Cognitive Sciences
Volume X, Issue X, pp. XXX-XXX

Opinion

EXPLOITING THE RELATIONSHIP OF NATURAL LANGUAGE AND COMPUTER SCIENCE: A NOVEL THEORETICAL APPROACH TO FAIRNESS

Charles E. Bailey [*]

Global Institute for Scientific Thinking, Lake Mary, Florida, 32746 USA

This paper links the challenges of understanding human behavior, including speech, cognition, emotions, fairness, empathy, and cooperation, to the concept of robustness (dynamic adaptability to disturbances in a changing and uncertain environment). This connection will be explored in light of a proposed corollary between natural (human) language and tools from computer science. The author contends that humans come closest to establishing fairness when they can share as accurate an understanding of circumstances and views as possible. To that end, this paper proposes that humans can best develop an accurate understanding with a usage-based natural language model that establishes a reference point of cognitive accuracy paralleling the classes and methods used in the software language of computer science. Such an approach allows for evaluating language for the relative accuracy of its representation of classes and methods using cognitive accuracy as a fitness function. The hypothesis is that natural grammar can be distinguished in a narrow sense by regional, cultural, or local use and more broadly across the human species by scientific, or *global use*. This, in turn, allows us to design and execute scientific experiments intended to identify human language usage representing an idealized scientific global optimum that can act as a potential reference point for maximizing harmony, cooperation, robustness and the perception of fairness. Further implications and the relationships between the variables and future work will be discussed.

---

[*] Correspondence should be addressed to: Charles Bailey, Director: Global Institute for Scientific Thinking, Lake Mary, Florida, 32746

Keywords: Fairness, cognitive accuracy, natural language, usage-based language, cooperation, computer science, novelty search

1. INTRODUCTION

Throughout the natural world we observe varying levels of competition, cooperation, and survival among and between organisms and species. From an evolutionary view, these struggles appear related to the overall fitness of individuals and populations, and to the survival of the fittest as proposed by Charles Darwin in *The Origin of the Species*. In human primates, it seems at first glance that cooperation is enigmatically intertwined with competition and conflict — so enmeshed that conflict and war seem to trump reasonable problem solving, often favoring competition and inhumanity over cooperation and humane treatment of others. Wars beget wars, and fairness and unfairness are bandied about as slogans promoting higher values, but more likely to represent expressions of implicit beliefs, ambiguous assumptions, biases, and demands characterized by prejudices, blaming, judging, punishing, retribution, war, and threats of war (Mitchell et al., 2006; Phelps & LaBar; Milgram, 2004; Beck, 1999; MacLean, 1990). How can this enigma exist in humans, the human animals that *know they know* (Risberg, 2006)? How can we regularly use terms that most people recognize as positive and desirable, such as fairness, love and peace, in ways that effectively run counter to our stated goals?

It seems that natural human language usage represents a major variable that biases how accurately and harmoniously humans interface across cultures and with the natural environment. Human language springs from and directly affects cognition, emotions, cultural values, and behavior, with implications for reliable communication (Bailey, 2006, 2007; Brown & Keeley, 2007; Chow & Cummings, 2007, Cummings & Miller, 2007; Decety, 2007; Gazzaley & D'Esposito, 2007; Ochsner, 2007, 2006, 2005; Phelps & LaBar, 2006; Adolphs, 2006; Cacioppo & Berntson, 2004; Logothetis, 2004; Phelps, 2004; Wright, 2004; LeDoux, 1996, 2002; Damasio, 2000; Mesulam, 2000; Panksepp, 1998; Ellis & Harper, 1997; Luria, 1981; Beck, 1976). How we use language affects not only how others understand us, but how we understand ourselves and our surroundings — the domain we live in. As a consequence, how we use language directly affects how effectively and equitably we can develop a shared sense of fairness. The more accurately we can globally understand and articulate the circumstances and concerns of others, the closer we can come to agreeing on the most desirable terms for the fairest possible interactions with them (Figure 1).

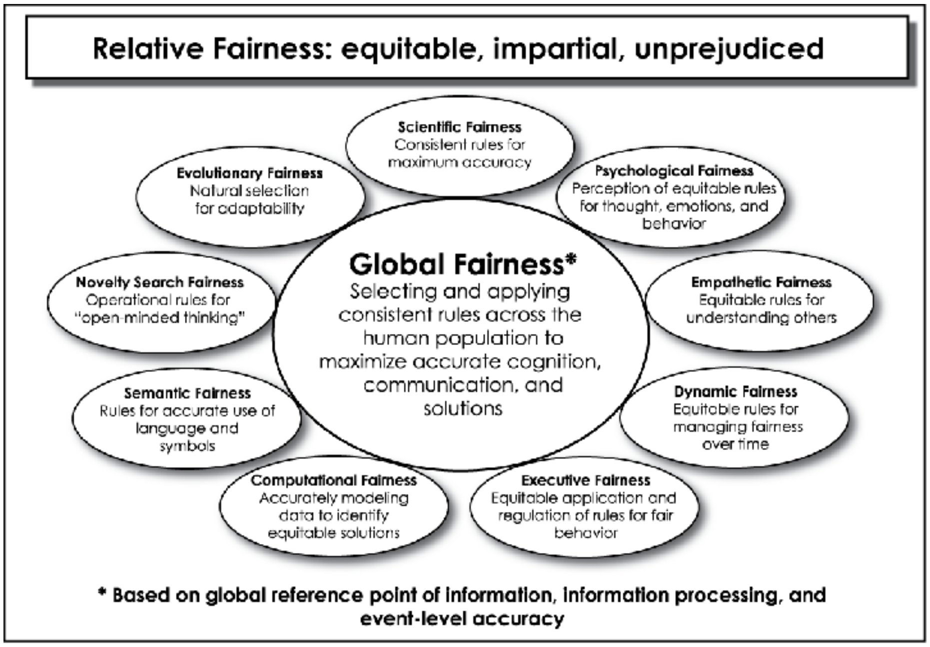

Figure 1. Relative fairness: equitable, impartial, unprejudiced.

2. MOTIVATION

There is a wealth of knowledge about human brain and behavior across many disparate fields. Unfortunately, it is often difficult to integrate this almost overwhelming mass of complex information as an understandable and useful whole. Scientists can moderate these constraints by agreeing on a common language with reliable semantics for expressing and communicating their insights from their own areas i.e. fundamental vocabulary and definitions, percepts and concepts, nomenclature, experimental design, conclusion, applications, etc. Some might contend that we already do this and have most of this information, and therefore might ask what we can learn if we explore this matter further. “We already know that. We already do that.” This begs the question; do we know all this already? Who already does this? What do we actually know? And just as important, do we apply what they know?

The concept of knowledge itself presents some concerns. People often say they know or believe something even in the absence of any underlying factuality. Humans tend to cling dearly to individual beliefs, whether substantiated or not, but usually insist on irrefutable evidence when asked to reconsider them or to consider the beliefs of others. This imbalance in what we require for evidence allows us to keep our own beliefs intact and unquestioned while dismissing contrary beliefs with barely a second thought. This makes it easy to find one small blemish as an excuse to reject the whole. This seems to affect most humans regardless of their level of education, even scientists. It makes it difficult for us to open-mindedly explore new information and new ideas, and imposes significant constraints on the use of information we do find.

Indeed, since much of this information is already available, this paper attempts to search for a common denominator that will practically facilitate its integration and application. There is a distinction between having knowledge, integrating and understanding knowledge, and applying knowledge. With this in mind, the challenge is to make use of this information — actually applying this knowledge with the hope of making the world a better place — a better place for us, our children, and future generations. We can apply our knowledge most effectively and equitably if we can identify a mutual basis for communication and agreement. Language is still evolving. By adjusting our use of human language to reflect a common scientific reference point of thinking as accurately as possible, cognitive accuracy may offer such a common denominator. We then face the logical implication: if we accept this common reference point in general terms, then can we benefit individually by applying it to our own specific beliefs and biases as well?

3. BACKGROUND

Humans have evolved the capacity to use language to a much higher degree than any known species, including nonhuman primates. We use language to represent information as knowledge and as a result we have accumulated a litany of words, ideas, beliefs, books, libraries, educational institutions, etc. (Nowak, 2006; Mesulam, 2002). Language not only enables humans to express information about thoughts and feelings, it shapes and influences those thoughts and subsequently our behaviors. Language endows us with the capacity to speculate about and confirm the thoughts of other humans. A bias toward accuracy helps transforms this capacity into ability. Language provides a symbolic abstract vehicle for categorizing information about objects, the behaviors of objects, and the relationships of objects to each other and to the natural world (Nowak, 2006).

Language, as a tool, has inspired us to develop technology, art, literature, economics and many other abstract and concrete human inventions. But, we see a disturbing tendency to apply these inventions within a context governed by hierarchal dominance and subordination, superiority and inferiority — as if our success results from our power to get ahead of others or to get back at them. We seem to have mastered the art of competition, punishment, war, and destruction — and employ them as if the survival of our species depends on it. But the harmonious cooperation that might actually ensure our survival continues to elude us in many substantive ways.

These baffling circumstances might be better understood using an evolutionary approach integrated with our knowledge of the human brain and how language accuracy biases our cultural belief systems and our behavior. As humans, if we can understand these relationships and their effects on our thoughts, emotions, behaviors and competitive tendencies, then perhaps we might find a recipe, or algorithm that enables us to use language as accurately as possible — thus enhancing cooperation, harmony, and at least the perception of fairness. If we ignore the effect of inaccurate language use and its inherent effect on emotions and behavior, we will likely continue to use this critical tool ineffectively and inexpertly in frustration and confusion, with little hope of achieving any more harmony that those who went before us.

Humans are inherently biased due to their individual histories, including cultural belief systems. In that respect, it may not be possible to eliminate bias in humans. If this is indeed the case then a reasonable alternative might be preferentially choosing the bias that provides the most robust solutions. This suggests that striving to promote an accuracy based bias could prove beneficial.

Due to its value of scientific fairness, science offers reliable principles for uncovering and managing bias, as well as an accuracy-based methodology with which to explore this conundrum (Browne & Keeley, 2007). Additionally, computer science itself offers powerful tools for finding solutions to difficult problems.

> In today’s world, high-level decision makers are struggling to develop policies for a wide range of complex adaptive systems...It is hard for me to imagine how intelligent decisions can be made without the kinds of insights our community (computer science) can provide. (De Jong, 2006)

4. EVOLUTIONARY FAIRNESS

While evolution has invested billions of years into the development of humans and our human brains, we only developed scientific thinking fairly recently. In a sense, the language of science represents evolution of human language to a higher degree of accuracy. Scientific knowledge has increased dramatically in the last hundred years and in the last decade knowledge about human brain and behavior seems to have increased exponentially. Science strives to understand the domain of the natural world as the dynamic and uncertain platform on which evolution runs experiments to find more adaptive solutions to the problem of survival. So it comes as no surprise that the concepts of evolutionary and scientific fairness share many similarities. We can define evolutionary fairness as achieving harmony or a congruent fit within and between organisms and the environment, i.e. equilibrium or homeostasis. Viewed optimistically, it is tempting to say that non-partisan natural selection yields winning solutions for the most adaptive organisms and species as a whole. From a more dispassionate view, selection seems biased, however, inadvertently, towards finding the best solutions by eliminating less adaptive ones.

Natural selection seems to do the dirty work by eliminating less fit organisms. In other words, selection is evolution's way of finding out how things work most successfully in the real world. Thus, overall, evolutionary harmony represents the expression and regulation of a set of *best solutions* for robust adaptation to the conditions of a dynamic and uncertain environment. In this view of evolution, natural selection can balance the aspirations of all to assure the survival and propagation of the fittest or most robust organisms. Robustness can be thought of as skills or behaviors that represent a high level of dynamic adaptability. In turn, robustness in evolution has been equated with cooperation, because in evolution, cooperation augments dynamic adaptability (von Dassow & Meir, 2004). "A biological system is robust if it continues to function in the face of perturbations" (Wagner, 2005).

5. SCIENTIFIC FAIRNESS

Science can be viewed as a form of human culture, with specific rules for scientific thought, and behavior. This system of beliefs represents scientific values used to classify information as knowledge. However, unlike most typical beliefs and values of other human cultures that resist knowledge assessment based on accuracy, the beliefs of science are generally stated probabilistically as scientific facts, based on accuracy, validity, reliability, error variance, confidence intervals, etc. Science uses observation, experimentation, and repeatable results to draw conclusions in support of the value of scientific fairness. This requires rules and methods designed to eliminate or minimize any inaccurate experimental bias, and to resist *a priori* favoring of one among all possible outcomes. The goal is to obtain the most accurate results in the form of sound conclusions and solutions that hold up to further testing without regard for the preference of the experimenter. Dynamically, scientific fairness favors accuracy and accuracy enables scientific fairness.

Thus, scientific conclusions are meant to represent reliable explanations of cause-effect relationships in nature regardless of how the experimenter might view these relationships. In other words, science searches for the most accurate descriptions possible of observed relationships among objects, and between objects and the natural world. When science establishes a relationship, it is so because of scientific observation and testing of evidence about the real world. We accept this as scientific fact, not just because someone *says it is so*, not even if that person has expert status, but because we can make the same observations and follow the same rules to see it for ourselves. This can be applied across cultures — whether the expert represents science or some other cultural belief system. Science says, “Show me.” Science uses fairness to find out how things *really* work in the real world, similar to evolution. This, in turn, creates a winning situation, when the results allow us to interface with others and with the environment more accurately and thus more effectively and predictably.

This doesn’t mean that science guarantees exactly correct conclusions. Rather, science continually makes adjustments to approximately correct scientific facts while updating and refining knowledge over time. In this respect scientific knowledge with a reference point of accuracy seems preferable to knowledge alleged to be factual but without an established reference point. Scientific fairness, qualifies as equitable because accurate knowledge shifts the bias from inaccuracy towards accuracy by providing a more congruent world view. We then achieve a greater capacity for more efficient and harmonious homeostasis overall. Does this apply to psychological fairness, and if so, how does it apply?

6. PSYCHOLOGICAL FAIRNESS

Viewed as a corollary to evolutionary fairness and scientific fairness, psychological fairness can be defined as harmony resulting from agreement on accurate evaluation of the relationship between feelings, attitude, and action. Think of it in terms of the relative congruency between an individual's thoughts, feelings and behavior. This definition can be extended to include equilibrium with the environment and with others as part of the environment. This suggests that accuracy is also important to *psychological* fairness based on what we have learned from *scientific* fairness.

When we look at psychological fairness from a congruency perspective, we can establish possible measures for evaluating the agreement between what an individual says and does. That is, we can observe the degree of agreement between what and how they think and how they behave, including speech. The person's behavior represents a measure of the equitable relationship between the rules they profess to use, the rules they *actually* use, and the rules they expect others to use. In other words, we evaluate the degree to which they practice what they preach. Do they inconsistently use one set of rules for themselves and another set of rules for others? As an observable phenomenon, congruency between rules and behavior represents a reference point for how reliably and equitably we interface with others.

Humans have different personal frames of reference, biased by their histories and inheritance of local cultural belief systems. These unique and often incongruent frames of reference, usually laden with hidden assumptions, can lead to the perception of inequitable outcomes, misunderstandings, and conflict — especially since most cultural belief systems usually depend, implicitly or explicitly, on a local body of assumed knowledge embedded with static, opaque, and often unreliable rules for operating in a dynamic uncertain world. Hence, decisions based on these rules are easily perceived as inequitable. When we apply these static rules to the inherent dynamic uncertainty of our environment, it usually predisposes us toward variable and ambiguous results. In effect, the application of the rules tends to vary from minute to minute, day to day, person to person, and culture to culture.

This makes congruency problematic, and tends to interfere with reliably achieving ongoing equitable solutions to the daily problems we face. On the other hand, people have been shown to resist inequitable outcomes (Fehr & Schmidt, 1999). This suggests that creating a transparent frame of reference biased toward accuracy and observable congruency would have potential benefits for psychological fairness. Science offers such a frame of reference based on the long held values of its world view: to fairly understand the natural world using a global belief system biased toward accurate reasoning and effectively characterized as cognitive accuracy.

7. EMPATHETIC FAIRNESS

In its strongest sense, empathy could be described as understanding not only the way others feel, but also how and why they think, evaluate, and behave the way they do. Understanding someone's situation implies that we have more accurate and detailed information about them, which in turn makes it harder to relegate them to a subjective abstract category that we can prejudicially dismiss. In any sense, understanding plays a role, hence the more accurately we view others the more likely we can empathize with them. This highlights the importance of the explicit ability of an individual or group of individuals to mutually transmit and receive clear and accurate information about their situations, expectations and their understanding of others. In other words, the ability to communicate specifically regarding what and how they feel, think, and behave becomes a critical factor in establishing fairness. The skill of seeing things from as many points of view and as accurately as possible maximizes the potential achievement of empathetic communication — even when we disagree with some of those views.

Understanding a person does not necessarily mean agreeing with them entirely, but, the perception of fairness has been shown to affect how individuals think about others (Singer et al., 2006). However, whether in relationships between intimates or relationships among nations, the ability to understand and empathize is a primary step in communication that promotes compromise and cooperation. Relationships are usually complex. In most relationships, it is difficult to solve problems and find collaborative solutions with an inaccurate view of the problem. Realistic rules for interactions based on more accurate view points provide more realistic options for effectively identifying and exploring the problem, and achieving more equitable solutions. Similarly, we might assume that accuracy and explicit empathetic understanding of problems in the natural world would allow for better and more cooperative solutions overall. This suggests that preferentially empathizing accurately with others enhances our ability to interact with them more fairly. As such, empathy fosters a transparent environment conducive to ongoing harmony and perceived fairness.

8. DYNAMIC FAIRNESS

We might then consider how to maintain the perception of fairness in a changing uncertain environment. Even if we have agreed on the general rules of fairness, it is hard to imagine that at any given moment we have an exact balance from each individual's perspective. Furthermore, even a moment of mutual agreeable balance will not likely last in the face of changing circumstances. In the natural world, humans face ever-changing dynamics that can rarely be predicted with total accuracy. Scientists deal with this uncertainty using the adaptive approach of assigning probabilities to various outcomes.

This suggests a relative dynamic concept of fairness that can be applied over time to changing conditions to achieve dynamic, best-fit, more accurate solutions. Defining the general rules of dynamic fairness in terms of probability offers a method for balancing different view points more accurately in an ongoing give-and-take manner to boost the goal of achieving the highest relative degree of perceived fairness over the long run. Methods of dynamic fairness rely on the probability of preferentially achieving a long-term equitable balance of perceived fairness using explicit policies and consistent rules for accuracy and clarity. These policies and rules are important because they manage the bias of information used for communication, planning, negotiation, sharing, empathy, consideration, cooperation, and compromise. These applications of fairness can be extended to computer science and natural language with a usage-based approach to establish the concept of robust fairness.

9. COMPUTATIONAL FAIRNESS

In our struggle to understand humans, many theories appear possible but over time often fail to stand up under the rigors of scientific scrutiny as meeting standards for being probable. Computer science offers useful principles for modeling complex adaptive systems, and tools for fairly evaluating the probability of scientific hypotheses. These tools are rooted in sound scientific principles using computer software capable of accurately representing problems and solutions within a particular context. Computers are essentially input-output devices with internal hardware and software. Computational fairness can be thought of as mitigated by the integrity of the software. The software runs on the hardware to map input to output by transforming input information about the problem to output data as the solution.

The software, in a sense, represents one or more recipes, or algorithms, for finding solutions to problems in a certain domain. Reliable software along with accurate and reliable data produces reliable information. In this sense, software biases the accuracy of the results. Software functions at many levels, somewhat analogous to natural language, including system software, application software, higher-level object-oriented programming language software, and so on. Computer languages and software allow computers to run experiments and find the most accurate solutions to a variety of real world problems from difficult engineering tasks for structural component design, to manufacturing and optimization strategies for production lines and information processing.

10. THE BLACK BOX: NO FREE LUNCH

Computers and the recipes they use for problem solving are sometimes viewed metaphorically as black boxes that perform implicit functions, in effect, mapping inputs to outputs using certain software algorithms that are invisible to those outside the box. Information in the form of data about a problem is entered as input. Then the problem is solved by the algorithm and transformed to output data as information representing the solution. The general contention is that no one algorithm is complex enough to solve all possible problems. Since diverse problems can appear at different times and locations, this adds an element of complexity to finding solutions. So for any given application, it is difficult to know in advance which algorithm is best or if the task is even doable (Rogers et al., 2003).

In a complex dynamic and uncertain problem space, if you choose to stick with only one algorithm that works great for some problems, it may be woefully inadequate for others. Each algorithm has task dependent pluses and minuses. What you pay for on one hand you tend to give up on the other — in other words, "there is no free lunch" (Wolpert & Macready, 1995). However, you might reduce the disadvantages of a single, sometimes inappropriate solution, by having other options with diverse features to choose from. This suggests that more accurate problem solving can result from probabilistically using overall open-ended policies for algorithm selection. This dynamically provides adaptive methods for openness and promotes flexibility and diversity.

11. EXECUTIVE FAIRNESS

This black box metaphor can also help us understand the human brain. As in computers, both the input to the human brain and the resulting output are influenced by our learned software recipes. The brain's executive function (left prefrontal cortex) shows a functional similarity to linear controllers used in many engineered systems. The left prefrontal cortex uses information expressed as language to apply regulatory control to suppress stimulus-bound reactions and enable voluntary behavioral selection (Striedter, 2005). This facilitates the capacity for reliable dynamic regulation of voluntary thought, and furnishes methods for reliably choosing appropriate recipes for input-output mapping. In turn, these accurate, adaptive feedback mechanisms yield more accurate, more effective behaviors. This executive regulatory function can be compared to the responsibilities of an executive of a company (Bailey, 2007).

The executive of an organization is responsible for establishing a set of reliable rules that govern employee behavior and interaction. Rules that are equitably applied multilaterally and consistently tend to be viewed as fair, i.e. everybody plays by the same rules. Subsequently, rules are easier to monitor and reinforce when they fairly present an accurate and reliable representation of expectations for reasonable behavior in that domain. Rules that regulate and promote adaptive communication and cooperative interaction between humans in a dynamic and often uncertain work environment provide a bonus. Such an environment is more conducive to real time adaptive feedback that can regulate and continually fine tune the systems input-output processing for better outcome performance.

The regulatory rules also may be viewed as executive policies that optimize solutions by allowing the flexibility to choose the best methods for approaching different tasks and actively making dynamic corrections to minimize error variance. We see a corollary in the human brain because humans can selectively interrupt the flow of the brain's automatic algorithms and appropriately redirect processing to more relevant methods as tasks change over time. This allows for maximizing our problem solving abilities by appropriately using all of our available resources including our algorithms and concomitant brain circuits and networks. The ability to flexibly choose and optimize the best algorithms for each application demonstrates plasticity — the dynamic capacity for adaptability and more robust solutions. The computer corollary is computation that calls on and regulates various subroutines to process a particular equation involved in problem solving. The more options available, and the more accurately the subroutines function, the more likely the success. Probabilistically, effective executive policies flexibly regulate the open-ended dynamic capacity to predictably find and optimize the most useful solution to specific problems as they occur.

And so it is with the executive policies that humans use to select the algorithms for applying to individual problems. If we rigidly insist on static closed-ended executive policies that supposedly solve all application problems, we run into difficulties with adaptively regulating error detection and correction due to change and uncertainty in our dynamic domain. On the other hand, we allow for plasticity by using reliable open-ended executive policies and algorithms. In other words we can variably and dynamically select, apply, adjust, and optimize the best possible algorithms to specifically solve individual problems as they arise and change over time. This allows for robust adaptation to the ever-changing problems we face in our dynamic uncertain world. This may not exactly add up to a free lunch, but it might optimize our chances of getting the lunch we would prefer (Wolpert & Macready, 1997). Given the uncertainty and variable predictability of the domain we live in, reliable adaptive flexibility offers a potentially robust executive policy strategy — even though it may not always provide the *perfect* solution.

> If a program works in only one way, then it gets stuck when that method fails. But a program that has several ways to proceed could switch to some other approach, or search for a suitable substitute. (Minsky, 2006)

12. ACCURATE REPRESENTATIONS

To execute the program algorithm and generate the most accurate solutions, the computer requires that the problem descriptions provide as accurate a representation as possible of the objects and methods relative to that domain. Objects can be represented as categories and subcategories such as classes, subclasses, and types of information. Methods represent behaviors or functions ascribed to the objects, somewhat like action abilities, or the special skills of an agent in a particular situation. An accurate representation of the objects and methods establishes a more accurate understanding of cause-effect relationships in that particular domain.

This suggests the importance of the accurate assessment of objects, behaviors, and the domain, because the relative accuracy of the problem definition significantly improves our ability to draw effective conclusions, that is, solutions that fit the problem at hand. Solutions derived under these circumstances can usually be tested for appropriate standards of accuracy. The corollary is that inaccurate or misleading representations of the domain, problem, objects or methods will bias the result by increasing the probability of misleading or ineffective solutions or conclusions. This parallels human cognition, where a faulty or inaccurate understanding of the problems we face inadvertently biases us toward the probability of undesirable outcomes, misguided conclusions, diminished predictability, and subsequently ineffective decisions.

13. ACCURATE REASONING

Metaphorically, we can say that our brain uses language, running on our inherited hardware, to perform input-output mapping and transformation of information for interfacing with our environment (Kaufman, 1993). The accuracy of our cognitive input-output mapping derives in large part from the accuracy of our software. Human software consists of our individual history of accumulated rules, recipes, and information that can be represented by language. Although language is often viewed as the result of cognition, it can also be viewed as a parallel function or a driver of cognition. In a somewhat circular manner, language use has a major effect on the ideas we *think* and on the brain circuits and networks we *think with* (Bailey, 2007). Language endows humans with greater overall capacity for abstract ideas (Striedter, 2005) and enhances our computational abilities, that is, the human ability to reason more accurately and effectively.

The accuracy of our ideas biases the results we obtain — because what and how we think is germane to our computation. When we improve the accuracy of the language we use to describe and characterize our circumstances, we predictably improve the relative probable accuracy of our results. How variations in language bias accuracy would seem to be an important element for theories about human thought and behavior. These elements imply important variables and parameters useful for optimizing the congruency of our interactions with others and with the environment — and for our ability to reason about and achieve fairer, more harmonious relationships. In this sense, as humans, we benefit from an approach to fairness biased toward cognitive accuracy, which increases our ability to reason reliably and adapt more robustly to the potential disturbances that we encounter in our dynamic uncertain world.

15. A USAGE-BASED APPROACH TO LANGUAGE: CLASSES, METHODS, AND FITNESS

Our analysis so far suggests benefits of using computer science terms and concepts to develop a scientific representation of usage based human language. We obtain these benefits from the accurate representation of the relationship between human thought, emotion, and behavior. As such, language becomes a useful tool for accurately evaluating human primates. And since language influences thought and behavior, the same tool can be used to understand the concepts of fairness, empathy, and values across cultures. Usage-based approaches for evaluating human language are gaining favor, and seem to offer practical utility for cross-cultural evaluations. But more importantly, a usage-based approach offers descriptive representations useful for reliable classifications of objects, nouns, information, and ideas; as well as behaviors, verbs, intention, prediction, and cause-effect relationships (Tomasello, 2005; Bloom, 2002; Kemmer & Barlow, 2000). This utility seems appropriate for accurately and effectively integrating human language, cognition, emotions, and behaviors with experiments in computer science.

If so, then considering human grammar in terms of usage and accuracy may prove beneficial for determining the overall language fitness and representations of classifications and methods. Individuals or groups use such language representations as they practice input-output mapping within the domain of the natural world. Then in a sense, accuracy represents a fitness measure for the functional use of language, referring mainly to the accuracy of the grammar and semantics in that domain. Such a fitness function tends to be a critical component in shaping solutions (Gruau et al., 1996). The choice of the fitness function determines *what* and *how* improvement is measured and biases the solutions we obtain. This highlights the importance of carefully selecting the appropriate fitness function and accurately representing the domain. Not surprisingly, the most fit language characteristics obtained from these methods represent the most accurate mapping from the problem to the solution and the most reliable cause-effect conclusions.

If there is, in fact, a relationship to accuracy, we might conjecture that different local or global linguistic subsets of human languages could relate to points on parameter gradients between endpoints that represent relative levels of bias. Table 1 contrasts functional features for points along a gradient describing relatively closed-ended and open-ended interaction between language and cognition. Features include: 1) flexibility; the ability to respond effectively in a wide variety of situations, 2) adaptability; the ability to update behavior to account for changes in the environment, 3) robustness; the ability to withstand disturbances and changes in circumstances without having to make substantive changes in behavior, and 4) perhaps cooperation and fairness. Since local and global cultural beliefs and values also depend on language, we might learn more about these language subsets and how they relate to accuracy in the domain of the natural world from further study of these interactions.

Table 1. Functional Interaction of Language and Cognition

Closed-Ended	**Open-Ended**
Ambiguous assumptions and dogmatic beliefs, unsupported by facts or evidence, but rigidly stated as unquestionable "truths of the Universe"; prohibitions against questioning and novelty search. Based on unexamined local classes, rigid ritualistic thought and behavior that promotes ambiguities & misinterpretation of the natural world	Unambiguous assumptions stated probabilistically as theories; hypotheses and conclusions supported by evidence, scientific testing, retesting, with mandatory questioning and encouragement of novelty search. Uses global scientifically examined classes, adaptable thought, and behavior that promotes transparent interpretation of the natural world
Rigid, maladaptive methods, subjective bias, vertical subordinate communication. Absolute, static bias: certain, determinate, static guarantees. Tends to get stuck and stay stuck due to deception, assumes "this is the right way and the only way, this is the only solution"	Flexible, adaptive methods, objective bias, collateral communication. Variable, dynamic bias: uncertain, non-determinate, dynamic probability. If stuck, tends to try something new and thus avoids deception, assumes "many ways and potentially better solutions"
Local limitations for problem interpretation, search, and solutions, limited when applied to global problems	*Global* capacity for problem interpretation, search, and solutions, adept application to local and global problems and solutions.

and solutions. Variable local labeling of objects, classes, and variable local methods	Consistent global labeling of objects, classes, and consistent global methods
Monotonic reasoning; cognition using dichotomous grammar, restricted executive function and problem search space	Non-monotonic reasoning; cognition using multivariate grammar, expanded freedom of executive function and problem search space
Veridical bias: true and false, either-or, absolute, concrete, black and white; constrictive and restrictive, non-contextual, parental, demanding, adversarial	Adaptive associative bias: abstract, gray, gradated; expansive and extensive, contextual, adult, requesting, cooperative
Predetermined certainty, all knowing, resulting in decreased vigilance and awareness: reflexive "implicit thoughtlessness"	Indeterminate relative uncertainty, inquisitive, resulting in increased vigilance and awareness: deliberate "explicit thoughtfulness"
Limited empathy, cooperation, understanding, and fairness due to local frame of reference representation based on unexamined inaccuracies	Expansive empathy, cooperation, understanding, and fairness due to global frame of reference representation based on examined accuracies
Semantic inaccuracy: vague, poorly defined word use, with more inaccurate overgeneralizations and absolute inferences: always, never, every, all, none, etc.	Semantic accuracy: specific, best definition and word use with more accurate graded generalizations and non-absolute inferences: frequently, often, occasionally, many, some, few, etc.
Rigid assertions implying limited, closed-ended forced choices, ignoring broader choices and consequences; imperative and prescriptive; I should, I must, I have to, I need to, and I have got to. "I am obligated." Offers expression of a weak or nonexistent argument expressed "as the solution" to the problem and tends to ignore the number or strengths of assumptions	Flexible assertions implying multiple open-ended choices, based on broader choices and consequences; preferential: I prefer, I would rather, I would like to, I choose to. "It is a choice." Offers expression of a strong argument relating to the solution of the problem and relies on minimizing the number and maximizing the strength of assumptions
Insensitive to inaccuracies of information (classes), of thought process (methods), and of event-level orientation; retroactive, learning "reactive" with focus on retrospective minimization of	Sensitive to accuracies of information (classes), of thought process (methods), and of event-level orientation; forward-thinking, learning "proactive" and "considerate" with focus on learning from

errors and distortion of cause-effect by rationalizing, justifying, judging, blaming, finger pointing, and punishing, i.e. variable irresponsibility and poor accountability, "driving in the rear view mirror"	mistakes using error detection and correction for improvement of cause-effect i.e. active, consistent responsibility and accountability of choices and consequences, "driving with eyes on the road"
Inaccuracies and faulty assumptions, faulty and inaccurate cause-and-effect conclusions, fosters solutions with weak or no arguments	Accuracies and reasonable assumptions, more plausible and more accurate cause-and-effect conclusions, fosters solutions with strong arguments
General unawareness of thought process accuracy, (methods) local cultural belief system, bias unawareness, promotes "blind sight"	General awareness of thought process accuracy (methods) and global scientific belief system, bias awareness, promotes "insight"

(Modified from Bailey, 2007; Bailey, 2006)

16. SEMANTIC FAIRNESS

Last but not least, semantic fairness allows for more accurate expression and evaluation of human thought, emotion, and behavior. As science strives for fairness, it relies on semantics to accurately express scientific concepts, hypotheses, experiments, conclusions, etc. Semantic fairness derives from the accurate representation, categorization and verbalization of information about observed relationships including objects, behaviors, mass, energy, and cause effect. This allows scientists to say what they mean and mean what they say as precisely as possible. The more accurately the natural world is defined, the more equitably it is understood. This establishes a global world view biased towards accuracy.

Semantic fairness makes it possible to accurately think, read, and speak abstractly about fairness itself. Thus, semantics support how we define and subsequently think about concepts such as fairness, equitable, bias, impartial, unprejudiced, etc. In this respect, semantic accuracy provides a reliable substrate for thinking more effectively, congruently, and realistically about the world we live in. This suggests that the fairest representation of the world is the most accurate one. Accurate semantic usage allows us to think and speak reliably about real relationships in the real world and provides us with a robust platform for open-ended problem solving and adaptation. With this understanding, semantic fairness offers credible resources for improving global communication, cooperation, and humane interaction.

17. DISCUSSION

From a scientific perspective, this approach allows for using natural human language to more accurately evaluate objects, behaviors, and cause-effect relationships within an appropriate awareness of the situation at a given moment, i.e. time and place. The contention is that it may be possible to generate pragmatic algorithms, including executive policy algorithms to model human cooperation, from computer science areas like artificial neural networks, artificial intelligence, artificial life, evolutionary algorithms, optimization, evolutionary computation, multiagent systems and games, etc.

Inspired by a yearning to simply understand the human brain and desires to develop a computer or robot with human brain-like capabilities — the field of computer science has witnessed a large expansion in research and practical applications. The quest for obtaining human brain-like capability dates back at least to Alan Turing's famous test for thinking machines (Turing, 1950). This is no small challenge considering the biological complexity of the human brain's estimated 100 trillion connections (Stanley et al., 2009 in press). In the fields of artificial intelligence, artificial life, and multi-agent systems and games, we have seen a dramatic expansion and evolution of experiments progressing towards a solution to this complex problem. There continue to be many language-based experiments in these areas producing fruitful results i.e. from artificial intelligence and conversational human-computer interfaces using human language (Allen et al., 2001), to artificial life environments with interactive human-like robots and agents (Nolfi & Floreano, 2000), to individual and group behaviors involved in multi-agent systems and games (Ostrowski et al., 2002), and to evolutionary computation.

Evolutionary computation maps the genome to the phenotype (behavior). This provides a powerful platform for exploring the development of new explanations for cognitive phenomena, testing the internal consistency and completeness of theories, testing the conditions under which phenomena do or do not occur, exploring the emergent properties of systems of agents in which the behavior of the system cannot be deduced from the behavior of a single agent, and deepening our understanding of what is, against a background of what might have been, and the study of the evolution of language and its emergent properties (Wiles & Hallinan, 2003).

Human language is often an overlying theme of these experiments, because language structure and function relates to human behavior and the interaction of complex adaptive systems involved in learning, culture, and biological evolution (Kirby, 2002). Since human speech is a behavior, we can evaluate language representation along with other human behaviors. This provides a natural conceptual test bed for potentially illuminating experiments using linguistic features to simulate human interactions. Even when these computer domains and experiments are not directly related to human language, they are usually at least indirectly related because programmers typically use some form of natural language when writing interactive programs in higher level object-oriented programming languages.

Useful knowledge could be gleaned from experiments evaluating fairness using multiple groups of agents. Fairness between agents seems to relate to cooperation with the concordance between the internal rule set agents follow for their behavior and their external rules for their interaction with other agents. We might compare the rule set agents employ for their own behavior, with the rule set those agents expect other agents to use. When we see relative concordance between what an agent says, how an agent behaves, and the behaviors the agent expects from others, we would deem that the rules were relatively fair. The assumption would be that when each agent in a group follows similar standards for cooperation, the group as a whole achieves, at least a perception of fairness. In theory, this defines a platform that values accuracy and reliability for communicating, negotiating, sharing, and cooperating. This fundamental platform provides useful representations for defining, evolving, and optimizing fair solutions, among agents, cultures, groups of cultures and their domain.

Further experiments might explore the relationship between the relative accuracy of language usage by agents and the resulting effect on predictability in their domain. Accurate communication between agents offers a vehicle for achieving harmony and cooperation among individuals and groups, by increasing the achievable level of empathy. This would imply a requirement for an accuracy-based reference point for the evaluation of language. This reference point has previously been introduced as cognitive accuracy. The rules for cognitive accuracy (Table 2) have been proposed in a prior manuscript: information accuracy, information processing accuracy, and event-level accuracy (Bailey, 2007).

Cognitive accuracy assumes the acceptance of human imperfection, contends that thoughts cause or significantly influence feelings and behavior, and notes that the use of open-ended flexible multivariate terms is generally more effective in a dynamic uncertain world than closed-ended rigid dichotomous terms. These rules allow for a more accurate classification of humans as Homo sapiens (knowing) while relying on scientific methods and critical thinking as tools for maximizing accurate thought and communication. This idealized scientific reference point offers potential as a reliable tool for achieving more predictable cooperation and perceived fairness, and seems to yield variables worthy of consideration and measurement. The arguments put forth so far suggest that fairness is a relative term. The degree to which an act is perceived as fair depends in part on the perspective of those doing the measuring. Relative fairness in effect depends on the reference point we use for measurement. In this light accuracy seems to provide a robust reference point providing better predictability. “What you get from a measurement depends on what you choose to measure” (Lindley, 2007).

Table 2. Cognitive Accuracy

- *Information accuracy (classes):* seeking and using objective information based on empirical observation; premise, deduction, conclusions, and testing

- *Thought process accuracy (methods):* making evaluations and decisions flexibly with critical thinking, multivariate terminology, and awareness of individual responsibility and accountability

- *Event-level accuracy (methods):* connecting and verifying both information and decisions in a time- and context-dependent manner to increase the relative probability of more accurate predictions of reasonable outcomes

- *Acceptance of human imperfection to enhance information (classes) and process (methods) accuracy.* We can accurately classify humans as flawed and fallible. Accepting our own flaws and fallibilities encourages us to accept others as human beings, albeit with flaws. This acceptance promotes horizontal, human-human, adult-adult, collateral communication. It also reduces inaccurate, absolute, dichotomous, or culturally biased classifications, ratings, and labeling, because we do not believe that any person is all bad or all good. Instead, we rate the behavior (method) rather than the person (class). Human acceptance minimizes inaccurate classifications, judgmental categories and cultural bigotry, while promoting realistic scientific belief systems, classes and methods, supported by healthy skepticism

- *Flexibility (methods) to enhance information and process accuracy.* Flexibility generally works better than rigidity as a method for the most accurate planning, problem solving, cooperating, and compromising in a dynamic uncertain world. Rigid, dichotomous, culturally determined terms like should, must, have to, got to, and need to, restrict options and diversity, while multivariate, preferential terms, such as I would prefer, I would rather, and to me, this seems best, multiply the diversity of possible choices and acceptable outcomes. Opinions and preferences replace absolute declarations of right and wrong. Generalities stated as assumptions or deductions represent a higher level of accuracy than generalities misrepresented as true or false facts when they actually represent subjective beliefs. Such generalizations amplify inaccuracies. Moreover, we improve the precision of our thought processes, decisions, and communications by using the most accurate word definitions, making specific rather than vague statements, using multivariate spectra or gradients for evaluations and revaluations, and avoiding faulty cause-effect conclusions (Browne & Keeley, 2007)

• *Awareness of the relationship between thoughts and emotions to enhance information, process, and event-level accuracy (representation, classes, and methods).* Normally, cognition has a significant causal relationship with our feelings, and realizing this enhances accurate assessment of individual responsibility and accountability for thoughts, feelings, and behaviors. Our thoughts, or the implicit rules behind our thoughts, cause or significantly influence our feelings whether we recognize the connection or not. Awareness of this relationship enables us to choose the healthiest and most reasonable thoughts to maximize our emotional and behavioral balance at a given time. Although we might initially react to the situation itself, we largely generate and sustain our emotional reactions to events by what we think or "believe" about them (Ochsner, 2006). We tend to sustain the emotion long afterwards through the action of implicit internal rules and appraisal habits that affect us almost continuously, generally without our awareness or deliberate direction. As Epictetus wrote in the Enchiridion almost 2000 years ago, "People are disturbed not by things, but by the views which they take of them" (Ellis & Harper, 1997)

(Modified from Bailey, 2007; Bailey, 2006)

18. CONCLUSIONS

Though this approach to accurately understanding humans may seem somewhat novel, it merits consideration. Having the ability to generate novel ideas and the option of adding an accuracy based reference point can lead to valuable new solutions when confronted with difficult and complex problems. This is especially so when searching with predetermined objectives fails to reasonably solve the problem. Objective search uses a prescribed objective and then searches for a solution by discovering and using stepping stones along the way to connect the dots between problem and solution. Unfortunately, objective searches can often lead to dead ends and deception, perhaps due to our subjective human tendency to find only what we are looking for i.e. confirmation bias. And unfortunately, once we know or think we know something, we tend to stop searching. Certainty closes the door on further discovery — it stops the search dead in its tracks.

This suggests that closed-ended thinking produces closed-mindedness due to the constraints of closed-ended *knowing* and an undue sense of *certainty* — thereby thwarting the search for better solutions. On the other hand, open-ended novelty search is a recent breakthrough in computer science that exploits uncertainty. Novelty search suggests an appealing alternative or complement to objective search, involving the development of new and unfamiliar solutions to problems without postulating *a priori* what the best solution might look like. With no pretention to know the solution, this open-ended approach keeps the door open for discovery. Research in machine learning shows that open-ended novelty search avoids deception and outperforms objective searches for finding difficult solutions (Lehman & Stanley, 2008). Perhaps counterintuitively, the more difficult the problem, the more novelty search outperforms objective search.

By dynamically and preferentially searching for novelty, open-ended novelty search mimics evolution's preference for variety and variability when searching complex environments. This allows for the discovery of newer and more robust solutions. Over the long run, a greater variety of potential solutions can be created by open-minded exploration for new ideas. Novelty search fairness provides a model for open-mindedness. Open-minded thinking mitigates the human tendency to get stuck in familiar ways of problem solving by providing alternatives — not putting all of our eggs into one basket, so to speak. This open-ended approach to search for novelty can conceivably discover more effective solutions over time. As humans, having novelty search in our tool box keeps us open to new options. This can best be accomplished by preferentially tempering what we know or think we know about our dynamic uncertain world with a touch of probabilistic thinking, and at least a pinch of healthy skepticism about the relative certainty of our knowledge.

With this in mind, we can see that probabilistic uncertainty promotes ongoing learning and provides a reasonable scientific interface with the natural world. It seems there is something to learn from computer science. This may represent an opportunity for discovering, validating, and sustaining a robust solution for fairness, empathy, and cooperation. In this sense, computer science supports the tools for, and a practical approach to resolving the dilemma of incongruent local cultural belief systems, rules, values, and biases as they relate to fairness, cooperation, and harmony. Perhaps defining robust fairness globally — in terms of optimal, dynamic, equitable, and cognitively accurate understanding of the problems we face — will lead to more predictable and equitable solutions. A novel approach with a reference point of cognitive accuracy may, in fact, enhance the probability of achieving this goal by providing an explicit, valid, reliable, and robust dynamic fitness function with global usefulness. If this is indeed the case, as scientist, the final step of ensuring the application in the natural world rest in our hands.

ACKNOWLEDGEMENTS

A special thanks to: Nora Miller for her invaluable editorial skills; Dr. Kenneth Stanley for taking time out of his busy schedule to share his insightful comments and suggestions; Joel Lehman and the EPLEX computer science group at the University of Central Florida for their thoughtful conversation; and especially Amy Hoover for her very patient tutoring.

REFERENCES

Adolphs, R. (2006). What is special about social cognition? In J. T. Cacioppo, P. S. Visser, & C. L. Pickett (Eds.), Social neuroscience, people thinking about people (pp. 269-285). Cambridge, MA: MIT Press.

Allen, J. F., Byron, D. K., Dzikovska, M., Fergonson, G., Galescu, L. & Stent, A. (2001). Towards conversational human-computer interaction. AI Magazine, 22, 4.

Bailey, C. E. (2006). A general theory of psychological relativity and cognitive evolution. ETC: A Review of General Semantics, 63:278-289.

Bailey, C. E. (2007). Cognitive accuracy and intelligent executive function in the brain and in business. In C. Senior & M. J. R. Butler (Eds.), The social cognitive neuroscience of organizations. New York, NY: Annals of the New York Academy of Sciences, 1118:122-141.

Bailey, C. E. (2007). Semantically mediated integration of cognition in Homo sapiens: Evolution, grammar, uncertainty, and cognitive accuracy. Cognitive Sciences 3 (1): 85-142.

Beck, A. T. (1999). Prisoners of hate: The cognitive basis of anger, hostility, and violence. New York, NY: HarperCollins.

Beck, A. T. (1976). Cognitive therapy and the emotional disorders. New York, NY: International University Press.

Bloom, P. (2002). How children learn the meaning of words. Cambridge, MA: MIT Press.

Browne, M. N. & Keeley, S. M. (2007). Asking the right questions: A guide to critical thinking (8th ED.). (pp. 119, 147) Upper Saddle River, NJ: Pearson Prentice Hall.

Cacioppo, J. T. & Berntson, G. G. (2004). Social neuroscience. In M. S. Gazzaniga (Ed.), The cognitive neurosciences III (pp. 977-985). Cambridge, MA: MIT Press.

Chow, T. W. & Cummings, J. L. (2007). Frontal-Subcortical circuits. In B. L. Miller & J. L. Cummings (Eds.), The human frontal lobes, 2nd Ed. (pp. 25-43). New York, NY: Guilford Press.

Cummings, J. L., & Miller, B. L. (2007). Conceptual and clinical aspects of the frontal lobes. In B. L. Miller & J. L. Cummings (Eds.), The human frontal lobes, 2nd Ed. (pp. 12-21). New York, NY: Guilford Press.

Damasio, A. R. (2000). A second chance for emotion. In R. D. Lane & L. Nadel (Eds.), Cognitive neuroscience of emotion (pp. 12-23). New York, NY: Oxford University Press.

Decety, J. (2007). A social cognitive neuroscience model of human empathy. In E. Harmon-Jones & P. Winkielman (Eds.), Social neuroscience: Integrating biological explanations of social behavior (pp. 246-270). New York, NY: Guilford Press.

De Jong, K. A. (2006). Evolutionary computation: A unified approach (pp. 232-3). Cambridge, MA: MIT Press.

Ellis, A. & Harper R.A. (1997). A guide to rational living (3rd Ed.). (pp. 39, 69) N. Hollywood, CA: Melvin Powers Wilshire Book Company (Original work published 1976).

Fehr, E. & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. The Quarterly Journal of Economics, 114, (3): 817-868.

Fuster, J. M. (2003). Cortex and mind, unifying cognition. New York, NY: Oxford University Press.

Gazzaley, A. & D'Esposito, M. (2007). Unifying prefrontal cortex function: Executive control, neural networks, and top-down modulation. In B. L. Miller & J. L. Cummings (Eds.), The human frontal lobes, 2nd Ed. (pp. 187-206). New York, NY: Guilford Press.

Gruau, F., Whitley, D. & Pyeatt (1996). A comparison between cellular encoding and direct encoding for genetic neural networks. In J. R. Koza, D. E. Goldberg, D. B. Fogel, & R. L. Riolo (Eds.), Genetic Programming 1996: Proceedings of the first annual conference. Cambridge, MA: MIT Press.

Kaufman, S. A. (1993). The origins of order: Self-Organization and selection in evolution. (p. 403) New York, NY: Oxford University Press.

Kemmer, S. & Barlow, M. (2000). Introduction: A usage-based Conception of language. In Michael Barlow & Suzanne Kemmer (eds.), Usage Based Models of Language. Stanford, CA: CSLI Publications.

Keller, H & Chasiotis, A. (2006). Evolutionary perspectives on social engagement. In P. T. Marshall & N. A. Fox (Eds.), The development of social engagement: Neurobiological perspectives (pp. 275-303). New York, NY: Oxford University Press.

Kirby, S. (2002). Natural language from artificial life. ALIFE. 8 (2): 185-218. Cambridge, MA: MIT Press.

LeDoux, J. (2002). Synaptic Self: How our brains become who we are. (pp. 198, 203) New York, NY: Russell Sage Foundation.

LeDoux, J. (1996). The emotional brain. New York, NY: Simon & Shuster.

Lehman, J. & Stanley, K. (2008). Exploiting Open-Endedness to Solve Problems Through the Search for Novelty. In Bullock, S., J. Noble, R. Watson, and M. A. Bedau (eds.) (2008). Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems. Cambridge, MA. MIT Press.

Lindley, D. (2007). Uncertainty: Einstein, Heisenberg, Bohr, and the struggle for the soul of science. (p. 155) New York, NY: Doubleday.

Logothetis, N. K. (2004). Higher cognitive functions. In M. S. Gazzaniga (Ed.), The cognitive neurosciences III (pp. 849-969). Cambridge, MA: MIT Press.

Luria, A. R. (1981). Language and cognition. (pp. 1-13, 205-9) J. V. Wertsch (Ed.).Washington, DC: John Wiley & Sons Inc.

MacLean, P. D. (1990). The triune brain in evolution. New York, NY: Plenum Press.

Mesulam, M.-Marsel (2002). The human frontal lobes: Transcending the default mode through contingent coding. In D. T. Stuss & R. T. Knight (Eds.), Principles of frontal lobe function (pp. 8-30). New York, NY: Oxford University Press.

Mesulam, M.-Marsel (2000). Principles of behavioral and cognitive neurology (2nd Ed.). (p. 93) New York, NY: Oxford University Press.

Mesulam, M.-Marsel (1985). Patterns in behavioral neuroanatomy: Association areas, the limbic system, and hemispheric specialization. In M. Mesulam (Ed.) Principles of Behavioral Neurology (pp. 1-70). Philadelphia, PA: F. A Davis.

Milgram, S. (2004). Obedience to authority (first published 1974). New York, NY: HarperCollins.

Milner, B, & Petrides, M. (1984). Behavioral effects of frontal-lobe lesions in man. Trends in Neuroscience, 7:403-407.

Minsky, M. L. (2006). The emotion machine: Commonsense thinking, artificial intelligence, and the future of the human mind. (p. 6) New York, NY: Simon & Schuster.

Mitchell, J. P., Macrae, C. N. & Banaji, M. R. (2006). Dissociable medial prefrontal contributions to judgments of similar and dissimilar others. Neuron, 50:655-663.

Nolfi, S. & Floreano, D. (2000). Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines. Cambridge, MA: MIT Press.

Norris, C. J. & Cacioppo, J. T. (2007). I know how you feel: Social and emotional information processing in the brain. In E. Harmon-Jones & P. Winkielman (Eds.), Social neuroscience: Integrating biological explanations of social behavior (pp. 84-105). New York, NY: Guilford Press.

Nowak, M. A. (2006). Evolutionary dynamics: Exploring the equations of life. (pp. 249-52) Cambridge, MA: Belknap Harvard University Press.

Ochsner, K. N. (2007). How thinking controls feelings: A cognitive neuroscience approach. In E. Harmon-Jones & P. Winkielman (Eds.), Social neuroscience: Integrating biological explanations of social behavior (pp. 106-133). New York, NY: Guilford Press.

Ochsner, K. N. (2006). Characterizing the functional architecture of affect regulation: Emerging answers and outstanding questions. In J. T. Cacioppo, P. S. Visser & C. L. Pickett (Eds.), Social neuroscience, people thinking about people (pp. 245-268). Cambridge, MA: MIT Press.

Ochsner, K. N., Bunge, S. A., Gross, J. J. & Gabrieli, J. D. E. (2005). Rethinking feelings: An fMRI study of the cognitive regulation of emotion. In J. T. Cacioppo & G. G. Berntson (Eds.), Social neuroscience: Key readings in social psychology (pp. 253-270). New York, NY: Psychology Press.

Ostrowski, D. A., Tassier, T., Everson, M. p. & Reynolds, R. G. (2002). Using cultural algorithms to evolve strategies in agent-based models. In D. B. Fogel, M. A. El-Sharkawi, X. Yao, G. Greenwood, H. Iba, P. Marrow & M. Shackleton (Eds.), Proceedings of the 2002 Congress on evolutionary computation CEC2002, pp. 741-746. IEEE Press.

Panksepp, J. (1998). Affective Neuroscience: The foundations of human and animal emotions. (p. 20) New York, NY: Oxford University Press.

Phelps, E. A. (2004). The human amygdala and awareness: Interactions between emotion and cognition. In M. S. Gazzaniga (Ed.), The cognitive neurosciences III (pp. 1005-1015). Cambridge, MA: MIT Press.

Phelps, E. A. & LaBar, K. S. (2006). Functional neuroimaging of emotion and social cognition. In R. Cabeza & A. Kingstone (Eds.), Handbook of functional neuroimaging of cognition (pp. 421- 453). Cambridge, MA: MIT Press.

Risberg, J. (2006). Evolutionary aspects on the frontal lobes. In J. Risberg & J. Grafman (Eds.), The frontal lobes, development, function and pathology (pp. 1-20). New York, NY: Cambridge University Press.

Rogers, S. K., Kabrisky, M., Bauer, K. & Oxley, M. E. (2003). Computing machinery and intelligence amplification. In D. B. Fogel & C. J. Robinson (Eds.) Computational Intelligence: The experts speak (p. 25). NJ: IEEE Press.

Singer, T., Seymour, B., O'Doherty, J. P., Stephan, K. E., Dolan, R. J. & Frith, C. D. (2006). Empathetic neural responses are modulated by the perceived fairness of others. Letters, Nature, 439 (1): 466-469.

Stanley, K. O., D' Ambrosio, D. & Gauci, J. (2009). A hypercube-based indirect encoding for evolving large-scale neural networks. Artificial Life Journal, accepted for 2009 publication.

Striedter, G. F. (2005). Principles of brain evolution. (pp. 322, 344) Sunderland, MA: Sinauer Associates.

Tomasello, M. (2005). Constructing a language: A usage-based theory of language acquisition. Cambridge, MA: Harvard University Press.

Turing, A.M. (1950). Computing machinery and intelligence. Mind, 59 (236): 433-460.

von Dassow, G. & Meir, E. (2004). Exploring modularity with dynamical models of gene networks. In Gerhard Schlosser & Gunter P. Wagner (Eds.). Modularity in Development and Evolution. (pp. 256-8) Chicago; University of Chicago Press.

Wagner, A. (2005). Robustness and evolvability in living systems. Princeton, NJ: Princeton University Press.

Wiles, J, & Hallinan, J. (2003). Evolutionary computation and cognitive science. In D. B. Fogel & C. J. Robinson (Eds.), Computational intelligence: The experts speak. (pp. 183-188) Piscataway, NJ: IEEE Press.

Wolpert, D. H. & Macready, W. G. (1995). No free lunch theorems for search. Technical Report SFI-TR-95-02-010. Santa Fe, New Mexico. Santa Fe Institute.

Wolpert, D. H. & Macready, W. G. (1997). No free lunch theorems for optimization. IEEE Transactions on Evolutionary Computation, 1:67-82.

Wright, J. H. (2004). In J. H. Wright (Ed.), Cognitive behavioral therapy. Washington, DC: American Psychiatric Publishing.