WikiSilo:
A Self-organizing, Crowd Sourcing System for Interdisciplinary Science

David Pierre Leibovitz (dpleibovitz@ieee.org),
Robert L. West (robert_west@carleton.ca) &
Mike Belanger (mike.belanger@carleton.ca)
Institute of Cognitive Science, 2201 DT, Carleton University
1125 Colonel By Drive, Ottawa, Ontario, K1S 5B6 Canada

Abstract

WikiSilo is a tool for theorizing across an interdisciplinary field such as Cognitive Science, and provides a vocabulary for talking about the problems of doing so. It can be used to demonstrate that a particular cognitive theory is complete and coherent at multiple levels of discourse, and commensurable with and relevant to a wider domain of cognition. WikiSilo is also a minimalist theory and methodology for effectively doing science. WikiSilo is simultaneously similar to and distinct, as well as integrated and separated from Wikipedia™. This paper will introduce the advantages of WikiSilo for use in the Cognitive Sciences.

Keywords: Crowd-sourcing; Interdisciplinary; Philosophy of Science; Self-organizing; Unification; Wikipedia.

Introduction

Cognitive Science is a challenging learning enterprise spanning numerous disciplines and populated by a multi-agent network of academics that communicate using conceptions of varying meanings encapsulating differing implicit assumptions and based on conflicting approximations as needed for lines of inquiry going in diverse directions. This produces a fracturing of activity into ever narrower and detailed silos encompassing different terminology and approaches to discovery. However, the goal of Cognitive Science requires being informed by all these diverse silos. It requires a form of integration and synthesis that will lead to increasingly useful and theoretically unified theories.

Science proceeds on several fronts (Millikan, 1923). Following highly specialized paths of inquiry is required, and silos facilitate this. The sea of possibilities – the promising sets of differing assumptions – is vast and requires much exploration. However, theorizing within the interdisciplinary domain of Cognitive Science should also involve integrating theories from different silos into unifying theories with increasingly larger domains of applicability. The question is, can we surmount the landscape of institutionally entrenched, isolated silos?

To achieve this, we argue that we need to have more explicit discussions relating Cognitive Science to Philosophy of Science, Epistemology, and Ontology. However, such efforts, although laudable, generally end up in their own silo and have little impact outside of it. To avoid this, we propose that such theories should be embodied in online systems for managing, contrasting, and relating different silos to each other, as shown in Figure 1, for the purpose of greater

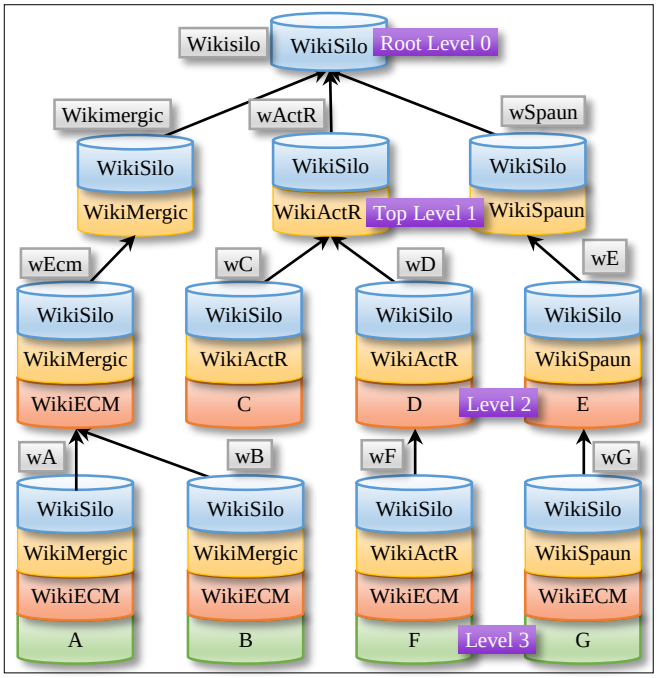

Figure 1: Twelve WikiSilo hierarchies arranged hierarchically – One root level and three top level

unification. To illustrate this we describe our own system for doing so, which is called WikiSilo, and which can be thought of as embodying our unifying philosophy. A working prototype is available online and a part of that system provides the hyperlinked details, which can be found throughout this paper. The prototype can be found at:

`http://en.wikimergic.org/wiki/Main_Page`

WikiSilo is meant to interact with Wikipedia™. Wikipedia is a crowd sourced tool for bringing together a diverse set of viewpoints and many of the existing silos within Cognitive Science are already reasonably *described* therein. However, although Wikipedia can mirror the pluralism within science, it is not set up to be used as a tool to create unification. WikiSilo is meant to provide a crowd sourced means of facilitating this goal.

The WikiSilo tool provides a platform that is meant to be used by individual theoreticians to explicate their assumptions and to theoretically explore their ramifications. The goal is to provide individual silos within WikiSilo with the means to demonstrate the coherence of a large and complicated theory that goes beyond what can be achieved in individual journal papers or on Wikipedia. In part, this is

achieved by a process that relates WikiSilo to Wikipedia and to journal papers (see Figure 2). Ideally, each silo within WikiSilo will make as strong a case as possible for the theory represented by the silo. However, WikiSilo is also designed to use crowd sourcing to drive the self-organization of a silo hierarchy that explicates both the synergies and the disagreements between silos.

Stand-alone benefits of WikiSilo

WikiSilo is best introduced in relation to Wikipedia™. Please note that there is absolutely no affiliation between Wikipedia and WikiSilo.

A WikiSilo provides your own private version of Wikipedia that has been pruned down to include only your research interests, and edited such that it reflects your own coherent view and interpretation of the world – of empirical evidence and rational argument. In other words, whereas Wikipedia articles must fairly describe the current pluralism within science, your WikiSilo is meant to describe your own unified account of your area. In addition, you can include and annotate alternative views with evidence or rationale as to why they are not, on balance, as good as your own. You are the research authority for your own WikiSilo and can control who can add content.

Used in this manner, a WikiSilo offers a number of advantages. First, complex systems and associated theories have many non-hierarchical interrelationships that are difficult to communicate effectively within the linearity of a scholarly paper or even within a book. This situation is made worse in an interdiscipline as each of the subdisciplines will be asking different research questions over the same content. A WikiSilo allows for non-linear reading as directed by the needs of an interested scientist and their line of inquiry as influenced by their subdiscipline. Moreover, articles can be written at many levels of detail and for different target audiences so that specific questions can be answered or details easily glossed over.

What helps communicate complex ideas can also help to analyze complexity in order to improve your theory. A WikiSilo can be used as a theorizing mind tool for uncovering your own connections and for slowly building up your unified/coherent theory of the world. In the beginning, a silo within WikiSilo will likely be neither coherent nor complete. WikiSilo is designed to facilitate a process to achieve these goals.

While a silo within WikiSilo can be set up to allow you as the *only* editor, it is still freely open for reading by others so all your research is easily accessible by anyone around the world online – always current and always complete. Typically it is extremely difficult to find all the up-to-date material about a theory in a single place. Portions are often distributed across multiple papers, across multiple journals, and over time. Old material may become outdated and there is no such indication within archives. Completeness is extremely problematic and expensive in an interdisciplinary field where your research institution might not have subscriptions to all the relevant journals. Most importantly,

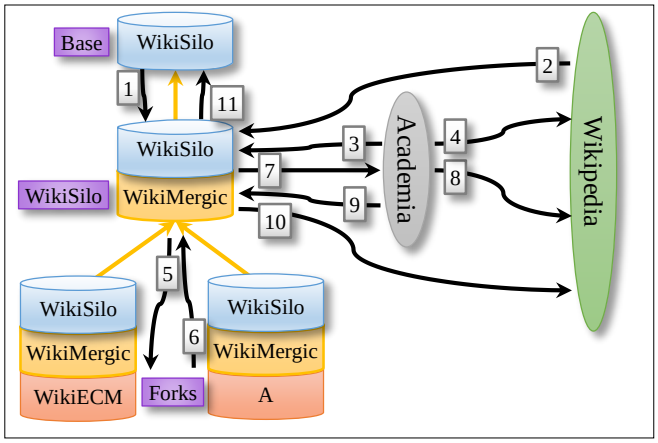

Figure 2: WikiSilo unifying lifecycle

your WikiSilo can have a level of detail that would not be possible in even a large book. This could include material on how to use your theories, best practices, lessons learned from case studies, material from workshops, tutorials, exercises, hyperlinks to related sites, etc.

All the previous benefits could be had with any number of simpler wiki based systems. However, a WikiSilo is partially integrated (but not affiliated) with Wikipedia and includes significant amounts of support content of interest to all scientists. You can have your own annotated copies of Wikipedia articles and your version will be automatically attributed to Wikipedia as required by its Creative Commons Attribution-ShareAlike License (CC-BY-SA). This implies that your content has a similar license, and anyone else using it, as implied by the hierarchy of Figure 1, must also attribute your work. Once your own work has been published in academia and referenced within your WikiSilo, you can simply copy your summary articles directly back into Wikipedia so that other scientists can more easily find and take advantage of your research. This is illustrated as step 10 of the WikiSilo lifecycle (Figure 2). If your unified theory becomes extremely popular, WikiSilo provides for multilingual support and has additional accessibility features derived from Wikipedia.

In short, silos within WikiSilo are meant as a tool to demonstrate, in detail, that you have a coherent account of cognition or a way of getting there. This has the side benefit of limiting critical commentary from anyone not using the fundamental theoretical positions within your WikiSilo hierarchy (Healy & Perry, 2000).

"If you want to comment on material in a WikiSilo, you must adopt the theoretical axioms of that WikiSilo. Otherwise, feel free to create your own WikiSilo with theoretical axioms you can accept.

Self-organizing benefits of WikiSilo

The previous section discussed some of the benefits of a WikiSilo when used as a disconnected island. While these advantages remain, they do not constitute unifying theory. In this section, we summarize how a hierarchical social network

of WikiSilos (Figure 1) helps to unify Cognitive Science by methodologically exploring and exploiting the theoretical space in a disciplined manner. Note that in our terminology, theories *of* science are called ontologizing theories, while theories *for* conducting science are called epistemizing theories. Thus WikiSilo theory is a minimal epistemizing theory.

Unifying bases and forks

Every WikiSilo can function as both a unifying base and a unifying fork as illustrated in Figure 3. For example, the WikiSilo named Wikimergic is a unifying fork off the base named Wikisilo – the root of all top-level WikiSilos. However, Wikimergic can also function as a unifying base to its own forks named WikiECM and wA. These WikiSilos are owned and controlled by the first author.

It is the primary purpose of every fork to unify its base by increasing the base's domain of applicability to cover additional phenomena. A base is not extended directly, but indirectly by creating a new unifying fork with additional constraints (or axioms) over its base. Once created, a fork does not modify its constraints, but explores their ramifications, by, for example, testing them within a unified cognitive model. *Sibling* forks off the same base, e.g., WikiECM and wA, compete to be the one to merit being merged back into their base as determined by the research authority that owns the unifying base. Note that a fork can be owned by another research authority – anyone must be allowed to fork off your WikiSilo as required by CC-BY-SA.

Once a fork has merged into the base, it is obsolete, as are all corresponding sibling forks. Other sibling forks that might be attempting to improve a completely different aspect of their base would need to be refreshed with the updated base. If a fork chooses to remain after its base has changed, then it must become a sibling to the old base by rehoming to the grandparent base and taking on all the unifying research that was occurring in its old base. Having different axioms, it is no longer relevant to its old base and cannot add commentary therein. Debate must occur at a common grandparent base, if it exists. Top-level WikiSilos have little in common as the Wikisilo root is currently empty, save for the minimal WikiSilo theory.

Perhaps a more concrete example is suitable. Imagine that a WikiSilo named WikiActR existed that housed the cognitive theory behind ACT-R (Anderson, 2007) and was managed by the research authority of John Anderson and his team. Most ACT-R based cognitive models are local models to explain a specific psychological phenomenon and are not integrated into a single unified model. These should still be retained within the base so that when ACT-R theory is improved, they can be re-tested for accuracy. That is, their role was to *use* ACT-R and not *extend* ACT-R. In contrast, models can extend the proper use or epistemology of ACT-R. For example, *threaded cognition* (Salvucci & Taatgen, 2011) is not enforced as part of the cognitive architecture but is a unifying way of doing multitasking. It could be one of several generic design patterns approved for use and

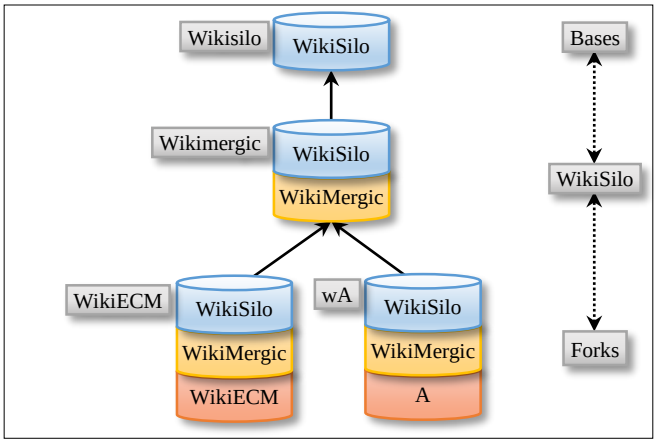

Figure 3: WikiSilos functioning as unifying bases and forks

documented within WikiActR. A few models do attempt to improve ACT-R (Chandrasekaran, Banerjee, Kurup, & Lele, 2011), and these would make good exploratory forks. However, by choosing Anderson's theory as a base and Anderson's authority as ruler of winning forks, researchers must a) refresh their fork whenever ACT-R is updated, and b) must be satisfied if a competing solution is finally adopted. If not, they can create a splinter ACT-R WikiSilo that absorbs their fork and stands as an alternative account, which of course would require a different name.

Comparing forks as a game of chess. Sibling forks should not necessarily be compared to each other in the short term, as each is a long term research proposition. The reason for this can be explained by an analogy to chess.

The main loop common to all automated games of chess is a problem space search (Newell & Simon, 1976). A complete search would take too long (Chase, Hertwig, & Gigerenzer, 1998). Thus all games of chess prune their search by the use of a fitness function. The simplest such heuristic simply adds up the values of the remaining pieces by their static value, for example a one for each pawn, and a twenty for the king. The difficulty is that nobody has ever produced a *linear* fitness function that converts the search landscape into one with a single maximum to be found. Theoretically, this is computable (given enough time and space) and easy to describe. If it is white's turn to play, and I now have say twenty possible moves for this turn, then a linear fitness function would tell me, for each of these moves, what is the guaranteed minimum number of turns before I would win, or zero if that move does not guarantee a win. Negative numbers do the same for black. Thus I would be guaranteed to win by playing any move with a positive number, and I would win quickest by choosing the smallest such number greater than zero. If I do not have any winning moves, I could prolong the game by playing the move with the lowest negative number. A check-mate is any move with a value of two, a win is any move with a value of one (that takes the king). With a linear fitness function, I would only compare a handful of current moves and would not need to look ahead.

For lack of a linear (or convex) fitness function, the search landscape has many local maxima, and one must play many turns ahead to evaluate the future in order to determine the best option for now. For example, sacrificing a piece to gain strategic long term advantage is not possible without looking ahead.

A ten-level WikiSilo hierarchy is akin to looking ten moves ahead in chess, or equivalent to attempting ten valley crossings (Weissman, Desai, Fisher, & Feldman, 2009) in pursuit of a more global maximum. One reason for using the WikiSilo hierarchy, is as an explicit, methodical and disciplined storage of research options, i.e., the current set of moves being explored.

Constraints with degrees of freedom. Every fork has an additional constraint over its base. Nevertheless, constraints do not reduce the degrees of freedom to none. For example, a cognitive architecture allows for many cognitive models. We define *theory* as a predictive system with no free parameters, or one where all parameters have been fixed to a particular set of values. A theory can be expressed mathematically or as a computational model. A *unifying* theory or model is one that remains static over an increasing domain of applicability. As all theories can always be extended, a unifying theory improves over time. However, at any one time, it remains fixed for an increasingly larger set of phenomena being explained.

A unifying theory improves over time is a research program that follows the Lakatosian model of theory development (Cooper, 2006). At any one time it can produce a parameter free theory or model. A WikiSilo is a perfect vehicle for exploring and expressing such evolution. For example an existing theory could be explored by three possible extensions, and this would be equivalent to creating three new forks off the current WikiSilo.

Theory death (not comparison)

Maslow famously wrote that "I suppose it is tempting, if the only tool you have is a hammer, to treat everything as if it were a nail" (1966). Eventually, of course, poor theories become exceedingly stretched, less parsimonious and more implausible. The current WikiSilo is the hammer, and every additional fork is an extension to unify additional phenomena given the nails available in the theory. Eventually a unifying theory may become too stretched. If an alternate fork can explain the same set of phenomena in a simpler manner, then the current fork will finally be abandoned.

In summary, theories cannot be compared as they are actually research programs undergoing change. A poor theory could eventually be improved to surpass a good theory, and a good theory might not be able to accommodate an additional phenomenon in its domain of applicability. Theories or forks are simply abandoned when they become too wieldy to improve, assuming of course, that one is indeed trying to use them to explain an ever larger set of phenomena with no free parameters.

However, it is not clear whether discovering new phenomena and explaining them with local theories leads to scientific progress. WikiSilo provides a language to discuss these issues.

Signal to noise ratio as a measure of scientific progress

The WikiSilo concepts of forks and bases allow an analytical definition for *theoretical progress* in both absolute and relative terms. We exemplify the *relative* measure first in relation to some arbitrary WikiSilo functioning as a unifying base somewhere in the global WikiSilo hierarchy. Relative to this base, the construction of additional exploratory forks is considered as increasing the overall theoretical noise. There is now an increasingly pluralistic set of conceptions emanating from the unifying base that ultimately will need to be arbitrated. If a fork is abandoned, and this is clearly indicated, then that constitutes a reduction of theoretical noise. Note that journal archives do not indicate the theoretical state of any of their older submissions, and that is an advantage of an online WikiSilo.

However, when a specific fork is deemed worthy of being merged into the base, then the base increases its domain of applicability and that is considered as an increase in the theoretical signal. Moreover, after the merge, all corresponding forks are obsoleted and this leads to a reduction in theoretical noise as well. Recall that other forks may exist to extend the base to cover different phenomena, and these will be updated.

The *absolute* measure of theoretical progress refers to the single root level WikiSilo called Wikisilo managed by the Wikisilo.org research authority. The only forks it could accept for merging would be ones that have more or less universal agreement as to their suitability. In the "hard" sciences that have undergone their Kuhnian revolution (Kuhn, 1962), there could be much content. For example, in Physics, relativized quantum mechanics and general relativity could be two unifying theories that would be accepted, and Newtonian mechanics would also be accepted as a useful and compatible approximation in a modestly reduced domain of applicability – away from light speed, astronomical/nano distances and dense objects.

However, within Cognitive Science, there are too many competing theories and approaches with exceedingly small unifying domains of applicability such that none would be accepted today into the single master WikiSilo. Cognitive Science has no central theory. It has many local ones but no overarching theory, not even a provisional one (Gigerenzer, 2010). On the other hand, the fecundity in creating exploratory forks has been enormous (Newell, 1973) leading to a geometric increase in theoretical *noise*. How to convert that noise to a signal is the main purpose of a WikiSilo. After all, not only do we want to increase the signal, we also want to improve the signal to noise ratio and reduce noise.

Millikan (1923) showed how theory and data need to walk hand in hand. Each additional phenomena ought to improve a unifying theory, rather than be related to a one-off, throw-away local theory or surrogate (Gigerenzer, 2009). Just like physics is divided into experimental and theoretical branches,

Cognitive Science should have a pure theoretical branch devoted to unification as well.

Coherence as software version control. The field of software engineering has developed a discipline to add new functionality (and fix bugs) via a set of coherent changes using a version control system. This allows a group of related files to be changed and tested over time, and more importantly, the entire set of changes to be merged into a release stream as an atomic or coherent update. Indeed the amount of time to develop a coherent set of changes with improved functionality can be significant. Most operating systems may have local bug fixes delivered weekly, but new versions of the operating system take years to develop.

Wikipedia has a page-by-page revision control system, but cannot group a set of changes together. As numerous individuals are independently updating pages, it becomes increasingly unlikely that any large theory can be described coherently (even if there is no pluralism). However, a WikiSilo (or unifying fork) functions as a version control system whereby a set of unifying changes that improve the base (or eliminated an inconsistency) can be developed and tested for coherence by the research authority. Each unifying fork can be thought of as an independent software design team developing the next big thing, and it is up to the authority of the unifying base to determine which competing team wins.

Instead of validating the coherence of software, WikiSilo allows for validating the coherence of ideas – of conceptions and theories. Instead of *checking-in* a software file for update, a concept from the unifying base is *checked-in* for improvement.

Unifying coherently

Steps 2-7 of WikiSilo lifecycle (Figure 2) involve unifying the content of the WikiSilo. While how to do so is beyond the scope of this article, the mechanics are as follows.

Articles that may be relevant are copied from the base, or from Wikipedia. They may be highly ambiguous and incoherent. New articles can be created at any time. However, the copied content must be changed into a coherent form, and this is accomplished by several means.

Coherence overlay at the page level. A WikiSilo can function as coherence overlay to base or Wikipedia pages by distinguishing them as

- theorizing content in green. This is your coherent set of conceptions and definitions clearly demarcated from the sea of interdisciplinary ambiguities
- obsoleting content in red. This is the set of incommensurable ambiguities you are attempting to eliminate but must use as they are currently part of the lingua franca.
- support content in blue. These are atheoretical conceptions.

Hyperlinking from journal articles. The theorizing, obsoleting and support concepts are hyperlinked within journal articles (if allowed). The purpose of having hyperlinks is to be able to be 100% explicit on the terminology. Any two researchers might use the same word with a slightly different meaning, and as long as their specific meaning is to be found within their specific WikiSilo, then the article can be fully understood. For example, the term *working memory* can be used in a *descriptive* sense or in an *explanative* sense, and it may not be obvious to readers, especially from another subdiscipline, what was intended, or what specific working memory theory was indicated.

Hyperlinks reduce the clutter of citations to every term, but even citations would not be adequate as the term might have numerous meanings within a journal article. Actual citations (as in this paper) are reserved for significant theoretical positions that must be attributed. Hyperlinks are for common terminology that today includes numerous ambiguities and implicit assumptions. Within WikiSilo all significant terminology should be explicated for a fully parameterized theoretical specification.

Annotating and abridging. Both obsoleting and support concepts can be annotated and abridged as required. Annotations allow one to indicate added content or ~~deleted content~~ within an existing page. Abridging allows for removing material that is irrelevant to your WikiSilo so as to focus attention on the important aspect. However, the hidden material will not imply deletion from the unifying base.

Conclusion

A WikiSilo is likened to a software version control system whereby time can be taken to make coherent a set of changes across the entire theory and determine their unifying impact. Version control systems can operate on a hierarchy of changes, as can WikiSilos. However, instead of operating on software files, a WikiSilo operates on theories.

A WikiSilo is also likened to a game of chess where each level in the WikiSilo hierarchy can be considered as evaluating a move one step further into the future as required when no linear fitness function is available. Theories cannot be compared as they are actually long-lived research programs. However, a WikiSilo can demonstrate that a network of conceptions is coherent, something that is exceedingly difficult to do otherwise. Moreover, a WikiSilo can limit irrelevant debate from elsewhere as they do not belong within your WikiSilo hierarchy.

Finally, the WikiSilo hierarchy is self-organizing via exploratory forks. The best forks are exploited by being merged into an increasingly unifying base leading to theoretical progress which can be formally measured based on a signal to noise ratio.

References

Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe? Science* (pp. x–290). Oxford University Press.

Chandrasekaran, B., Banerjee, B., Kurup, U., & Lele, O. (2011). Augmenting Cognitive Architectures to Support Diagrammatic Imagination. *Topics in Cognitive Science*, *3*(4), 760–777. doi:10.1111/j.1756-8765.2011.01156.x

Chase, V. M., Hertwig, R., & Gigerenzer, G. (1998). Visions of rationality. *Trends in Cognitive Sciences*, *2*(6), 206–214. doi:10.1016/S1364-6613(98)01179-6

Cooper, R. P. (2006). Cognitive Architectures as Lakatosian Research Programs: Two Case Studies. *Philosophical Psychology*, *19*(2), 199–220. doi:10.1080/09515080500462388

Gigerenzer, G. (2009). Surrogates for Theory. *APS Observer*, *22*(2), 2–3.

Gigerenzer, G. (2010). Personal Reflections on Theory and Psychology. *Theory & Psychology*, *20*(6), 733–743. doi:10.1177/0959354310378184

Healy, M., & Perry, C. (2000). Comprehensive criteria to judge validity and reliability of qualitative research within the realism paradigm. *Qualitative Market Research: An International Journal*, *3*(3), 118–126. doi:10.1108/13522750010333861

Kuhn, T. S. (1962). *The Structure of Scientific Revolutions* (3rd ed., pp. xiv–212). University of Chicago Press.

Maslow, A. H. (1966). *The Psychology of Science: A Reconnaissance* (First Edit., p. 168). HarperCollins.

Millikan, R. A. (1923). Nobel Lecture: The Electron and the Light-Quant from the Experimental Point of View. *Nobelprize.org*. Retrieved January 30, 2014, from http://www.nobelprize.org/nobel_prizes/physics/laureates/1923/millikan-lecture.html

Newell, A. (1973). You Can't Play 20 Questions with Nature and Win: Projective Comments on the Papers of this Symposium. In W. G. Chase (Ed.), *Visual Information Processing: Proceedings of the 8th Symposium on Cognition* (pp. 283–308). New York: Academic Press.

Newell, A., & Simon, H. A. (1976). Computer science as empirical inquiry: symbols and search. *Communications of the ACM*, *19*(3), 113–126. doi:10.1145/360018.360022

Salvucci, D. D., & Taatgen, N. A. (2011). Toward a Unified View of Cognitive Control. *Topics in Cognitive Science*, *3*(2), 227–230. doi:10.1111/j.1756-8765.2011.01134.x

Weissman, D. B., Desai, M. M., Fisher, D. S., & Feldman, M. W. (2009). The rate at which asexual populations cross fitness valleys. *Theoretical Population Biology*, *75*(4), 286–300. doi:10.1016/j.tpb.2009.02.006