Office of Technology Assessment, U.S. Congress. (1978). *Assessing the efficacy and safety of medical technologies*. <http://www.fas.org/ota/reports/7805.pdf>
- Pearl, J., & Mackenzie, D. (2018). *The book of why: The new science of cause and effect*. Basic Books.
- Pease, A., Budzynska, B., Lawrence, J., & Reed, C. (2014). Lakatos games for mathematical argument. In S. Parsons, N. Oren, C. Reed, & F. Cerutti (Eds.), *Proceedings of the Fifth International Conference on Computational Models*. <http://comma2014.arg.dundee.ac.uk/res/pdfs/07-pease.pdf>
- Prakken, H. (2006). Formal systems for persuasion dialogue. *Knowledge Engineering Review Archive*, 21(2), 163–188. <http://www.cs.uu.nl/groups/IS/archive/henry/dgreview.pdf>. 10.1017/S0269888906000865
- Quinn, J., & Bederson, B. B. (2011). Human computation: A survey and taxonomy of a growing field. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1403–1412). 10.1145/1978942.1979148
- Siitari, K., Martin, J., & Taylor, W. W. (2014). Information flow in fisheries management: Systemic distortion within agency hierarchies. *Fisheries*, 39(6), 246–250. 10.1080/03632415.2014.915814
- Stratton, K., Ford, A., Rusch, E., & Clayton, E. W. (Eds.). (2012). *Adverse effects of vaccines: Evidence and causality*. National Academies Press.
- Stratton, K., Wilson, C. B., & McCormick, M. C. (Eds.). (2002). *Immunization safety review: Multiple immunizations and immune dysfunction*. National Academies Press.

- TruthSift. (n.d.a). *On controversial topics who is more often right, the majority or a minority?* [diagram]. <https://tssciencecollaboration.com/graph/On%20Controversial%20Topics%2C%20Who%20is%20More%20Often%20Right%2C%20the%20Majority%20or%20a%20Minority/466/0/-1/-1/0/0#lnkNameGraph>
- TruthSift. (n.d.b). *Does Jane Doe have breast cancer given a positive test?* [diagram]. <https://tssciencecollaboration.com/graph/Does%20Jane%20Doe%20have%20cancer%20%20given%20a%20positive%20test/964/0/-1/-1/0/0#lnkNameGraph>
- TruthSift. (n.d.c). *It would be far better to replace journal refereeing with TruthSift.* [diagram]. <https://tssciencecollaboration.com/graph/It%20would%20be%20far%20better%20to%20replace%20journal%20refereeing%20with%20TruthSift/819/0/-1/-1/0/0#lnkNameGraph>
- Wilson, A. M., & Likens, G. E. (2015). Content volatility of scientific topics in Wikipedia: A cautionary tale. *PLOS One*, 10(8). 10.1371/journal.pone.0134454

Received April 28, 2022; accepted July 21, 2022.