

Neuroscientific Insights on Biblical Myths

SIMPLIFYING HEURISTICS VERSUS CAREFUL
THINKING: SCIENTIFIC ANALYSIS OF MILLENNIAL
SPIRITUAL ISSUES

by Daniel S. Levine and Leonid I. Perlovsky

Abstract. There is ample evidence that humans (and other primates) possess a knowledge instinct—a biologically driven impulse to make coherent sense of the world at the highest level possible. Yet behavioral decision-making data suggest a contrary biological drive to minimize cognitive effort by solving problems using simplifying heuristics. Individuals differ, and the same person varies over time, in the strength of the knowledge instinct. Neuroimaging studies suggest which brain regions might mediate the balance between knowledge expansion and heuristic simplification. One region implicated in primary emotional experience is more activated in individuals who use primitive heuristics, whereas two areas of the cortex are more activated in individuals with a strong knowledge drive: one region implicated in detecting risk or conflict and another implicated in generating creative ideas. Knowledge maximization and effort minimization are both evolutionary adaptations, and both are valuable in different contexts. Effort minimization helps us make minor and routine decisions efficiently, whereas knowledge maximization connects us to the beautiful, to the sublime, and to our highest aspirations. We relate the opposition between the knowledge instinct and heuristics to the biblical story of the fall, and argue that the causal scientific worldview is mathematically equivalent to teleological arguments from final causes. Elements of a scientific program are formulated to address unresolved issues.

Daniel S. Levine is a professor in the Department of Psychology, University of Texas at Arlington, Arlington, TX 76019-0528; e-mail levine@uta.edu. Leonid I. Perlovsky is Visiting Scholar at Harvard University and technical advisor for the Air Force Research Laboratory, SN, Hanscom AFB, MA 01731; e-mail leonid@seas.harvard.edu.

[*Zygon*, vol. 43, no. 4 (December 2008)]

© 2008 by the Joint Publication Board of *Zygon*. ISSN 0591-2385

Keywords: amygdala; anterior cingulate cortex; beautiful; biblical story of the fall; brain; causality; cognitive science; creativity; decision making; dorsolateral prefrontal cortex; effort minimization; emotions; evolutionary adaptations; frontal lobes; heuristics; knowledge instinct; original sin; psychology; risk; neural networks; sublime; teleology

IS THERE A DRIVE TO KNOW?

Leonid Perlovsky (2001; 2006a) has proposed that humans possess a knowledge instinct (KI)—that is, a drive to create realistic internal models of the world. This involves developing mental representations that are as consistent as possible across different hierarchical levels of the brain.

The idea of a KI is controversial because it seems quite different from standard biological instincts such as eating, sex, and aggression that human beings share with other animals. Also, the KI runs counter to the long-established biological principle of effort minimization (EM).

The tradeoff between the KI and effort minimization is considered in the ensuing sections of this article. Yet clearly the KI has the survival value of leading us to generate as accurate an internal representation of the world as possible. At the same time, the KI is closely involved with important aspects of the quality and meaningfulness of life beyond mere survival, such as play, curiosity, and beauty.

What is the scientific evidence for the existence of a KI? The brain systems involved in such an instinct have been but dimly suggested. Yet there is a wide range of results both in experimental psychology and neuroscience that support the overall notion.

Many psychological studies demonstrate that both humans and monkeys possess an exploratory drive, a curiosity about their environments regardless of the instrumental value of the knowledge they acquire about it. This has been known at least since the work of Harry Harlow in the middle of the last century, who showed that monkeys will learn to solve various kinds of mechanical puzzles when no external reinforcement is provided other than the puzzle itself (Harlow 1953).

Also, it is well established that humans seek to reduce cognitive dissonance (Festinger 1957; Harmon-Jones and Mills 1999). For example, people often experience dissonance when they discover inconsistency between their own behaviors and professed beliefs, and this dissonance can generate physiological signs of distress, such as a skin conductance response. The dissonance can be reduced in several ways: changing behaviors, changing beliefs, or finding a higher-level synthesis of the two (Abelson 1959; Hardyck and Kardush 1968; Harmon-Jones et al. 1996). In support of the notion of a knowledge instinct, there are data suggesting that people tend to find changing beliefs less satisfying than synthesis. Social psychologists Roger Elkin

and Michael Leippe (1986) showed that in some cases such attitude change does not relieve the physiological discomfort caused by the dissonance, suggesting that the attitude change has not been internalized at an unconscious level. We return to their work later with a suggestion for a further experimental test.

The KI is what drives both intellectual and aesthetic creativity (Perlovsky 2001; 2006a). We believe it also is essential for the difficult interpersonal negotiations necessary for creation of a humane society, one in which each individual considers others' interests as well as his or her own.

Yet there also is abundant evidence that humans, like most biological organisms, do not think or act optimally all the time (see Levine and Elsberry 1997 for a partial review). In particular, we often do not make decisions using all of the information available to us. We tend to employ heuristics (shortcuts or rules of thumb) that enable us to employ minimal cognitive effort in decision making. Heuristics typically use a small but seemingly important part of the available information in combination with previously developed beliefs. Often they result in solutions that may not be the best but are adequate, minimize cognitive effort, and are made quickly.

Does the use of heuristics run counter to the KI? Or are both useful in different situations? We next review some of the psychological data about decision-making heuristics. Then we draw some analogies between these different ways of thinking (the KI and EM) and aspects of a biblical story. We then discuss what brain processes might be involved in both knowledge maximization and effort minimization and also discuss some possible neural bases for individual differences in the tendency to employ the knowledge-maximizing versus the effort-minimizing mode of thinking.

HEURISTICS OR EFFORT MINIMIZATION

From the work of the pioneering eighteenth-century mathematicians Jakob Bernoulli and Thomas Bayes through the late twentieth century, the dominant notion in the psychology of human decision making was based on rational optimization. The belief was that each decision maker had an internal, and self-consistent, *subjective utility function* and made all choices involving risk by choosing the alternative for which the mathematical expectation of utility was the largest. But all that changed with the work, starting in the late 1960s, of Daniel Kahneman, winner of the 2003 Nobel prize in economics, and Amos Tversky, who would have shared that prize had he been alive (Tversky and Kahneman 1974; 1981).

Tversky and Kahneman found that in many choices relating to gain and loss estimation, preferences run counter to rational optimization and lack self-consistency over different linguistic framings of the choice. For example, subjects asked to consider two programs to combat an Asian disease expected to kill 600 people tend to prefer the certain saving of 200

lives to a 1/3 probability of saving all 600 with 2/3 probability of saving none. However, subjects also tend to prefer a 1/3 probability of nobody dying with a 2/3 probability of 600 dying to the certainty of 400 dying. The choices are identical in actual effect but are perceived differently because of differences in frame of reference (comparing hypothetical states in one case with the state of all being alive, in the other case with the state of all dying). Tversky and Kahneman explain their data by noting that “choices involving gains are often risk averse while choices involving losses are often risk taking” (Tversky and Kahneman 1981, 453).

Tversky and Kahneman were led to descriptions of a large repertoire of simplifying heuristics that decision makers characteristically employ. Three of these heuristics that are now widely recognized are *availability* (selective enhancement of memory representations that come to mind easily), *representativeness* (deciding on the likelihood of category membership based on similarity of descriptive characteristics to those of the category), and *anchoring and adjustment* (estimating numerical quantities by starting on a suggested “anchor” value and then going up or down from that).

Also it has been shown that in multiattribute decision making, humans frequently do not weigh all the attributes in proportion to their importance. Rather, they tend to simplify their search of alternatives by focusing on one or two major attributes, and this tendency is actually *increased* by the presence of information about a large number of other attributes (Payne, Bettman, and Johnson 1993).

Most psychologists believe, however, that heuristics have evolutionary value despite sometimes leading to errors and information loss. Heuristic simplification is particularly useful when a decision must be made rapidly on incomplete information, or when the stakes of the decision are not high enough to justify the effort of thorough deliberation. An example is buying a box of cereal in a supermarket (Levine 1997).

The use of heuristics such as availability, representativeness, and anchoring and adjustment has been explained as minimization of cognitive effort (Montgomery and Svenson 1976). How do humans, other animals, or artificial neural systems decide when to search for maximum knowledge and when to simplify? And how do individuals, normal or pathological, differ in their criteria for choosing knowledge or simplification?

ADAM AND EVE—KI AND HEURISTICS

The origin of the controversy between the KI and heuristics can be traced to the first pages of the Bible, to the story of Adam and Eve. In the twelfth century Moses Maimonides in his “Guide for the Perplexed” (Maimonides [1190] 1956) analyzed the relationship between KI and heuristics. He was asked by his student: “Why did God, on one hand, give Adam the mind and free will, while on the other, forbid him to eat of the tree of knowl-

edge? Did God not want Adam to use his mind?” Maimonides answered that God gave Adam the mind to think for himself what is good and what is bad (we associate this ability with the KI). But Adam succumbed to temptation and ate from the tree of knowledge, thereby taking a “short-cut” and acquiring ready-made heuristics, that is, rule-of-thumb knowledge to guide him so his choices did not require hard thinking. Maimonides explained that Adam’s story describes our predicament. Whereas God’s ultimate commandment is to use the KI, it is difficult and we are not completely capable of doing it, especially when thinking about the highest values. Adam’s story describes the workings of our mind: struggle between the KI and EM. EM provides the surety of millennial cultural support, but it may not suit our individual circumstances. The KI may lead to doubts and uncertainties, but if successfully used it leads to the satisfaction of being more conscious about our choices.

Maimonides’s interpretation of the biblical story adds another dimension to the previously discussed differences between the KI and EM. Mathematically, it is possible to formulate a mind’s utility function so that the KI and EM are brought close to each other. This utility function can account for the survival value of quick decisions and also for the limited amount of any individual experience, for uncertainty in observation of data, and for minimizing the worst-case losses (such as preventing death) versus maximizing average gain. The utility function even can account for the fact that the future is unknown and therefore individual experience should be integrated with culturally accumulated knowledge. But Maimonides hints at something different, something more fundamental than correct formulation of a utility function. He suggests that “original sin” determining the basic imperfectness of humankind is related to how we do or do not use our ability for knowledge and for making conscious choices.

We now relate this biblical account to theoretical understanding and experimental data about the mind and brain, consider neural mechanisms of the KI and EM, relate them to conceptual and emotional intelligence and to consciousness and the unconscious, and try to add scientific interpretation to this millennial old mystery.

MECHANISMS OF THE MIND

The basic mind mechanisms making up operations of the KI have been described mathematically by Perlovsky (2006a, b). Here we give a conceptual summary of this description, which is compatible with a large number of experimental findings in cognitive science and psychology. The mind is organized as a heterarchy of multiple levels; it is not a strict hierarchy because of feedback connections between levels (but we refer to a hierarchy for simplicity). Perception of simple objects occurs at lower levels. Higher levels are responsible for cognition of general and abstract concepts, which

unify lower-level knowledge, and at the top of the hierarchy there are concepts unifying our entire knowledge. As we argue later, these are concepts of the meaning and purpose of life and existence.

At every level, mechanisms of perception and cognition involve interactions between top-down and bottom-up signals. Top-down signals are generated by representations or mental models based partly on memories stored at the higher levels. Bottom-up signals are generated by models recognized and activated at a lower level; at the lowest level these come from sensory organs. There are many levels between the retina of the eye and regions of the cortex that recognize objects, but for simplicity we refer to object recognition as a single level, as if our eyes perceive objects due to object-models. Object perception occurs when an object-model matches a subset of signals in the retina corresponding to the object. But stored memories never exactly match new sensory signals because “the same” object is never really the same; angles, lighting, and surroundings are always different than in the past. Therefore memories-models always have to be modified to match signals. The matching mechanism is driven by the KI. The knowledge instinct therefore is necessary for every act of perception.

When computers appeared in the 1950s, mathematicians thought that soon they would be able to create artificial intelligence far surpassing the human one. It did not happen. Why? The first attempts at creating intelligent computers assumed that matching models to sensory data consisted of selecting the correct model from a store of memories. Many years of mathematical analysis led to a different conclusion: Stored representations-memories-models are not exact but vague, so that they approximately match various objects in a variety of conditions. The KI drives the process of matching models to signals; in this process models become crisp and clear and resemble the objects. Mathematically the KI maximizes a measure of similarity between the models and signals. Normally we are unconscious of this process and about vague models. We become conscious only about the final result, the crisp models matching sensory signals; this is called the *resonance* (between the model and signals). Still, it is easy to experience directly the vague memories-models. Imagine a familiar object with closed eyes. This imagination is vaguer than a crisp perception of a model-matched-to-object. In summary, the KI evolves¹ vague and less conscious models into crisper and more conscious models that better match objects in bottom-up signals.

This process of modifying models and matching them to bottom-up signals is repeated at higher levels of the mind, where complex cognition takes place. Abstract general models are like internal eyes of the mind² that “perceive” general concepts among lower levels’ bottom-up signals. This process is amazingly close to the idea of homunculus often discussed in prescientific literature about the mind. The homunculus was conceived as a little mind inside our mind that perceived our perceptions and made

them available to our mind. The fundamental difference is that the scientific explanation does not need an infinite chain of homunculi inside homunculi. Instead, there is a hierarchy of the mind models with their conscious and unconscious aspects; recent neural-network theories and brain imaging results are pointing the way toward a biological description of these models. The conscious differentiated aspect of the models decreases at higher levels in the hierarchy; they are less certain and more vague. At the top of the hierarchy there are vague and mostly unconscious models of the meaning of our existence.

The KI is manifest in emotional responses to matches and mismatches at every level of perception and cognition. Stephen Grossberg and Daniel Levine (1987) developed a neural-network model in which satisfaction or dissatisfaction of instincts is felt as positive or negative emotions. The word *emotions* is used for several related but different processes, including physiological responses, cognitive appraisals, affective communications, and neural signals. We use it here for neural signals that connect recognition and decision-making brain regions with instinctual regions; they indicate to instinctual mechanisms potential satisfaction or dissatisfaction of instinctual needs. For example, a food object-model will send positive emotional signals to the eating instinctual region of a hungry person; the food object will receive a priority in the recognition process.

How do we feel satisfaction of the KI? At lower levels of object perception, emotional signals satisfying the KI usually are below threshold of conscious registration. However, if objects around do not correspond to our models, we immediately feel an intense disharmony (between reality and our expectations). Thriller movies exploit this property of our perception, showing objects and situations that do not correspond to our models-expectations, thereby creating strong negative emotions. Perceptions at lower levels are not much different in this respect from other autonomous functions. For example, we do not consciously perceive the working of our stomach mechanisms as long as everything is okay. But we become acutely emotionally aware when our stomach does not function properly. At higher levels of cognition, we feel both positive and negative KI-related emotions consciously. For example, we feel happiness when we solve a complex problem that has occupied our mind for days and frustration when we are unable to solve it.

Emotions related to higher levels of the KI involve processes of thinking and understanding more than bodily activities such as sex or eating do. For this reason and this reason alone they can be called *spiritual*.³ Since Kant's work these spiritual emotions have been called *aesthetic* emotions. We emphasize that these emotions are not reserved for artistic activities but accompany every act of perception and cognition. At every level the KI drives the learning of model-concepts to become crisper and more conscious. But at the higher levels, the model-concepts are intrinsically more

vague and more difficult to make conscious; hence, more cognitive effort is required for understanding them, for adapting them to life, and for making them more conscious.

Correspondingly, we feel more aesthetic emotions when we succeed in this. When the highest models of the meaning and purpose of existence are adapted and made more conscious, we feel the presence of beauty. Nothing in the world around us can be directly perceived as giving purpose and meaning to our lives. In fact, random circumstances often seem to destroy purposiveness. Nevertheless, to be able to endure life, to concentrate will, and to achieve satisfaction, we have to believe that meaning exists. And from time to time we experience this fleeting, but so dear, emotion of the beautiful, when the KI is satisfied at the highest level.

The mechanism of sublime feelings is similar to beauty. Whereas the beautiful is related to the highest models of cognition, the feeling of the sublime is related to the highest models of behavior. A significant part of behavior is similar to cognition in that it is governed by models. We understand how to act by matching models of behavior to circumstances. Understanding the meaning and purpose is not sufficient; we also would like to realize this understanding in our lives. To do this we need to develop corresponding models of behavior. These models are vague and uncertain, like the corresponding models of cognition. We do not know for sure what the best way is to achieve the highest meaning and purpose in our lives. (There is an opposite point of view, the belief that one *can* know and can choose at will the best course of behavior. In the fourth century this was called the Pelagian heresy by the Christian church.)

There is a mathematical reason why choices of the beautiful and sublime cannot be made crisp, clear, and completely conscious: These choices require evaluation and selection from infinite sets. Recognition of a simple object, as discussed, requires matching the object model to a subset of signals originating from the object. This subset has to be selected among many other subsets. There are about 10,000 signals that each eye retina receives ten times per second. Therefore perception requires thousands and millions of signals to be matched to models of thousands of objects. But let us forget for a moment about these large numbers and consider a choice of just 100 object-models to be matched to subsets of 100 signals. The number of these subsets is 100^{100} . This number is larger than the sum of all interactions among all elementary particles in the entire life of the Universe (Penrose 1989, 326–45). Thus, choices of the beautiful and sublime require evaluation of a physically infinite number of subsets (and therefore involve an infinite amount of information). Perception of objects around us is helped by these objects actually being there, but cognition of abstract high-level concepts does not have such firm grounds.

Each of us has definite values that limit our search among all these signals from the environment and from our own minds. Hence, models can

partly illuminate our choices of the beautiful and sublime. But our mathematical argument shows that our highest aesthetic aspirations cannot be fully reduced to neat computational formulas. Thus, modern science supports insights from many religions that ultimately the choices of beautiful and sublime are beyond clear and conscious human reach. This resonates with Maimonides's belief that we must strive for the choice of the beautiful and sublime, because we are given the KI but cannot completely fulfill this demand especially in complex matters of important life choices. We are bound to resort sometimes to shortcuts, heuristics. We have to use collective wisdom accumulated over millennia.

This conclusion does not resolve the matter of original sin, however. What is so sinful about not being able to do the impossible? Why did the early Christian church declare the quite reasonable teachings of Pelagius to be heretical? Why is Maimonides's interpretation of the Adam and Eve story still controversial? Why is Friedrich Nietzsche's concept of the *superman*—the ideal man who will be able to decide what is good and bad by thinking for himself (Nietzsche [1885] 1999)—often scorned or misunderstood?

Such questions dramatize the constant tension in human cultures between maximizing knowledge in order to enhance human potential and restricting knowledge in order to preserve known or comfortable social relationships. All of culture, including religion, is produced by our brains acting in society. Hence, we can obtain some insights into the sources of this tension, and indeed of original sin itself, from theories about how the brain is organized to perform cognitive functions.

We develop here a theory of brain organization that encompasses both the knowledge instinct and effort minimization, as well as differences between individuals in the tendency toward either KI or EM. In order to understand this complex and intricate set of phenomena we must search for what evolutionary biologist Edward Wilson (1998) calls *consilience*: a set of principles and system understandings that is self-consistent across different disciplines and the different levels of understanding those disciplines foster. Because the common language of mathematics often enhances the possibility of this kind of interdisciplinary consistency, we turn now to mathematical neural-network models to sharpen our understanding about the KI/EM tradeoff.

NEURAL-NETWORK THEORIES

What are neural networks exactly? They have become a widely used computational modeling tool in neuroscience and experimental psychology as well as a device for engineering applications of intelligent systems. Yet they have been surrounded by a certain amount of hype and misconception. The term *neural network* does not yet have a universally accepted definition. Perhaps the closest is the following:

. . . a neural network is a system composed of many simple processing elements operating in parallel whose function is determined by network structure, connection strengths, and the processing performed at computing elements or *nodes*. . . . Neural network architectures are inspired by the architecture of biological nervous systems, which use . . . processing elements operating in parallel. (DARPA 1988; emphasis added)

Neural-network nodes, the elements in network diagrams, need not be interpreted as single neurons but have *activities* that are idealized frequencies of neuron firing. The connections between nodes have *weights* that are idealized strengths of synapses between different neurons. The ultimate aim is to make these networks as biologically realistic as possible. Sometimes nodes correspond to brain areas or specific cell types in those brain areas. At other times, when not enough is known about brain processes or when modeling at a functional level is desired, nodes correspond to cognitive entities such as the memory of a specific word, the tendency to approach a specific object, or the intensity of a specific drive or emotion.

The history of computational neural networks as used in cognitive and behavioral neuroscience, as they have gradually moved toward greater biological realism as more data has become available, is reviewed elsewhere (Levine 2002; 2007a). One of the best mathematical theories for interactions between multiple processing levels in the brain is *adaptive resonance theory (ART)*, developed by Gail Carpenter and Stephen Grossberg. In particular, ART has been widely used as a network for linking together representations of categories and of their attributes. A very brief review of ART follows. Much more detail appears in Grossberg 1999.

In its simplest form (Figure 1), the ART network consists of two interconnected layers of nodes (neurons), called F_1 and F_2 . F_1 (“bottom”) is assumed to consist of nodes that respond to input features. F_2 (“top”) is assumed to consist of nodes that respond to categories of F_1 node activity patterns. Learning takes place at both bottom-up and top-down connections (synapses) between the two layers.

The F_1 nodes do not directly interact with each other, but the F_2 nodes competitively inhibit each other’s signals. Such competition is a common device in neural networks for making choices in short-term memory. In this version, the simplest form of choice (winner take all) is made—only the F_2 node receiving the largest signal from F_1 becomes active. To compute the signal received by a given F_2 node, the activity of each F_1 node in response to the input pattern is weighted by the strength of the bottom-up synapses from that F_1 node to the given F_2 node, and all these weighted activities are added.

Inhibition from the F_2 level to the F_1 level shuts off most neural activity at F_1 if there is mismatch between the input pattern and the active category’s prototype. Only with a sufficiently large match are enough of the same F_1 nodes excited by both the input and the active F_2 category node, which is needed to overcome nonspecific inhibition from F_2 .

If match occurs, F_1 activity is large because many nodes are simultaneously excited by input and prototype. If mismatch occurs, F_2 reset shuts off the active category node as long as the current input is present. The criterion for matching is that some function representing the degree of match between top-down and bottom-up patterns must be greater than some positive constant r , which is called the *vigilance* of the network.

How is the degree of abstractness or generalization controlled? In the ART model, a low value of the vigilance parameter makes the network learn broad categories, whereas high vigilance makes the network learn more specific categories.

In Carpenter and Grossberg 1987 the vigilance r is a uniform parameter of the network. It is expressed as a gain of a signal from F_1 to the orienting node. This can be adapted easily to attribute-selective vigilance (see Leven and Levine 1996) by setting different values of the gain from each node at F_1 representing a different stimulus attribute. Also, multiple ART modules can be combined into hierarchical networks with three or more levels of abstraction (Carpenter and Grossberg 1990). The top level of one ART module can be the bottom level of another ART module. Also, an ART network can be extended to include modules representing relationships between other modules, as Nilendu Jani and Levine (2000) did in modeling the learning of analogies.

A neural architecture developed by Perlovsky (Perlovsky and McManus 1991; Perlovsky 2001), *Neural Modeling Fields (NMF)*, is similar to ART in many respects. NMF can be considered an extension of ART toward

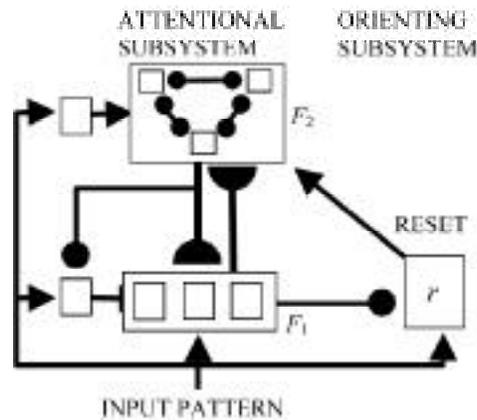


Fig. 1. ART 1 architecture. Short-term memory is encoded at the feature level F_1 and category level F_2 , and learning at interlevel synapses. The orienting system generates reset of F_2 when bottom-up and top-down patterns mismatch at F_1 , that is, when the ratio of F_2 activity to input activity is less than the vigilance r . Arrows denote excitation, filled circles inhibition, and semicircles learning. (Adapted from Carpenter and Grossberg 1987, with permission of Academic Press.)

ideas of the KI, concept-models, and mental hierarchy. NMF uses complex models at the higher level F_2 . The original state of these models is vague and fuzzy. The KI drives the process of model improvement, model adaptation to patterns in bottom-up signals. NMF uses a mathematical measure of similarity between top-down signals coming from models in F_2 and bottom-up signals in F_1 . Mathematically, the KI maximizes the similarity. Aesthetic emotions are changes in this similarity measure and take active part in adapting models to signals, making models more clear, crisp, and conscious. Models that were successfully matched to patterns in bottom-up signals achieve a resonant state, a high degree of certainty, and are available to consciousness. These models are activated and send signals to the next level of more abstract and general concepts. This architecture is shown schematically in Figure 2.

Mathematical description of the operations of the KI within the hierarchical structure of the mind is a matter of ongoing and future research (Perlovsky 2004; 2006a, b, d). Here we outline some details beyond those discussed in the previous section. It is sometimes convenient to refer to the

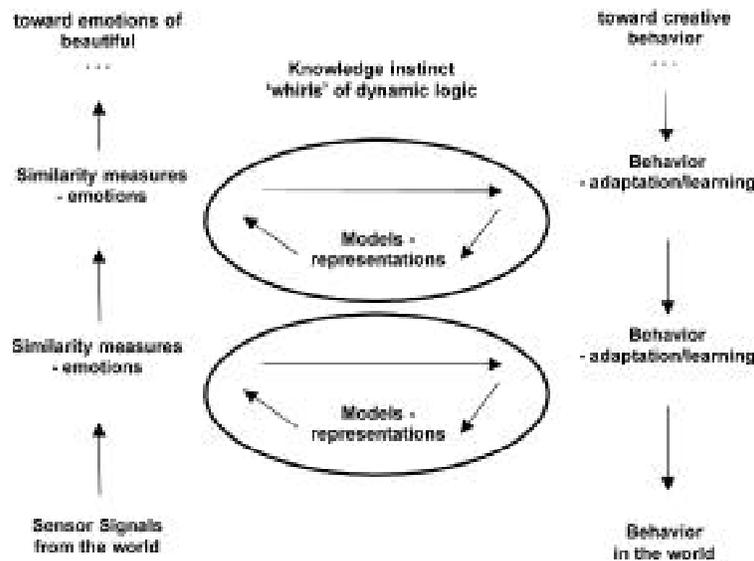


Fig. 2. Hierarchical NMF system. At each level of a hierarchy there are models, similarity measures, and actions (including adaptation, maximizing the knowledge instinct - similarity). High levels of partial similarity measures correspond to concepts recognized at a given level. Concept activations are output signals at this level and they become input signals to the next level, propagating knowledge up the hierarchy. Each concept-model finds its mental meaning and purpose at a higher level. (Signal flows leading to adaptation of behavioral models are not shown; they are similar to cognitive models).

two neuronal population levels in Figure 1 as a single processing level, where the KI improves and creates new models, which better match patterns in the bottom-up signals. This process of creating more specific, more diverse models increases specificity and diversity of knowledge. Following Carl Jung, we call this process *differentiation*.

At a single level, the purpose⁴ of each model is to satisfy the knowledge instinct by finding patterns in the bottom-up signals and adapting to these patterns. There also are meanings and purposes related to bodily instincts; for example, food objects can be used to satisfy needs for food and desires for eating. In this essay we limit our discussion to spiritual⁵ needs, to the knowledge instinct.

Models acquire deeper meanings and purposes at higher hierarchical levels. The pure aesthetic feeling of harmony between our knowledge and the surrounding world at lower levels, as discussed, is below the threshold of conscious registration in our minds. We do not feel much joy from understanding simple objects around us. But we do enjoy solving complex problems that required a lot of time and effort. This emotional feel of harmony from improving/creating high-level concept-models is related to the fact that high-level concepts unify many lower-level concepts and increase the overall meaning and purpose of our diverse knowledge. Following Jung, we call this process *synthesis* (Jung [1922] 1971). Jung emphasized that synthesis is essential for psychological well-being.

Synthesis, the feel of overall meaning and purpose of knowledge, is related to the meaning and purpose of life, which we perceive at the highest levels of the hierarchy of the mind. At those high levels, as discussed, models are intrinsically vague and undifferentiated, not only in terms of their conceptual content but also in terms of differentiation between the conceptual and emotional. At the highest levels of the mind the two are not quite separable. This inseparability, which we sometimes feel as a meaning and purpose of our existence, is important for evolution and survival. If the hierarchy of knowledge did not support this feeling, the entire hierarchy would crumble, which was an important (possibly the most important) mechanism for collapses of former civilizations. The KI demands satisfaction at the lowest levels of understanding concrete objects in the environment and also at the highest levels of the mind hierarchy, understanding of the entire knowledge in its unity, which we feel as meaning and purpose of our existence. Hence the KI paradoxically drives us toward both differentiation and synthesis (Perlovsky 2006c).

Neuroscience provides suggestions for where vague model-concepts come from. At lower levels of the mind hierarchy, vague models rely on inborn properties of the brain (Grossberg 1988; Zeki 1993); their differentiation and adaptation are grounded in patterns in sensory signals. These mechanisms of perception we share with animals. But where do abstract and general models originate? Clearly, babies are not born with abstract ideas

or cognitive concepts in their minds such as those of rationality, abstractness, or purpose. These ideas also cannot be formed by selecting useful combinations of simpler ideas or concepts, as we do with objects that are directly perceived in the surrounding world. As discussed earlier, there are too many combinations of simpler ideas/concepts; evaluating them and selecting the useful ones is too slow a process. The evolution of language enabled human beings to form and select abstract, high-level ideas in ways that are inaccessible to other animals. In our minds there are two parallel and interacting hierarchies for cognition and for language, as illustrated in Figure 3.

Ability for learning higher levels of the hierarchy is closely related to ability for language. The reason is that otherwise there is no ground for learning cognitive models—there are no abstract concepts that could be directly perceived in the world. The only ground for learning abstract cognitive concepts is language concepts, which are learned from surrounding language and culture at many hierarchical levels. In the integrated NMF system, abstract cognitive models at higher levels in the hierarchy are grounded in abstract language models.

The double hierarchy of Figure 3 integrates language and cognition. A data stream constantly comes into the mind from all sensory perceptions; every part of this data stream is constantly evaluated and associated with cognitive and language models. At the beginning, the models are vague, fuzzy; cognitive models vaguely correspond to uncertain undifferentiated sensory perceptions. Language models vaguely correspond to sounds. This is approximately the state of the mind of a newborn baby. First, models of

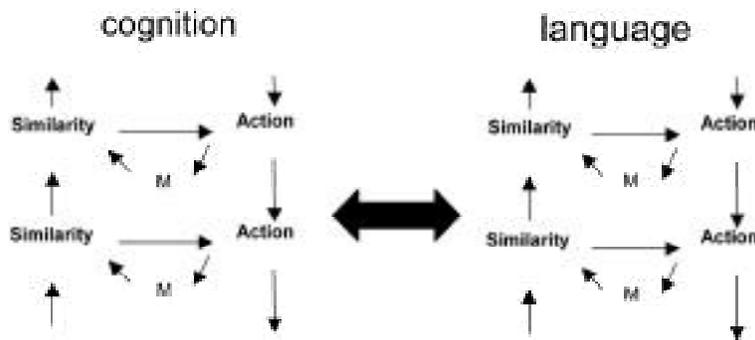


Fig. 3. Hierarchical integrated language-cognition NMF system. At each level in a hierarchy there are integrated language and cognition models. Initial models are fuzzy placeholders, so integration of language and cognition is subconscious. Associations between signals and models depend on language, cognitive models, and signals. Therefore language model learning helps cognitive model learning and vice versa. Abstract cognitive concepts are grounded in abstract language concepts.

simple perceptions differentiate; objects are distinguished in visual perception. Language sounds are differentiated from other sounds. Until about one year of age, perception models corresponding to simple objects become crisper at a faster rate than language models. Between the first and second year of life the speed of adaptation of language models tremendously accelerates and overtakes learning of cognitive models. By the age of 5 or 7, a child knows a tremendous number of language models (words, rules of grammar), which have attained differentiated, crisp status. At 5 or 7 a child can talk about much of the entire cultural content. But it will take the rest of his life to associate language models with cognitive models, adapt them to real-life situations, and acquire highly differentiated crisp cognitive models.

Language models are acquired from the surrounding culture ready-made, differentiated, without any effort. This is so because of a special inborn mechanism in the mind, which Steven Pinker called the language instinct (1995). But the language instinct does not connect language models to reality. Language models provide just a starting point for developing cognitive models. Cognitive models have to be adapted to the world. Cultural knowledge is acquired when cognitive models attain the same differentiation as language models differentiated in culture. This process, as we discussed, is driven by the KI and involves aesthetic emotions.

Language models are crisp and conscious, and as long as we speak without thinking, like children, we may perceive ourselves as conscious and in control of our thoughts and life. But when we apply knowledge contained in language to real life, we may experience uncertainty in our thoughts and behavior corresponding to the vague and unconscious state of our cognitive models. For example, every 5-year-old knows about good guys and bad guys, but who can claim at 40 or 70 that he or she applies this knowledge in his life without error? Philosophers and theologians have argued about good and evil for thousands of years, and these arguments are likely to continue forever. Crisp and clear language models give heuristics for use in real life. Thinking for oneself requires using the KI to develop cognitive models that are adapted to one's own life experience (ideally, including the entire cultural heritage).

Returning to differentiation and synthesis, emotional investment in each concept decreases with an increase in the number of concepts, and a drive for differentiation and creating more concepts subsides. Emotional investment in a concept is a measure of the meaning and purpose of this concept within the mind system—that is, a measure of synthesis. Thus, the drive for differentiation requires synthesis. Synthesis leads to differentiation, whereas differentiation decreases synthesis. The KI demands both differentiation and synthesis, but these mechanisms contradict each other. Here, along with the dichotomy between KI and EM, may be another root of the fundamental contradiction in human nature that some call original sin.

People differ from one another in their ability to connect language and cognition. Many individuals are good at talking, without fully understanding how their language concepts are related to real life. On any subject, they can talk one way or another without much emotional investment. Yet successful synthesis of language and cognition involves synthesis of the emotional and conceptual contents of the psyche. The other side of these relationships is too much synthesis. High-value concepts (related to family life, or political cause, or religion) are so important to us and so emotional that we cannot coldly analyze and thereby differentiate them. Too high a level of synthesis invests concepts with too much emotional value, so that differentiation is stifled (Perlovsky 2007). People resort to heuristics, to ready-made knowledge acquired with language.

In summary, the choice between increase of knowledge and minimization of cognitive effort, between the KI and EM, Maimonides connected to original sin. The Bible identifies it as the fallen condition of humankind, the source of the world's miseries. Buddhism sees the source of human unhappiness as *tanha*, loosely translated as "desire" or "attachment" (Smith 1958) but, from a scientific perspective, meaning self-absorbed deficiency-based emotions leading to overreliance on EM heuristics. The KI involves individual effort for increasing knowledge and aesthetic emotions; at the highest levels of the mind hierarchy it involves the beautiful and sublime. It also involves the conscious and the unconscious, the conceptual and the emotional, language and thinking. There is a difference between the "fallen," bodily emotions involved in EM, in using language without thinking, and aesthetic emotions related to the KI.

What are the relationships between the high and low, between aesthetic emotions of the beautiful and sublime, and primitive emotions of minimizing cognitive effort? And what are the roles of these mechanisms in the entire system of the mind? Answering these questions requires dedicated scientific programs, and some of these are underway. The previous part of this article demonstrated that these questions can be formulated scientifically. In the next section we make a step toward more detailed formulation of such a scientific program and relate theoretical considerations of the mind mechanisms to psychological experiments and to brain regions, both those that we share with other animals and those uniquely human.

RELEVANT BRAIN REGIONS

We turn now to some cognitive neuroscience data that do not fully answer the questions about KI versus EM but begin to address them. After a review we suggest a few experiments yet undone that may further increase our understanding of knowledge maximization versus effort minimization, between "high" and "low."

The tendency that Tversky and Kahneman (1974; 1981) found for decision makers to be susceptible to linguistic framing is not universal. A

significant minority of adults (along with most young children) are not susceptible to the distortions of rational decision making caused by framing effects.

A functional magnetic resonance imaging (fMRI) study (DeMartino et al. 2006) showed significant differences in brain region activation between individuals who were and were not susceptible to framing effects. This study used a monetary analog of Tversky and Kahneman's "Asian disease" problem. Subjects had to choose between a sure option and a gamble option, where the sure option was expressed either in terms of gains (keep £20 out of the £50 they initially received) or in terms of losses (lose £30 out of the initial £50).

As in the Asian-disease problem, the majority of subjects chose the sure option with a gain frame and the gamble option with a loss frame. Yet there were significant minorities of subjects who chose the gamble with a gain frame or the sure option with a loss frame, in violation of the usual heuristics. fMRI measurements showed that the heuristics violators had more activation than the heuristics followers in two major areas of the frontal lobes and adjacent cortex—the orbitofrontal cortex (OFC) and anterior cingulate cortex (ACC). Conversely, those subjects whose choices were consistent with the framing heuristic had more activation in the amygdala, the area below the cortex that is most involved with primary emotional experience.

The OFC, ACC, and a third prefrontal region, the dorsolateral prefrontal cortex (DLPFC), are part of what is called the executive system of the brain (Pribram and Luria 1973). Some of the functions of these regions are reviewed in Levine 2007b and are summarized here as follows.

Damage to the OFC in humans often leads to decision-making deficits and socially inappropriate behavior, as in the famous nineteenth-century patient Phineas Gage (Damasio 1994). These clinical observations, together with animal lesion studies, suggest that the OFC forms and sustains mental linkages between specific sensory events, or motor actions, and positive or negative emotional states. Long-term storage of positive or negative emotional values is likely to be at connections between the OFC and amygdala (Schoenbaum et al. 2003).

The DLPFC is a working-memory region and is more closely connected with the hippocampus, the site of memory consolidation, than with the amygdala. It is involved in information processing at a higher level of abstraction than the OFC. For example, OFC lesions in monkeys were found to impair learning of changes in reward value within a stimulus dimension, whereas DLPFC lesion impaired learning of changes in which dimension was relevant (Dias, Robbins, and Roberts 1996).

The ACC is activated when a subject must select or switch among different interpretations or aspects of a stimulus (Posner and Petersen 1990). Also, in an attentional task with emotional distractors, ACC was the unique

area whose activation increased to both target (task-relevant) and distracting stimuli (Yamasaki, LaBar, and McCarthy 2002). Recent theories of ACC function have emphasized its role in detection either of potential response error or of conflict between signals promoting competing responses (Botvinick et al. 2001; Brown and Braver 2005).

Hence, executive regions of the cortex are more readily activated when knowledge motivations are engaged than when simplifying heuristics are employed. This is closely related to the long-established distinction in cognitive psychology between controlled and automatic processing (Schneider and Shiffrin 1977). What is the variable that changes between these two modes? The data suggest that the interplay between the two modes can be studied by means of some neural-network parameter that varies both between individuals and between domains in the life of the same individual. Such a parameter is vigilance, used in adaptive resonance theory (Carpenter and Grossberg 1987).

Figure 4 shows a three-level hierarchical ART network for knowledge encoding and processing. This figure should be taken as a skeletal diagram of relevant processes rather than a complete neural network for those processes that will take years to develop. In the network of Figure 4, simple heuristics involve feedback between the amygdala and OFC and do not

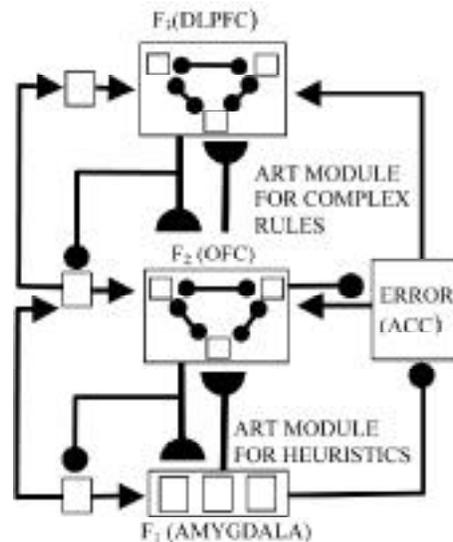


Figure 4. Proposed multilevel brain network that can either minimize effort or maximize knowledge. With low vigilance, the ART module combining F_1 and F_2 (amygdala and OFC) makes decisions based on simple heuristics. With high vigilance, discontent with outcomes of simple decisions generates activity in the orienting (error) module (ACC). ACC activity in turn may generate a search for more complex decision rules at F_3 (DLPFC).

engage the other two prefrontal executive areas (the error detector at ACC and the complex working memory analyzer at DLPFC). The individual with higher vigilance in the pursuit of knowledge, whether this vigilance is general or domain-specific, is sensitive to mismatches between the results of those heuristics and logical truth. This leads in turn to engagement of the other executive regions.

In the fMRI experiment (DeMartino et al. 2006), the more vigilant subjects had increased neural activity in OFC and ACC, but not as much in DLPFC as Figure 4 might suggest. The framing-influenced monetary decision task, however, may not be complex or abstract enough to engage the DLPFC. One partial test of our theory would be an fMRI study of better versus worse decision makers on a more complex decision task, one involving competition between different abstract principles, which therefore is more likely to engage complex working memory manipulations attributed to the DLPFC. An example relates to confusion between frequency and probability (Pacini and Epstein 1999), where subjects are shown two bowls containing red and white jelly beans and need to choose which bowl gives them the best chance of randomly choosing a red jelly bean when the ratios are, say, 8 out of 100 in one bowl and 1 out of 10 in another bowl. Our prediction is that those who correctly choose the higher probability (1 out of 10) will show greater DLPFC activation than those who incorrectly choose the higher frequency. One of us (Levine) is currently collaborating with Daniel Krawczyk of the University of Texas Southwestern Medical School on running brain imaging experiments to test that prediction. Also, more mathematical modeling needs to be done to specify the variables that actually correspond to activations of specific regions and how those relate to variables in the dynamical system equations that define model networks such as ART.

What might be a possible neural mechanism for aesthetic emotions? If the highest levels of processing are engaged, we propose that the error detector in the ACC of Figure 4 could become acutely sensitive to matches and mismatches of actual events with both cognitive and emotional expectations. The anterior cingulate traditionally was considered part of the limbic system, the network of regions below the cortex (including the amygdala) that is involved in emotional expression (MacLean 1990). The ACC is known to be involved in contributing to the subjective pain perception, with the intensity of pain correlating with ACC activation (Posner and DiGirolamo 2000). The connections of the anterior cingulate portion of the brain's executive system with more primitive subcortical and viscerally connected emotional areas (in the limbic system, hypothalamus, and mid-brain) means that aesthetic emotions have a partial common substrate with the more primitive emotions related to survival.

The higher up we go in the hierarchy, the closer we are to the beautiful and sublime, the easier, it seems, to succumb to the temptation to stop

thinking and to use ready-made concepts acquired from the culture: language concepts connected not to individual thinking but to Mom-and-Dad prohibitions, to amygdalar emotions triggered by previous failures, when we tried to think and got burned. Figure 4 should be understood as a scheme for structuring each one of the hierarchical levels in Figure 2.

The loops in the brain involving the anterior cingulate along with deeper areas of the “visceral brain” (MacLean 1990) provide an explanation for the close kinship between the most complex products of the knowledge instinct (art and science, for example) and our basic emotional sensibilities. They also help to explain experimental data showing that cognitive dissonance can lead to physiological signs of emotional arousal, even when that dissonance has no long-term consequences for the subject or anyone else (Harmon-Jones et al. 1996).

This suggests another possible behavioral experiment to test our theory. Elkin and Leippe (1986) noted that subjects given the choice to write an essay that disagreed with their previously held beliefs, and yet writing the essay anyway, often changed their beliefs to relieve the dissonance between their beliefs and actions, but still showed physiological signs of arousal or discomfort. What would happen if instead of giving the subjects an opportunity to change their beliefs, we gave them a chance to arrive at a mental synthesis between their beliefs and their behavior (for example, to tell themselves that their beliefs have not changed but they enjoyed exploring arguments in favor of the opposite position)? Would that eliminate the physiological discomfort, or even promote positive affect from enjoying the intellectual challenge? Our hypothesis is that mental synthesis would provide physiological comfort, at least in subjects who are high in their enjoyment of effortful cognitive activity, which in experimental psychology has been formalized by a construct called *need for cognition* (Cacioppo et al. 1996).

CAUSALITY VERSUS TELEOLOGY

Teleology explains the universe in terms of purposes. This usually is associated with religious purpose, as it suggests an ultimate Designer must exist. Therefore, teleology is a hot point of debates between creationists and evolutionists. Is there a purpose in the world? Evolutionists assume that the only explanation is causal. Newton’s laws gave a perfect causal explanation for the motion of planets: A planet moves from moment to moment under the influence of a gravitational force. Similarly, today science explains motions of all particles and quantum fields according to causal laws, and there are exact mathematical expressions for fields, forces, and their motions. Causality explains what happens in the next moment as a result of forces acting in the previous moment. Science is associated with causal explanation and opposed to teleological explanations in terms of purposes.

The very basis of science, it seems, is on the side of causality, and religion is on the side of teleology.

However, the contradiction between causality and teleology does not exist at the basic level of fundamental physics. The laws of physics, from classical Newtonian laws to quantum superstrings, can be formulated equally as causal or as teleological. An example of a teleological principle in physics is energy minimization. This posits that particles in each moment “know” their purpose: to move so as to minimize energy. The most general physical laws, including those governing causal dynamics, motions of particles, quantum strings, and superstrings, are formulated as minimization of a mathematical expression called the *Lagrangian* (Feynman and Hibbs 1965). A particle under force moves from point to point as if it knows its final purpose, to minimize the Lagrangian.

The KI is mathematically similar to these general physical laws; evolution of the mind is guided by maximization of knowledge. For the first time in a complex system, the mind, teleological principle is mathematically equivalent to causal dynamics. Cognitive effort minimization might interfere from time to time, but still evolution of the mind proceeds toward more knowledge. Models of human decision making went from one extreme of the artificial-intelligence models of the 1960s that included long-term planning but not emotion, to the other extreme of emotion-based models from the 1980s and 1990s that dealt only with short-term EM heuristics. Only recently, as more cognitive data about brain regions has become available, have decision models started to integrate long-term planning and emotional influences (Funahashi, Lee, and Rushworth 2006; Levine 2006).

Both the KI and Lagrangian result in equivalency of causal dynamics and purposive dynamics. We would like to emphasize that a general scientific assumption is that all biological laws ultimately are founded in physics. This, however, does not guarantee equivalency of causality and purpose. This equivalency exists in physics only for elementary interactions (of a few particles). For complex systems, *statistical* physics rules. It gives rise to the second law of thermodynamics, stating that less probable states evolve into more probable states—that is, entropy always increases. This defines the “arrow of time.” Also, entropy increase may sound like a teleological principle, but it is not; dynamic causal laws cannot be inferred from entropy increase. According to statistical physics and entropy, the end state of the Universe is thermal equilibrium, “thermal death.” This conclusion prevailed in the nineteenth and the first part of the twentieth century, but contemporary physics suggests that when effects of gravity are considered, thermal death is not inevitable. Yet cosmological theory was not yet able to incorporate equivalency of causality and purpose.

In this regard the KI is a revolutionary principle. For the first time it states that for a very complex system, the human mind, causality and purpose

are equivalent. Instead of the rule of entropy, arrow of time, and thermal death, the human destiny is ruled by increase of knowledge. This may be a scientific interpretation of a mysterious biblical statement that time belongs to the “fallen world” of matter and that in the redeemed world “time will be no more” (Revelation 10:6 KJV).⁶ One does not have to choose between scientific explanation and teleological purpose; causal dynamics and purpose-driven dynamics (teleology) are mathematically equivalent.

CONCLUSIONS

Figure 4 is a first approximation to a network that can incorporate the differences in vigilance, and therefore in knowledge seeking, between individuals or between domains for the same individual. It is built on previously established models of simpler processes such as perception and categorization.

When Albert Einstein said “Seek simplicity and distrust it,” he was comparing experimental observations with the predictions of relatively simple rules for the physical universe. Einstein found mismatches with Newtonian theory in areas that seemed minor in the large scheme of the theory (such as black body radiation). Yet his level and selectivity of vigilance, in some sense, was high enough to generate a search for a major new theory based on those seemingly minor mismatches.

Within each of the brain areas discussed here, there are multiple layers of cortex, cell types, and chemical neurotransmitter systems not included in Figure 4. More realistic neural models of decision processes (see Brown, Bullock, and Grossberg 2004; Levine, Mills, and Estrada 2005) include such biological details. Yet the schema of Figure 4 appears to capture the gist of some basic processes important for knowledge networks to operate in a complex and nonstationary environment.

How can a biological or artificial network be trained to develop the level of vigilance appropriate to each task and context? Explaining differences in thinking between Einstein and other investigators in terms of a single parameter, vigilance, is a vast oversimplification; any theory always mismatches data in many ways. Deciding which mismatches are fundamental and which are not will require much more than a single number. This is a subject for further investigation. There are a variety of neural-network training methods already in the literature, and some may be adaptable to interhuman interactions such as psychotherapy, education, parenting, and mentoring. Practical solutions to this problem of training may require a differentiated KI, a theory that is now only being created. But the best answer is what Artur Rubinstein said, according to an anecdote, when a tourist in New York asked him how to get to Carnegie Hall: “Practice, practice, practice!”

NOTES

1. Throughout this paper, usually the words *evolution* and *evolving* refer to evolution in the mind. We rarely refer to genetic evolution. The reason is that with the emergence of language, cultural evolution greatly outpaced genetic evolution. Richard Dawkins (1976) writes in this regard about emergence of new replicators (words and ideas) more powerful than old replicators, genes.

2. This analogy between models and eyes is a literal description of the fact that mathematical mechanisms of perception of general concepts among lower-level concepts are the same as perception of visual features and objects in the retina of the eye. Contrary to this, the following analogy with a homunculus should not be taken literally for scientific explanation. It is given for those who appreciate analogies instead of exact mathematical or neural descriptions and can be ignored if this analogy does not help or is distracting. Those interested in exact scientific mathematical description of these mechanisms should consult Perlovsky 2006a, c and further references therein.

3. In this article we adhere to a scientific view that there is no difference in substance between matter and spirit. Our spiritual experiences are results of material processes in the brain. We emphasize that this is commensurate with monotheistic views and contrary to the dualistic position that spirit and matter are of different substances. A dualistic view cannot be reconciled with this essay or with contemporary science. Our attempt here to unify spirit and matter, body and mind takes into account all known scientific data in psychology, neurobiology, and mathematical mechanisms unifying experimental data.

4. Every element or mechanism within a living being is a result of long evolution and evolved for a purpose. Purpose and function are related but not exactly equivalent. Whereas purpose is related to the question Why? and to explanation from final causes, function is related to the question How? and to “mechanical” or dynamic causes. When we discuss specific neural mechanisms, we sometimes refer to functions, or mechanisms of functioning. Since Aristotle and Kant it has been well appreciated that understanding of living beings requires understanding of final causes. Scientific explanation requires unifying dynamic and final causes. We attempt such unification throughout this essay; in particular here we refer to a purpose in a relatively simple mechanism, where both causal and purposeful aspects are relatively easy to trace. We return to this question in the final section of the article.

5. We remind the reader that in this article *spiritual* refers to mechanisms of the brain (see note 3).

6. King James Bible translation: “there should be time no longer.” Many other authoritative translations: “time will be no more” (for example, Henry 2003). We would like to emphasize that except for this sentence, the content of this section is exact science. We do not think that every word of the Bible has to be amenable to exact scientific interpretation. Nevertheless, when mysterious intuitions of biblical authors have scientific parallels, pointing out those parallels can add to our understanding of the Bible. This is the reason for this sentence.

REFERENCES

- Abelson, Robert P. 1959. “Modes of Resolution of Belief Dilemmas.” *Conflict Resolution* 3:343–52.
- Botvinick, Matthew, Todd Braver, Deanna Barch, Cameron Carter, and Jonathan Cohen. 2001. “Conflict Monitoring and Cognitive Control.” *Psychological Review* 108:624–52.
- Brown, Joshua, and Todd Braver. 2005. “Learned Predictions of Error Likelihood in the Anterior Cingulate Cortex.” *Science* 307:1118–21.
- Brown, Joshua, Daniel Bullock, and Stephen Grossberg. 2004. “How Laminar Frontal Cortex and Basal Ganglia Circuits Interact to Control Planned and Reactive Saccades.” *Neural Networks* 17:471–510.
- Cacioppo, John, Richard Petty, J. A. Feinstein, and W. B. G. Jarvis. 1996. “Dispositional Differences in Cognitive Motivation: The Life and Times of Individuals Varying in the Need for Cognition.” *Psychological Bulletin* 119:197–253.
- Carpenter, Gail, and Stephen Grossberg. 1987. “A Massively Parallel Architecture for a Self-Organizing Neural Pattern Recognition Machine.” *Computer Vision, Graphics, and Image Processing* 37:54–115.

- . 1990. "ART 3: Hierarchical Search Using Chemical Transmitters in a Self-Organizing Pattern Recognition Architecture." *Neural Networks* 3:129–52.
- Damasio, Antonio. 1994. *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: Grosset/Putnam.
- DARPA (Defense Advanced Research Projects Agency). 1988. *Neural Network Study*. Alexandria, Va.: AFCEA International Press.
- Dawkins, Richard. 1976. *The Selfish Gene*. Oxford, U.K.: Oxford Univ. Press.
- DeMartino, Benedetto, Dharshan Kumaran, Ben Seymour, and Raymond Dolan. 2006. "Frames, Biases, and Rational Decision-Making in the Human Brain." *Science* 313:684–87.
- Dias, R., Trevor Robbins, and Angela Roberts. 1996. "Dissociation in Prefrontal Cortex of Affective and Attentional Shifts." *Nature* 380:69–72.
- Elkin, Roger, and Michael Leippe. 1986. "Physiological Arousal, Dissonance, and Attitude Change: Evidence for a Dissonance-Arousal Link and a 'Don't Remind Me' Effect." *Journal of Personality and Social Psychology* 51:55–65.
- Festinger, Leon. 1957. *A Theory of Cognitive Dissonance*. Evanston, Ill.: Row, Peterson.
- Feynman, Richard, and R. A. Hibbs. 1965. *Quantum Mechanics and Path Integrals*. New York: McGraw-Hill.
- Funahashi, Shintaro, Daeyeol Lee, and Matthew Rushworth, eds. 2006. Special issue on Neurobiology of Decision Making. *Neural Networks* 19(8).
- Grossberg, Stephen. 1988. *Neural Networks and Natural Intelligence*. Cambridge: MIT Press.
- . 1999. "The Link between Brain Learning, Attention, and Consciousness." *Consciousness and Cognition* 8:1–44.
- Grossberg, Stephen, and Daniel Levine. 1987. "Neural Dynamics of Attentionally Modulated Pavlovian Conditioning: Blocking, Interstimulus Interval, and Secondary Reinforcement." *Applied Optics* 26:5015–30.
- Hardyck, Jane A., and Marcella Kardush. 1968. "A Modest Modish Model of Dissonance Reduction." In *Theories of Cognitive Consistency*, ed. Robert P. Abelson, Elliot Aronson, William J. McGuire, Theodore M. Newcomb, Milton J. Rosenberg, and Percy H. Tannenbaum, 684–92. Chicago: Rand McNally.
- Harlow, Harry. 1953. "Mice, Monkeys, Men, and Motives." *Psychological Review* 60:23–32.
- Harmon-Jones, Eddie, Jack Brehm, Jeff Greenberg, Linda Simon, and David E. Nelson. 1996. "Evidence that the Production of Aversive Consequences is Not Necessary to Create Cognitive Dissonance." *Journal of Personality and Social Psychology* 70:5–16.
- Harmon-Jones, Eddie, and Judson Mills, eds. 1999. *Cognitive Dissonance: Progress on a Pivotal Theory in Social Psychology*. Washington, D.C.: American Psychological Association.
- Henry, Matthew. 2003. *Concise Commentary*. Nashville: Thomas Nelson.
- Jani, Nilendu, and Daniel Levine. 2000. "A Neural Network Model of Proportional Analogy-Making." *Neural Networks* 13:149–83.
- Jung, Carl. [1922] 1971. *Psychological Types*. Collected Works, vol. 6, Bollingen Series XX. Princeton, N.J.: Princeton Univ. Press.
- Leven, Samuel, and Daniel Levine. 1996. "Multiattribute Decision Making in Context: A Dynamic Neural Network Methodology." *Cognitive Science* 20:279–99.
- Levine, Daniel. 1997. "Don't Just Stand There, Optimize Something!" In *Optimality in Biological and Artificial Networks?*, ed. Daniel S. Levine and Wesley R. Elsberry, 3–18. Mahwah, N.J.: Erlbaum.
- . 2002. "Neural Network Modeling." In *Stevens' Handbook of Experimental Psychology*, Volume 4: *Methodology*, ed. John Wixted et al., 223–69. New York: Wiley.
- . 2006. "Emotion and Decision Making: Short-term Reactions versus Long-term Evaluations." *Proceedings of International Joint Conference on Neural Networks*, July, 195–202.
- . 2007a. "Neural Network Modeling of Emotion." *Physics of Life Reviews* 4:37–63.
- . 2007b. "How Does the Brain Create, Change, and Selectively Override Its Rules of Conduct?" In *Neurodynamics of Higher-level Cognition and Consciousness*, ed. Robert Kozma and Leonid Perlovsky, 163–81. Heidelberg: Springer-Verlag.
- Levine, Daniel, and Wesley Elsberry, eds. 1997. *Optimality in Biological and Artificial Networks?* Mahwah, N.J.: Erlbaum.

- Levine, Daniel, Britain Mills, and Steven Estrada. 2005. "Modeling Emotional Influences on Human Decision Making under Risk." *Proceedings of International Joint Conference on Neural Networks*, August.
- MacLean, Paul. 1990. *The Triune Brain in Evolution: Role in Paleocerebral Functions*. New York: Plenum.
- Maimonides, Moses. [1190] 1956. *The Guide for the Perplexed*. Trans. M. Friedlander. Unabridged Dover Edition.
- Montgomery, Henry, and Ola Svenson. 1976. "On Decision Rules and Information Processing Strategies among Multiattribute Alternatives." *Scandinavian Journal of Psychology* 17:283–91.
- Nietzsche, Friedrich. [1885] 1999. *Thus Spake Zarathustra*. New York: Dover.
- Pacini, Rosemary, and Seymour Epstein. 1999. "The Interaction of Three Facets of Concrete Thinking in a Game of Chance." *Thinking and Reasoning* 5:303–25.
- Payne, John, James Bettman, and Eric Johnson. 1993. *The Adaptive Decision Maker*. Cambridge: Cambridge Univ. Press.
- Penrose, Roger. 1989. *The Emperor's New Mind*. Oxford: Oxford Univ. Press.
- Perlovsky, Leonid. 2001. *Neural Networks and Intellect: Using Model Based Concepts*. New York: Oxford Univ. Press.
- . 2004. "Integrating Language and Cognition." *IEEE [Institute of Electrical and Electronics Engineers] Connections* 2:8–12.
- . 2006a. "Toward Physics of the Mind: Concepts, Emotions, Consciousness, and Symbols." *Physics of Life Reviews* 3:23–55.
- . 2006b. "Neural Networks, Fuzzy Models and Dynamic Logic." In *Aspects of Automatic Text Analysis (Festschrift in Honor of Burghard Rieger)*, ed. R. Köhler and A. Mehler, 363–86. Heidelberg: Springer.
- . 2006c. "Modeling Field Theory of Higher Cognitive Functions." In *Artificial Cognition Systems*, ed. A. Loula, R. Gudwin, and J. Queiroz, 64–105. Hershey, Pa.: Idea Group.
- . 2006d. "Music—The First Principle." In Jury Dimitrin, Musical Theater e-journal, http://www.ceo.spb.ru/libretto/kon_lan/ogl.shtml.
- . 2007. "Neural Dynamic Logic of Consciousness: The Knowledge Instinct." In *Neurodynamics of Higher-Level Cognition and Consciousness*, ed. Robert Kozma and Leonid Perlovsky, 73–108. Heidelberg: Springer-Verlag.
- Perlovsky, Leonid, and M. Margaret McManus. 1991. "Maximum Likelihood Neural Networks for Sensor Fusion and Adaptive Classification." *Neural Networks* 4:89–102.
- Pinker, Steven. 1995. *The Language Instinct: How the Mind Creates Language*. New York: Harper Perennial.
- Posner, Michael, and Gregory J. DiGirolamo. 2000. "Attention in Cognitive Neuroscience: Overview." In *The New Cognitive Neurosciences*, 2d ed., ed. Michael Gazzaniga, 623–31. Cambridge: MIT Press.
- Posner, Michael, and Stephen Petersen. 1990. "The Attention System of the Human Brain." *Annual Review of Neuroscience* 13:25–42.
- Pribram, Karl, and Aleksandr Luria. 1973. *Psychophysiology of the Frontal Lobes*. Oxford: Academic.
- Schoenbaum, Geoffrey, Barry Setlow, Michael Saddoris, and Micaela Gallagher. 2003. "Encoding Predicted Outcome and Acquired Value in Orbitofrontal Cortex during Cue Sampling Depends upon Input from Basolateral Amygdala." *Neuron* 39:855–67.
- Schneider, Walter, and Richard Shiffrin. 1977. "Controlled and Automatic Human Information Processing: I. Detection, Search, and Attention." *Psychological Review* 84:1–66.
- Smith, Huston. 1958. *The Religions of Man*. New York: Harper & Brothers.
- Tversky, Amos, and Daniel Kahneman. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science* 185:1124–31.
- . 1981. "The Framing of Decisions and the Rationality of Choice." *Science* 211:453–58.
- Wilson, Edward. 1998. *Consilience*. New York: Knopf.
- Yamasaki, Hiroshi, Kevin LaBar, and Gregory McCarthy. 2002. "Dissociable Prefrontal Brain Systems for Attention and Emotion." *Proceedings of the National Academy of Sciences* 99:11447–51.
- Zeki, Semir. 1993. *A Vision of the Brain*. Oxford: Blackwell.

