

Amplifying the Wisdom of the Crowd, Building and Measuring for Expert and Moral Consensus

James L. Carroll¹, and Brent Allsop²

¹Physics Division, Los Alamos National Laboratory. jlcarroll@lanl.gov

²Canonizer.com. brent.allsop@canonizer.com

Abstract—*Scientific and cultural progress mostly takes place in the marketplace of ideas. The internet has largely changed how this marketplace functions. Although it has brought many improvements, it has also created a new set of problems. We propose a web service that can be used to crowdsource the process of knowledge acquisition and of knowledge summary, mediate disagreements, and improve discourse on controversial issues, thereby improving the efficiency of the marketplace of ideas. Our approach involves a specific technique for integrating: a forum; a wiki organized into separate camps, with camps organized into a hierarchical structure; a survey system based upon camp support; a mechanism for dynamically reorganizing the structure of the camp hierarchy while mediating user disagreements; and a customizable mechanism for determining how votes are weighted based on credentials or expert assignment. A version of this approach has been implemented, released as open source, and a beta test has gone live. Initial use seems to validate the merits of this approach.*

Keywords: crowdsourcing, consensus, information summary, wiki, voting

1. Introduction

Scientific progress is made as ideas compete for acceptance in the marketplace of ideas. The internet has increased access to, and participation in, this marketplace. This has been largely beneficial, but it has also introduced several significant problems. Currently, there are many inefficiencies in the marketplace of ideas, which has led to information overload, increased confirmation bias, discourteous discourse, and an increase in conspiracy theories.

Many web based approaches have been introduced in recent years to alleviate some of these concerns by promoting communication, information ranking, and information summary. Many employ crowdsourcing to leverage the wisdom of the crowd. However, each approach has its own drawbacks. Here, we propose a new hybrid approach to a crowdsourced, web based, knowledge acquisition service.

This work was partially supported by the Advanced Radiography Science Campaign at Los Alamos National Laboratory. LA-UR-12-21202.

Our proposed solution involves a specific integration of: a wiki organized into separate camps, with camps organized into a hierarchical structure for information summary; an email list / forum for communication integrated into the hierarchical structure; a survey system based upon camp support; a mechanism for dynamically reorganizing the structure of the camp hierarchy while mediating user disagreements; and a customizable mechanism for determining how votes are weighted based on credentials or expert assignment. When integrated and organized correctly, these elements function together as a coherent whole, which we believe has the potential to solve many of the inefficiencies in the marketplace of ideas. This approach can significantly improve discourse, especially surrounding controversial scientific or moral subjects like global warming, evolution, sex education, abortion, consciousness, or religion.

The remainder of the paper is organized as follows: In Section 2, we discuss the importance of the marketplace of ideas in scientific epistemology. In Section 3 we describe several of the most serious inefficiencies in the current marketplace of ideas, especially ones related to web based discourse. In Section 4 we overview several of society's current approaches for dealing with these problems, including web based services designed for knowledge collection and discussion. We will discuss the benefits and drawbacks of each approach. In Section 5, we give the details of our proposed knowledge synthesis system, known as "Canonizer." In Section 6, we describe several illustrative case studies drawn from an initial live beta test of the system. We believe that these case studies will show that our approach has many benefits for producing more effective discourse on controversial issues, including the ability to illuminate areas of agreement, and focus further research on the remaining unanswered questions. In Section 7, we conclude and propose further work.

2. The Marketplace of Ideas

Most of the earliest philosophers saw the study of logic as the quest to discover a mechanism whereby all that is true could be reliably determined, while all that is false could be reliably rejected. With the advent of modern science, a desire for objectivity and empiricism was added to the process. For many years there was great hope that the scientific

method itself provided the very mechanism sought by the early philosophers. However, since Hume's destruction of empiricism[1], much of this optimism has come to an end. Today we realize that observations are necessarily statements of what happened in a few places at a few times, while scientific theories are statements about what will happen in all places every time.[2] Making this leap from specific instances to general statements about reality requires generalization, which Mitchell[3] and Wolpert[4][5][6][7][8] have soundly demonstrated is impossible to do in a completely objective and unbiased manner.[9] Furthermore, from the various impossibility and in-computability theorems, we have also learned that there may be many things that are true, but which can not be formally proven.

These realizations do not mean that the scientific method is useless, or that we can never learn or know anything about the world in which we live. Modern approaches view the scientific method more pragmatically, noting that it reliably produces theories that are useful. Useful theories are the ones that make the most accurate predictions, in the widest set of circumstances, with the least effort.[2] These memetic theories then compete for general acceptance in a "marketplace of ideas." [10] The most useful theories survive, replicate, and eventually achieve general acceptance and consensus. It is often presumed that other human beliefs and practices such as religion, philosophy, ethics, and governmental institutions all compete for acceptance in a similar manner.[11]

An alternative philosophy of the epistemology of science that is popular among those that want to avoid an appeal to "usefulness" and to the market place of ideas, involves an appeal to statistical inference. According to this philosophy, science may not be able to prove that any general hypothesis is correct, but it can select among them by computing the probability of each hypothesis given the available data.[12] This approach can still be pragmatic in that utility theory can be used to select the action with the highest expected utility (the one that is most likely to "work" the best) given the uncertainties over hypotheses and the utilities of different outcomes.

Although this approach is very promising, it does not completely restore the objectivity of science,[13] since any deductive application of the laws of statistics to an inference problem requires the use of Bayes Law,[9] which states that:

$$P(\textit{Belief}|\textit{Data}) = \frac{P(\textit{Data}|\textit{Belief})P(\textit{Belief})}{P(\textit{Data})}$$

$P(\textit{Belief})$ is known as a *prior*, and is often subjective in nature. Unfortunately, it is possible to believe anything regardless of any amount of data if you choose the right pathological prior.[14] And the "right" way to choose priors is far from clear.[12] Nevertheless, for most reasonably chosen priors, different initial beliefs will converge to a common answer given sufficient data. It is also possible to use the prior to overcome most of the common problems of

generalization by attempting to choose a prior that embodies the effective bias of Occam's razor.[9] This approach can serve as a mechanical technique for arriving at both good practical hypotheses and for arriving at eventual consensus given sufficient data, much as the ancient philosophers hoped.

Unfortunately, it is not possible to be personally aware of all available data. Nor is it possible to be absolutely sure that any given data set has not been faked or otherwise altered. Nor is it possible for any one person to become enough of an expert to effectively apply this technique to the full range of subjects we regularly encounter in modern society. Therefore, in order to determine what we should believe on a wide range of topics, we must eventually place our trust in witnesses and experts who have experienced and evaluated the data on our behalf, and who have come to reasonable conclusions on our behalf. Practically, we can not evaluate all the primary data directly, but must instead treat the witnesses of various experts as our data instead. Because supposed experts in different fields often disagree, we must determine which experts we trust, and which we will ignore. This realization inevitably leads us back to some variant on the marketplace of ideas, where ideas and hypotheses compete for general acceptance.

It would appear that there currently is no way for the quest for truth to avoid the marketplace of ideas completely.

3. Inefficiencies in the Marketplace of Ideas

There are many potential inefficiencies in the marketplace of ideas, but we will focus on the following central issues, especially as they relate to the internet: 1) the limited access to information, 2) the un-manageable volume of information, 3) the reliability of information, 4) the need to determine who to consider an expert, 5) the need to determine what the experts believe, 6) the difficulty in changing popular notions that are resistant to change, 7) the incivility, repetition, and lack of focus in discourse, and 8) the difficulty in overcoming confirmation bias. By exploring these inefficiencies, we can begin to envision what their solution might look like.

1) The limited access to information has been the largest inefficiency in the marketplace for the majority of human history. Major milestones in the sharing and spread of information include the invention of language, the invention of writing, followed by the invention of the printing press. It was only at this point that it became possible for the average person to have access to any reasonable percent of the world's information. The rise of the internet and the world wide web represent yet another significant paradigm shift in this regard. Now an even greater percent of the world's information is even more widely accessible, and at the simple touch of a button. However, the internet's improvements to the availability of information have come

at a cost.

2) The un-manageable volume of information is the most obvious problem with the data available on the internet. Just as the printing press put books into the hands of the masses, so the internet has put the printing press into the hands of the masses, and the volume of information naturally grew accordingly. This presents the challenge of determining which information is most important. Today, we are drowning in information, but starved for knowledge. Furthermore, as human knowledge becomes ever more specialized, it is increasingly difficult for any one person to be adequately informed on all subjects. Even the smaller body of peer reviewed literature is growing far faster than it can be read, even by experts in their respective fields.

3) The reliability of information is increasingly difficult to determine. Before, there were editors and publishers who acted as gatekeepers to provide some limited measure of control over what was published. Now anyone can publish anything. The challenge is to differentiate between accurate information and hoaxes. It is tempting to judge ideas based upon their intuitive plausibility, or based upon how well they match with an apparent consensus. However, brilliant revolutionary ideas will run counter to current consensus almost by definition, and to the layman, they are often indistinguishable from crazy useless ideas. Traditionally, expert opinion has served to determine whether information can be trusted. But this leads to an additional challenge.

4) The need to determine who to consider an expert must be addressed if we are to trust their opinions. Unfortunately, the internet allows anyone to claim to be an expert, even if they are not. We have replaced the challenge of determining which information to believe with the new, equally difficult, challenge of determining which experts to believe. But even if we can determine who to consider an expert on each topic, we are left with an additional challenge.

5) The need to determine what the experts believe must be addressed if the expert's opinions are to be useful. Different populations, demographics, and fields of experts often have wildly differing opinions on controversial issues, so there is no one single answer. For example, the majority of climatologists believe that anthropomorphic global warming is real and significant, while this conclusion is far less certain among meteorologists, other scientists, and the general public. Furthermore, the peer review and tenure process rewards primary research more than public education, so there is insufficient incentive for experts to help the general populace understand what they believe and why.

6) The difficulty in changing popular notions that are resistant to change must be addressed for the marketplace to converge towards truth. Take for example the oft refuted adage that "we only use 10% of our brains," but which most people still believe. Sometimes these theories persist through confirmation bias, but often they simply persist because the general populace is unaware of the contrary consensus of

the experts, which is why determining and presenting what the experts believe is so important. Other times, they persist because the experts fail to communicate effectively to the masses, and are unaware of what information the masses might find convincing.

7) The incivility, repetition, and lack of focus in discourse is another major inefficiency in the market. The internet provides a mechanism for easy debate, where controversial ideas are constantly argued, but often with little success. The more controversial a topic is, the more likely debate upon that topic will digress into repetition and incivility, while producing no useful movement towards consensus. Godwin's law is the adage that any discussion in any open forum will eventually descend to the level where someone gets compared to Hitler or the Nazis. Unfortunately, this adage often proves accurate, even in scientific circles.

Even when outright incivility is avoided, most of these debates suffer from a lack of proper focus. They are often repetitive, dealing with issues that have long since already been resolved. Alternatively, they can focus on areas of disagreement, without discovering potential areas of agreement.

8) The difficulty in overcoming confirmation bias can be increased by internet interaction. Confirmation bias is our well known psychological tendency to seek information that confirms our pre-existing beliefs instead of seeking out information that challenges these beliefs. Before, people often had to share their ideas with those that were geographically close to themselves. This forced people to interact with others with a certain amount of diversity of opinions. However, in the internet age, anyone can seek out and find a community of people who believe whatever they themselves do, no matter how impossible or unlikely their beliefs may be. By allowing us to find a community to interact with that mostly shares our own views, access to the vast amount of information on the internet may actually be isolating us from opinions that differ from ours. This problem is becoming a bigger, more widely recognized, concern.[15]

This not only causes us to reinforce our own false beliefs, it can also prevent us from being an effective advocate, even for those of our belief that are correct. It is clear that there is a strong consensus among those that are best at effective communication and conflict mediation that one of the most important elements is to seek first to understand, and only then seek to be understood.[16][17][18] It appears that it is only after we understand another's position well enough to defend it better than they could, that we are best prepared to change their minds. Furthermore, is often the case that once we understand someone else's position that well, we realize that they were not as wrong as we once thought.

These problems permeate much of our society, including scientific, religious, political, and philosophical debate. Many of society's problems, including economic stagnation, poverty, war, and religious extremism, can be traced in one way or another to breakdowns of one form or another in the

marketplace of ideas. This makes the search for solutions all the more urgent. We are in desperate need of an “immune system for our memes”[19] that can directly address the above issues.

4. Current Mechanisms to Address Inefficiencies

Society has recognized many of the above problems, and has evolved several knowledge synthesis systems which attempt to address these concerns, including: 1) search engines, 2) peer review 3) wikis, 4) forums, bulletin boards, and email lists, and 5) survey systems. Evaluating the benefits and drawbacks of each of these solutions can motivate the creation of a new system that can leverage their existing benefits, while removing some of their weaknesses.

1) Search engines have vastly improved our ability to find the information we want amid the sea of information available. However, it is unclear whether we should delegate our determination of what matters to some search engine’s page rank algorithm. Modern page rank algorithms are often customized to give different information to individual users based on what the algorithm thinks that the user wants. This can further exacerbate the confirmation bias problem. Furthermore, being able to find the right primary information does not in and of itself provide the concise summaries of information that we need. The search approach largely succeeds because others have already provided these summaries using other tools.

2) Peer review is the most important information evaluation mechanism in the scientific community, and it is best at producing and evaluating primary research, but is poor at summarizing the *huge* volume of existing work. It serves as a primary filter, designed to differentiate good ideas from bad, by only allowing the best ideas to pass the review process. But this approach is not perfect. Whether a paper is accepted can often depend on luck with regard to who is chosen to review the paper. Also, the process is often slow, and the results are often not publicly available.[20]

In the past, there has been no good measure of consensus or impact. This situation has improved recently as web services have been created that allow researchers to determine the number of times a given paper has been cited. Unfortunately, this still does not measure the consensus opinion of the scientific community, since papers are often cited in refutations. In many fields, some of the most cited papers are the most controversial ones, not necessarily the most widely accepted ones.

3) Wikis provide a nearly opposite approach. Instead of intentionally limiting what is published like peer review does, this approach encourages everyone to participate, and their edits directly compete in a marketplace of edits. Wikis are best at leveraging the power of crowdsourcing to efficiently produce surprisingly accurate summaries of vast

bodies of information[21], but they are less effective at publishing and producing primary research. They can combat the volume problem, and the summary problem, while helping us discover which papers, researchers, experts, and ideas are most important. However, this approach suffers from a lack of trust. Their mechanism for producing trust, namely competition in the marketplace of edits, does not guarantee that any instance of a page at any one time is accurate. Most wiki systems compile each topic’s information into a central page that is supposed to represent the single aggregated and summarized truth. For controversial topics, this encourages edit wars. Although wikis such as Wikipedia produce good summaries of information, they often fail to determine expert consensus in controversial situations (even when it exists) because any attempt to do this violates Wikipedia’s neutrality point of view requirements.

4) Forums, bulletin boards, and email lists are powerful tools for communication, discussion and debate. But they suffer from the problems of incivility, repetition, and lack of focus that we discussed earlier. They do not provide any mechanism for determining which participants should be considered an expert, nor do they help us to determine what the experts believe. Information from these sources is often highly un-reliable.

5) survey systems are currently societies’ primary method for determining what different groups believe. Problems involve determining which demographic to survey, and how to acquire a representative sample. Web surveys can be especially inaccurate, because they often target single demographics. Many people do not care what the general populace thinks, but only what the experts in a given field think. But there is no standard mechanism for determining who the experts are in any given field in order to target surveys just to them, nor is there any way to adjust vote weights accordingly. It is always possible to find some people who will believe anything, and this is why projects like the “Global Warming Petition Project”[22] have limited value. Furthermore, surveys represent a static snapshot of a subset of opinions at one time. While this can be extraordinarily useful, they do not naturally track the shifting opinions that is the hallmark of progress in human knowledge, nor do they track the *reasons* behind such changes. Yet determining what data was convincing enough to cause a given researcher to change their minds is essential information.

These partial solutions seem to point the way forward towards a system that can use these approaches in concert to leverage their collective advantages, while minimizing their individual drawbacks. The eventual solution should be open, should refer to peer reviewed literature, but should use crowdsourcing to summarize and prioritize these results while avoiding edit wars, should allow communication and discussion like that found in forums, should determine who the experts in different fields are, should dynamically (and automatically) determine what different groups believe,

and should provide a system that encourages consensus, information summary, and progress. And it should integrate these pieces into a coherent intuitive whole.

5. The Canonizer Solution

No one solution can solve all of the inefficiencies in the marketplace of ideas, but we believe that some progress can be made. Our proposed solution involves the tight integration of several different complementary systems. The method of integration and organization is what allows each of these otherwise desperate solutions to function together as an effective whole. Our proposed solution involves:

- A wiki to leverage the wisdom of the crowd for knowledge collection and summary
 - Wiki pages are organized into camps with differing opinions, instead of single pages for each topic
 - Camps are organized into a hierarchical structure to encourage agreement
- A mechanism for dynamically reorganizing the structure of the camp hierarchy while mediating user disagreements to encourage agreement and highlight areas of continuing disagreement
- A survey system based upon camp support for determining consensus, and
- A mechanism for customizing how votes are weighted based on credentials or expert assignment to allow users to explore different demographic effects, or to allow users to give extra weight to experts in different fields if they so choose.
- A forum which also functions as an email list for discussion

Wiki Camp Pages: Our wiki system avoids the danger of edit wars by organizing information into a hierarchical set of “camps” for each controversial topic. There is less incentive to edit the page of a camp you disagree with directly. Instead, you can edit your own camp’s page to make an argument for why you believe that your opponent’s camp is incorrect. Instead of trying to remove an argument that you find unconvincing from another camp’s page, you can instead reference the argument, and explain why it is not convincing in your own camp’s page.

Camp pages also serve as an opportunity to summarize the relevant peer reviewed literature. Thus, the system is not intended to replace peer review, rather it is intended as a mechanism for leveraging the wisdom of the crowd to effectively compile, summarize, and prioritize the vast peer reviewed literature.

Hierarchical Structure: Camps are arranged into a hierarchical structure, with the base camp representing the question to be answered, discussed, debated, addressed, or summarized. The hierarchical structure pushes areas of disagreements towards the leaf nodes, and pushes areas of agreement towards the root (see Figures 1 and 2). Suggested

changes to camp structure are sent out on the email list, and only go live after a weighting period if no camp member objects within that time frame.

Camp Support: Both the problems of determining who the experts in a field are, and determining what different demographics believe is solved through camp support, which functions as a sort of survey system. Supporting a camp (voting for the camp), also supports all the parent camps up the hierarchy. This allows us to assess the sort of consensus that currently exists among the (potentially biased) subset of users interested in an individual question or problem enough to join a camp. Users can also create other camp hierarchies which provide surveys to determine who the experts are in various fields. Users can then select their own voting schemes to determine how votes are weighted, potentially based upon the results of these expert determining polls. This flexibility allows users of the system to ask questions like “what percent of general users believe in global warming?” “what percent of users with PhD’s believe in global warming?” “what percent of users with PhD’s in climatology believe in global warming?” “what percent of users that have been selected as experts in climatology believe in global warming, weighted by how many people consider them to be an expert?” etc. This flexibility in how users support is weighted is part of what makes the survey system so useful.

Forum and Email List: The forum also functions as a mailing list, and is intimately integrated with the camp hierarchy, as messages are sent to the supporters of camps. Joining a camp joins the camp’s mailing list. This makes it possible to precisely target your audience. Because the forum is integrated with the camp hierarchy, the camp pages quickly begin to look like condensed summaries of the best arguments that were made in the forums. This makes forum discussions less likely to be repetitive and confrontational. With the related summaries in mind, discussion can now revolve around new issues, not yet covered in the summaries. If a topic that has already been covered resurfaces, people generally point the individual to the relevant summary pages. It can also reduce confirmation bias, because the most effective camp summaries grow out of debates with competing camps, and require first reading your opponent camp’s summaries in order to write an effective rebuttal.

6. Case Studies

“<http://canonizer.com>” is a live beta test for our proposed system which is being used to find bugs and iron out user interface issues. The user base in the beta test is still in its infancy, so many scaling questions remain. However, some observations can still be made about this approach’s effectiveness as a tool for generating discussion and consensus among this small group of initial users.

The majority of Canonizer’s current users are philosophers interested in the complex, difficult, religiously and ideo-

logically charged problem of consciousness and subjective experience. This initial emphasis is largely due to the philosophy background of one of Canonizer’s initial programmers. Several world class leaders in this field have participated to varying degrees, including Steven Lehar, Stuart Hameroff, and a growing number of others. Obviously, world class participation has significantly contributed to the education and development of everyone.

We believe that Canonizer’s use in this field illustrates Canonizer’s ability to encourage the avoidance of confirmation bias, summarize opinions found in an increasingly large body of literature, and demonstrate the usefulness of its expert weighted voting system.

In the study of consciousness, there are a huge number of published peer reviewed articles. One online bibliography recently passed 20,000 published papers on the “Philosophy of Mind” alone.[23] Thus, the field is in desperate need of concise summaries of the literature which we are beginning to provide.

It is also difficult to determine from this vast body of literature whether there is any emerging consensus. David Chalmers et al. did a traditional philosophical survey in 2009[24] after he first heard about our survey project, which seemed to show a majority leaning away from behavioralism.¹ But this survey only targeted philosophers. The opposite feeling is generally accepted in the neuroscience and computer science communities. This divide has been illustrated by the authors of this papers’s own extensive debates upon the subject, where Allsop has a philosophy background and Carroll has a computer science background.

Given the initial philosophical bias of its users, theories of consciousness related to David Chalmer’s “property dualism” theory seemed to dominate, agreeing with the traditional survey. Most remaining disagreement was between David Chalmer’s Functional Property Dualism vs Stuart Hameroff’s Material Property Dualism. However, Canonizer’s demographic system made it clear that most users were biased towards philosophy. With the later introduction of more computer scientists and neuroscientists, it was possible to see that this dominant support for property dualism was not shared by this new community of users. Canonizer’s system of determining experts in different fields makes it easy to see this difference of opinion between fields by alternating the expert voting system from “mind experts” (the philosophy group), to “computer experts” while watching how the apparent “consensus” changes.

Canonizer has also vastly improved forum discourse, as participants constantly pointed at existing arguments and summaries in the camp pages. Canonizer also reduced confirmation bias, as it forced the computer scientists to read and respond to the philosopher’s “property dualism” camp pages, while forcing the philosophers to read and respond

¹See the question related to philosophical zombies.




Fig. 1: Initial camps for the global warming question.




Fig. 2: Later camps for the global warming question, growing consensus.

to the computer scientist’s “functionalist” camp pages.

This project is still far from comprehensive, only including a small handful of real experts. But already, Canonizer contains a more concise set of descriptions of the best theories than is available from many other locations. And things continue to progress as ever more people participate.

Another example involves the highly charged debate about global warming, which demonstrates how the hierarchical structure encourages the search for common ground. Initially users created two camps, one that advocated the idea that global warming represents a clear an immediate danger to the planet, while the other advocated the idea that global warming was a complete fabrication and a hoax. At first, there appeared to be no common ground (see Figure 1). However, as a greater variety of people joined the discussion, it became clear that some people who belonged to the hoax camp actually believed that the earth was warming, but for natural reasons. Others thought that the earth was warming for human generated reasons, but that the severity of the situation was being exaggerated. Others thought that the situation may be severe, but not sufficiently severe to warrant a cap and trade solution, where the proposed cure could be economically worse than the disease. A proposed change to the hierarchy (see Figure 2) created a camp that embodied some of the common ground discovered in the discussions, namely that anthropomorphic global warming was real, with critical and not critical sub camps. There now appears to be great consensus that anthropomorphic global warming is indeed real. This moves the debate forward, and allows us to recognize and to focus on the point of greatest remaining uncertainty, namely the severity of the problem. Similar consensus building results were found in the controversial sex education question.

7. Conclusions

Scientific and cultural progress depends upon ideas efficiently competing in the marketplace of ideas. But the marketplace has several significant inefficiencies that we work to address. Our approach involves a specific technique for combining several different elements into one functioning whole, designed to improve discourse and understanding. We have implemented a first draft of this approach, and released the code under an open license. We have also allowed the approach to be tested in a beta test environment, where it has already been beneficial, and even attracted several important experts in their respective fields.

7.1 Future Work

Future work will involve simplifying the process whereby users can self identify their educational level, and primary and secondary fields of interest. Further improvements to the look and feel, and user interface have also been suggested, and will be explored. Of most importance will be evaluating the performance of the system as the user base grows. For example, it remains to be seen how the system and growing community will be able to respond to trolls, fake accounts intended to manufacture false camp support and other such distractions that will inevitably surface.

References

- [1] Hume, *Hume's Treatise of Human Nature*. New York: Stanford University Press, 1739-1740.
- [2] K. R. Popper, *The Logic of Scientific Discovery*. London: Routledge, 1959.
- [3] T. M. Mitchell, "The need for biases in learning generalizations," Rutgers Computer Science, New Brunswick, New Jersey, Tech. Rep. CBM-TR-117, 1980. [Online]. Available: citeseer.ist.psu.edu/mitchell80need.html
- [4] D. H. Wolpert and W. G. Macready, "No free lunch theorems for search," Santa Fe Institute, Tech. Rep. SFI-TR-05-010, 1995.
- [5] D. H. Wolpert, Ed., *The Mathematics of Generalization*. The Addison-Wesley Publishing Company, 1995.
- [6] D. H. Wolpert, "The lack of a priori distinctions between learning algorithms," *Neural Computation*, vol. 8, pp. 1341-1390, 1996.
- [7] D. H. Wolpert and W. G. Macready, "No free lunch theorems for optimization," *IEEE Transactions on Evolutionary Computation*, vol. 1, no. 1, pp. 67-82, April 1997.
- [8] D. H. Wolpert, "The supervised learning no-free-lunch theorems," NASA Ames Research Center, Tech. Rep. 269-1, 2001.
- [9] J. L. Carroll, "A bayesian decision theoretical approach to supervised learning, selective sampling, and empirical function optimization," Ph.D. dissertation, Brigham Young University, March 2010. [Online]. Available: <http://james.jlcarroll.net/publications/>
- [10] R. Dawkins, *The Selfish Gene*. Oxford University Press, 1989.
- [11] R. Hofstadter, *Social Darwinism in American Thought*. Boston: Beacon Press, 1944.
- [12] E. T. Jaynes, *Probability Theory: The Logic of Science*. Unighted Kingdom: University of Cambridge, 2003.
- [13] J. O. Berger and D. A. Berry, "Statistical analysis and the illusion of objectivity," *American Scientist*, vol. 76, 1988.
- [14] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian Data Analysis, Second Edition*. Boca Raton and London and New York and Washington, D.C.: Chapman Ad Hall, Texts in statistical Science Series, 2004.
- [15] I. Poon, "Is the internet lying to us?" <http://www.physorg.com/news/2010-11-internet-lying.html>, November 2010.
- [16] S. R. Covey, *7 Habits of Highly Effective People*. Free Press, 2004.
- [17] A. Institute, *The Anatomy of Peace, Resolving the Heart of Conflict*. San Francisco: Berrett-Koehler Publishers. Inc, 2006.
- [18] M. Marquardt, *Leading with Questions*. Jossey-Bass, 2005.
- [19] K. E. Drexler, *Engines of Creation, the Coming Era of Nanotechnology*. Anchor Books, 1986.
- [20] M. Eisen, "How traditional publishing hurts scientific progress," <http://www.wired.com/wiredscience/2012/03/michael-eisen-open-science/>, March 2012.
- [21] J. Howe, "The rise of crowdsourcing," *Wired Magazine*, vol. 14.06, June 2006. [Online]. Available: http://www.wired.com/wired/archive/14.06/crowds_pr.html
- [22] "Global warming petition project," <http://petitionproject.org/>.
- [23] "Philpapers online research in philosophy," <http://philpapers.org/>, 2012.
- [24] "The philpapers surveys," <http://philpapers.org/surveys/>, 2009.