RELEVANCE, MEANING AND THE COGNITIVE SCIENCE OF WISDOM

By

John Vervaeke and Leo Ferraro

Let us begin the study of wisdom by noting that it involves some kind of cognitive improvement that affords the living of a good life. When we use the term "cognition" or "cognitive", it should be broadly construed as the terms are used in cognitive science, meaning thinking, reasoning, memory, emotion and perception. There are factors such as good fortune that can improve life, but wisdom centres on a kind of self-transformation of cognitive processing that enhances the quality of life in some comprehensive manner. Philosophers (especially ancient philosophers) have devoted a lot of time to addressing the related questions of what is wisdom and what it is to live a good life. Recently psychologists have also broached the topic because of the central role of cognitive processes in wisdom (Sternberg 1990, Sternberg 200?, Brown 2000). Neuroscientists have also begun to explore the topic as they have forayed into explaining higher cognitive processes and wisdom seems to involve higher cognitive processes such as self-regulation and problem solving (Meeks and Jeste 2009; Goldberg 2005). It stands to reason that cognitive science, which attempts to create theoretical links between philosophical, psychological, and neuroscientific theories (by making use of information processing ideas drawn from the fields of machine learning and artificial intelligence), could have a lot to say about wisdom. In a sense, wisdom is a quintessential cognitive science topic: cognitive science offers both a diverse and integrated theoretical perspective that makes it uniquely suited to investigating and explicating a phenomenon as cognitively complex as wisdom.

In this chapter we will explore what cognitive science contributes to the understanding of the nature and development of wisdom by explicitly drawing upon philosophical, psychological and neuroscientific theories that are integrated through information processing/machine learning ideas. In doing so we expressly are not attempting some simple reduction of wisdom but rather we hope to enrich the construct in a way that will make it more empirically tractable. To do this we hope to tie our theoretical account of wisdom to independently established constructs within psychology, machine learning, and neuroscience, and to thereby provide a process theory that explains how one undergoes processes of self-transformation that results in wisdom, as well as explicating and specifying those processes to some degree.

Why a process theory? As an analogy, let us consider the study of epistemology: Chisholm (1982) distinguishes between two broad methods for studying knowledge. One is to define knowledge and

then to use that definition to formulate the question "How is knowledge acquired? Which psychological processes result in knowledge as their product?" The second method is to define a process by which knowledge is acquired and to use that procedural specification as the basis for a definition of what knowledge itself is. Chisholm's point was not to set these two methods in competition with each other, but rather to reveal their fundamental interdependence. Any definition of a product implicitly presupposes the process by which it is generated. Any theory of a process likewise relies upon an implicit concept of the resulting product. In coming to understand wisdom, simply knowing the features of wisdom is insufficient. Such an understanding is fundamentally dependent upon our understanding of the processes by which someone becomes wiser. Over the past 30 years, there has been a sustained and fruitful effort to produce an account of what wisdom is. Researchers such as (Sternberg 1998; Baltes and Staudinger 2000; Ardelt 2004) have produced various accounts of what the fundamental features of wisdom are. However, by and large, the processes by which an individual comes to be wise have remained relatively unexamined. In short, none of the current theories offer guidance in the project of personal self-transformation with the goal of wisdom. This is consistent with the modern western scientific approach, which favours beginning an intellectual endeavour with a taxonomy. However, in the ancient world, wisdom was understood in a medicinal/therapeutic framework. In this way, the alleviation of suffering and the enhancement of flourishing were central to Hellenistic understanding the nature of wisdom (Hadot 1995; 2004). In the ancient world, any attempt to generate an account of wisdom was situated within the context of an account of foolishness, an account of flourishing and, most importantly, how to navigate from one to the other through a process of personal self-transformation. A theory of wisdom guided one in the process of personal development, whereby one came to recognize the causes of foolishness in one's own life and to cultivate the skills and virtues both to alleviate that foolishness and to produce a flourishing life.

This ancient approach embodies Chisholm's insight that process and product are tightly intertwined. We propose that the modern scientific study of wisdom needs to complement its current taxonomic/product approach with a developmental/process approach. As such, we will seek to emulate the ancient tradition: we will provide a psychological theory of foolishness, a theory of the psychological conditions for the possibility of flourishing and our theory of cognitive and affective self-transformation that leads one from a life dominated by foolishness to a life pervaded by flourishing.

What kind of cognitive processes might be central to wisdom? McKee and Barber (1999) after reviewing both *a priori* and empirically based accounts of wisdom point to the central feature of seeing through illusion and by implication seeing into reality. Wisdom seems to involve a special kind of insight, and this is born out by the intuition that it is not odd to say " Sam is not that educated but he is very wise" nor is it odd to say "Sarah is not that artistic but she is very wise;" however, it

does seem quite odd to say "Sam is not very insightful but he is very wise." In a recent neuroscientific review of the literature Meeks and Jeste (2009) explicitly note the importance of insight when they say that one of the central components of wisdom is self-reflection because it "is an essential prerequisite for insight, which is commonly included in many researchers' concept of wisdom." (Meeks & Jeste 2009, p. 360). However, McKee and Barber note that it is not just any kind of insight that constitutes wisdom. Rather it is some kind of insight through illusion and into reality. Meeks and Jeste indicate that the insight involved in wisdom also has importantly to do with self-reflection/self-understanding. The connection here is that illusion is some form of self-deception, so seeing through illusion and into reality involves important insight into one's own cognition and how it might be impeding contact with reality. This type of insight, therefore, involves and enables the self-transformation needed to dispel the illusory processing and facilitate an enhanced interaction with reality. Let's call these *depth* insights because they involve seeing deeper into our cognition, i.e. into the patterns and processes of one's own learning and perception, in order to break through misleading appearances to an underlying (deeper) reality. Such insights also involve increased abilities of self-understanding and self-transformation, i.e. one becomes a *deeper* person.

Three questions immediately arise. The first is what is insight and why is it so cognitively central? The second is what is the nature of the special kind of depth insight found within wisdom and how is it related to the more mundane kind? Finally, is the special type of depth insight sufficient for being wise, and if not how is it integrated into other cognitive processes to help develop wise individuals?

In order to understand the central role of insight and how it could have an impact comprehensively on all of one's life, we need to see insight as a specific and explicit phenomenological experience of a more pervasive and often implicit cognitive process that is central to cognition, viz. *relevance realization* (Vervaeke et. al. 2009). In order to both explicate the nature of relevance realization and to demonstrate its cognitive centrality let us consider problem solving. Historically, we have tried to mechanize problem solving as the following: an initial state, a goal state, a set of operators for moving between states and some set of path-constraints that limit their application. This results in what is conceptualized as a problem-space, which immediately falls prey to the issue of combinatorial explosion: even a relatively small set of parameters can result in a vast number of paths through the problem-space, a number that exceeds human computational capability. This brings into focus what is sometimes referred to as the finitary predicament (Cherniak 1990): our finite computational abilities preclude simple search-space procedures as the means by which we approach problem solving (and cognition as a whole); there is simply too much information, too many possibilities, too many contingencies for us to process. As such, algorithmic strategies of exhaustive search are generally doomed to failure. Rather, our ability to solve problems, to navigate combinatorially dense problem spaces, is contingent upon our ability to constrain that space. This puts the focus on the problem formulation aspect

of problem-solving, rather than the execution of the solution – we need to be able to construct problems in a manner that sufficiently constrains the set of options to one that is computationally manageable. Again, this is the role of relevance realization: the ability to ignore vast numbers of options (hopefully poor ones) and focus on a small set of potentially fruitful ones.

Thus far, in using the classical search-space framework for describing problems, we are still discussing a massively simplified domain, that of well-defined problems. Even with the combinatorial pitfalls thus far encountered, we have not yet fully articulated the difficulties we face in the act of problem solving. Above, we articulated problems in terms of clear initial and goal states, with known sets of operations and constraints allowing us to navigate from one to the other. However, such well-defined problems are rare in our day-to-day existence. Rather, we are most often faced with ill-defined problems: problems in which any or even all of the above parameters are mysterious to us. Consider the problem we (the authors) are currently attempting to solve: writing a good chapter for a scholarly anthology. Our initial state? The blank page. The goal state? A good chapter. But what constitutes a "good chapter"? There is no homogeneous class of "good chapters" that we can look to as a clear goal state; every chapter, good or otherwise is largely unique (certainly, we hope to offer something resembling a unique contribution). Moreover, the set of operations and constraints do not seem to offer much utility in solving our problem: we know a lot of things we should not do (i.e., do not plagiarize, adhere to the word-limits set forth, etc.), but those proscriptions offer little help in guiding us towards positive action in the service of our goal. As mentioned above, problem *formulation* is the key: the ability to resolve a nebulous intent into a specified problem. In the same way that we must use relevance realization to constrain our set of solutions in a well-defined problem, we must employ this machinery to constrain our problem formulation from the ill-defined set of goals and challenges we face every day.

As such, successful problem solving is contingent upon our relevance realization abilities, which is aptly demonstrated when we consider insight problems. The most famous insight problem is the classical nine-dot problem (e.g Weisberg and Alba 1981). The nine-dot problem presents an array of dots: three rows of three dots evenly spaced and aligned. The problem is to connect all of the dots with four connected straight lines. This presents a seemingly unsolvable problem, as most initial attempts at a solution fall prey to a poor problem formulation: the array of dots is taken to signify a square, whose edges form a boundary. And within that boundary, there is no solution to the problem. However, without that boundary, the solution is simple (simple enough to generally elicit negative reactions from those who are shown the solution after giving up on the problem, often with claims that the solution is a "cheat").

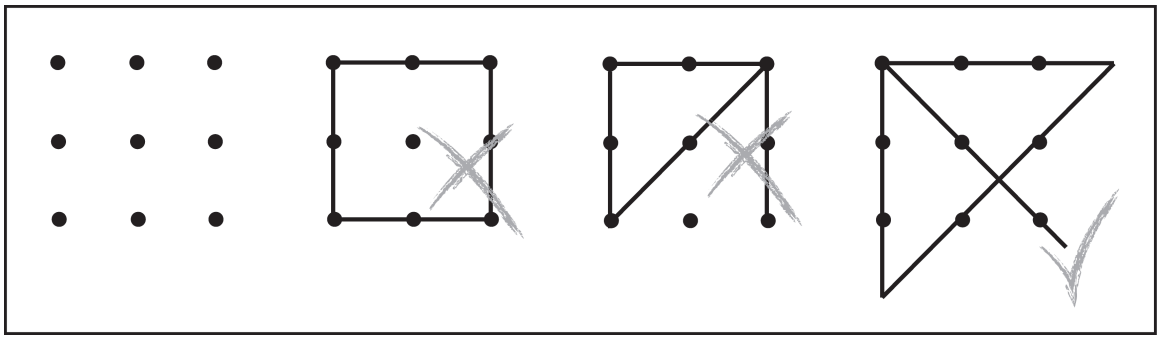

This is the genesis of the often-misused phrase "think outside of the box": the common formulation of the problem precludes its solution. However, by interacting with one's relevance realization machinery, one can reframe the problem in a way that enables the solution. Another good example of this is the mutilated chessboard problem. Consider a standard chessboard: an 8x8 array of alternating black and white squares. It is trivial to see that the entire board can be perfectly covered with a set of 32 dominos (each of which is a rectangle precisely the size of two adjoining squares).

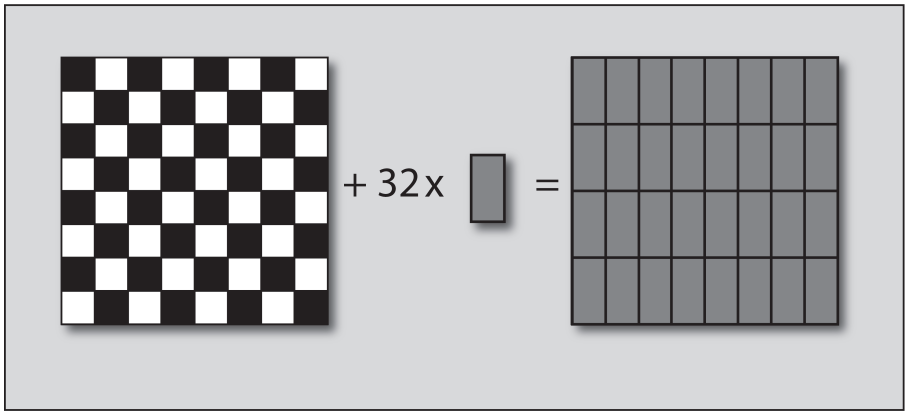

However, if two diagonally opposite corners are removed, can the board be perfectly covered by 31 dominos? Well, mathematically, initially, it seems so: the board has 64 squares and is coverable with 32 dominos. Remove two squares and we are left with 62 squares to be covered by 31 squares. And yet, attempts to produce a pattern of dominos that perfectly covers this mutilated chessboard fail over and over. Over and over, people tasked with this problem try various covering strategies and generally come to believe that it cannot be done, but are unable to provide a proof of this impossibility other than their own repeated failure.

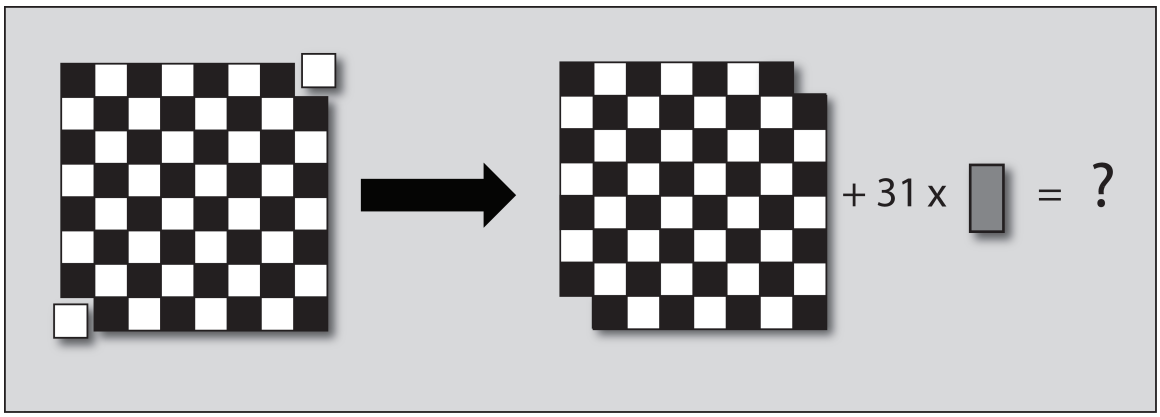

This is due to a poor problem formulation: the board is seen as a grid with simple arithmetic properties, which then result in large number of possible covering configurations. On the other hand, if one considers the colours of the chessboard, the solution becomes extremely simple: a given domino must cover both a white and black square, but removing diagonally opposite corners results in the loss of two squares of the same colour. As such, the board cannot be covered as specified, because an unequal number of black and white squares remain. Note how the framing of the problem is essential in its solution: treating it as a covering problem results in intractability, while focusing on the colours of the squares results in a clear, simple solution. Insight problems are good examples of the central nature of relevance realization and also reveal the value of being able to purposefully interact with that relevance machinery.

If we step back, we can consider problem solving in increasingly abstract scales; simple interaction with the world is the messiest, most ill defined problem of all. Acting in the world involves the consideration of the consequences of those actions in diverse domains from the simple mechanical, to the social, all the way up to the moral and ethical. This consideration of consequences and outcomes is what distinguishes one's actions from mere behaviour – we can talk about the *behaviour* of quasars, but we do not generally consider them to be actors with an intended goal. The consideration of consequences is an extremely broad domain of combinatorial entanglement (Dennett 1987): any action will result in an uncountably large set of consequences. A small subset of these will comprise our intended goals, while the rest fall under the broad heading of *side effects*. The difficulty lies in the appropriate treatment of these consequences – the relevance of our goals is clear, but only some of the side effects are worth considering, while others are immaterial to the matter at hand. As has been noted before, this set of relevant side effects is not a stable, homogeneous class and is extremely contextually sensitive. How then, do we manage to navigate our existence in the face of this sea of information? We cannot evaluate it all, moment to moment, to arrive at a principled determination of what matters and what does not; to attempt to do so would paralyze us into inactivity in even the simplest of circumstances. Rather, we must be able to *ignore* the vast majority of this information. This is the heart of relevance realization: to be able to usefully ignore information in a contextually sensitive manner so as to enable our actions. Taken together, combinatorial

complexity in problem solving, the largely ill-defined nature of most problems and the consideration of consequences serve to indicate the centrality of our relevance realization machinery.

Thus, relevance realization is central; it involves the ability to frame our cognition and, most importantly, to do this *flexibly*, i.e. to be able to reframe. This ability to reframe how we find things relevant is sometimes experienced as the "aha" moment called insight. However, if the above arguments are correct, there is a more implicit form of insight (i.e. relevance realization and reframing) that is pervasive and crucial to our successful learning and interaction with the world. However we should note that there are three kinds of relevance that need to be distinguished. There is also a central feature of relevance realization that needs to be significantly emphasized. Finally, there is also a very important theoretical feature of relevance realization that needs to be briefly discussed because of its important implications for the theory of wisdom.

First let us discuss the three kinds of relevance. We need to note that there is a difference between two pieces of information being relevant to each other, e.g. how the words of this sentence are all *co-relevant* to each other, and how the sentence is relevant to you the reader i.e. how *important* this sentence is to you. Co-relevance is how pieces of information belong together while importance is how they belong to you. Importance signifies information that is relevant to satisfying your goals while co-relevance is about patterns in information that help one to find and make use of important information. These patterns help one to overcome our pervasive problems in managing the information we get from the world. This information is generally partial, polluted with irrelevant information, and is (as we've seen) too vast to exhaustively search. By finding patterns of co-relevance one can facilitate the discovery of important information. Some of the patterns we realize are patterns in events. This patterning gives us the ability to intervene in causal processes, i.e., it affords us *knowing how* to interact with the world. This is our *procedural* knowledge. We can also find patterns of patterns, i.e., higher-order patterns. Especially important are finding patterns that are invariant across many different contexts and are multiply realized in many different causal processes. Such patterns indications causal conditions as general principles. Knowledge of such principles that constrain and enable how events unfold is knowledge of facts. This is our *knowing that* something is the case. It is our *propositional knowledge*[1].

In addition, to co-relevance and importance there is a third aspect to the phenomenology of relevance. As inherently social creatures we need to be relevant to others. We need to say and do things that are relevant to others. Not only does information need to belong together and to us, we also need to emit and transmit information that can belong to others, i.e. we need to participate in

[1] Note that here we are invoking a version of dual processing theory, as reviewed in Stanovich 2002.

and belong to a group. Loneliness and alienation are powerful experiences of suffering that indicate to us that we are somehow lacking in this dimension of relevance. This is the form of relevance that carries us beyond ourselves. Let us call this dimension *transcendence*. The knowledge of these patterns of participation and identification is our *perspectival* knowledge[2]. Perspectival knowledge is not just knowledge of how events are relevant to each other or how facts are relevant to each other. It is knowledge of how facts are relevant to events (constrain and enable sequences of events) and how events are relevant to facts (how events can change constraining and enabling conditions). This enables one to narrtively know what it is like to be in a particular situation, i.e. to participate in it with a particular role/ identity. This is to have a perspective. It enables one to be relevant to the world.

So, to summarize, there are three dimensions to relevance realization: co-relevance, importance, and transcendence. These abilities to realize relevance support the finding of different kinds of patterns. There are patterns in events, which afford procedural knowledge; patterns in facts, which afford propositional knowledge; and, finally, patterns in participation, which afford perspectival knowledge.

Science is largely the project of finding co-relevant patterns of facts that are expressible as propositional knowledge. The kind of insight found in science is theoretical and explanatory insight. How might wisdom differ from scientific knowledge, and what kind of insight is specific to it? Since wisdom is tied to both self-transformation and the cultivation of a good life, it stands to reason that wisdom must centre upon the procedural knowledge that realizes the important information that affords one the ability to intervene in the causal processes of self-transformation and the construction of a good life. However, this is not to say that propositional knowledge will play no role. The wise person must be able to acquire and use the propositional knowledge of causal principles and conditions that factually constrain (and enable) the processes of self-transformation and good life cultivation. So the wise person must have good insight into the important information that facilitates beneficial causal intervention into the self and towards the cultivation of a good life. The wise person must also have good insight into those causal principles and facts that bear upon these processes of self-intervention. Moreover, her or she must crucially have insight into how the procedural and propositional knowledge are perspectivally relevant to each other. Specifically, how the procedural knowledge of self-intervention and the propositional knowledge of factual constraint interact and how to advantageously manage that interaction.

Expertise is largely about finding co-relevant patterns of events that are expressible as sophisticated procedural knowledge. The kind of insight found in expertise is *intuition* (Hogarth

---

[2] We owe this idea to work currently being done with Greg Katsoras.

2006). How might wisdom differ from expertise? One significant difference is that expertise is largely value neutral. One can use one's expertise for good or for evil, yet wisdom seems to be inherently normative. It can only be about making life good. Also, expertise seems to be limited to specific domains with redundant features that support specific practice and provide feedback. (Vicente and Wang 1998; Ericisson and Towne 2010). Wisdom seems to be of a higher order. It does not apply to specific domains, and in this way it is more like intelligence and relevance realization (see below on intelligence and relevance realization) than expertise. There does not seem to be any aspect of life that is not threatened by foolishness and self-deception or that cannot afford opportunities for contributing to a good life. Wisdom seems to be much more like a cognitive style (see below) that can apply broadly. However, this is not to say that procedural knowledge will not play a significant role in wisdom. It is just that this procedural knowledge of the wise person must be integrated with propositional knowledge about important factual constraints and causal principles that govern the transformations of selves, and that this integration always takes place in light of the perspective of a good life for human beings. Wisdom involves the perspectival knowledge of saying and doing things that are highly relevant to creating good lives for human beings. Wisdom therefore involves seeing the world comprehensively in such a way that one can regulate one's actions into alignment with realizing a good life. Wisdom involves the cultivation of character, which is a procedural system more comprehensive than expertise and that is intricately integrated with perspectival knowledge of what it like to be self leading a good life.

All this talk about kinds of knowledge may give the impression that wisdom is largely an intellectual affair. However, this would be to misconstrue the nature of the relevance realization process. Relevance realization is always a matter of the selective direction of attention, the appraisal of value, and the rationing and commitment of processing resources. Relevance realization is simultaneously attentional, affective, and motivational. Relevance realization largely concerns how you *care about and care for* information. Since it involves selection, judgment of value, and the rationing of resources it can be evaluated according to standards of rationality. However, the rationality involved will not be the rationality of theoretical argumentation but a rationality of construal and caring, i.e. it will be a type of rationality that prominently figures depth insights.

The third overall point about relevance realization concerns an important theoretical feature. In a very deep sense we cannot theoretically define the content of relevance. Science requires homogenous and invariant classes of features in order to support broad inductive generalizations. Any set of entities that does not form a homogenous or stable class cannot be the basis of a scientific definition. So, for example, there cannot be a science of things that happen on Tuesday or a science of white things. In a similar manner, any information that we find relevant also does not form some stable or homogenous class. Things we find relevant one minute can be completely irrelevant the next. The classes of things we can find relevant are extremely

heterogeneous. We can find things that happen on Tuesdays relevant, or all white things relevant. This has two important consequences (See Vervaeke et al. 2009 for a more developed version of the following argument). The first is that there cannot be a theory of relevance, there can only be a theory of relevance *realization*. However, this is no cause for despair: consider that the class of features that make any creature fit, in a Darwinian sense, is neither stable nor homogenous. There is no theory of fitness. Natural selection is a theory of how fitness is continually being realized and re-realized in the world. Fitness is inherently pluralistic. What it is differs greatly from time to time and place to place, on many different scales. In contrast, the processes that generate it are universal, so it is not an arbitrary or purely relativistic phenomenon. In the same way what counts as relevant differs from time to time, and place to place, but the processes that realize it are universal and therefore non-arbitrary and non-relativistic. It is important not to confuse the universality of processing and the relativity of content. While such a theory admits of many possible instantiations of wisdom, it nevertheless affords discriminatory evaluation – it is possible to determine what is and is not wise by the application of the process model. As such, while not privileging a specific individual or tradition as defining wisdom, this theory nevertheless allows for the principled distinction of the wise from the unwise necessary for any scientific account of wisdom.

Since relevance cannot be pre-specified nor made constant in terms of its content, this means that, like the process of evolution, it must be a self-organizing and self-defining process (See Lewis 2000 for a discussion of the nature of self-organizing systems within development.) Evolutionary processes are constantly redefining *from within evolution* what it means to be fit. Relevance realization must constantly be redefining *from within relevance realization* what it means for information to be relevant. This in turn means that if relevance realization is central to both cognition and wisdom, then both cognition and wisdom must also be self-organizing and self-defining processes to some significant degree. It also means that wisdom is an inherently pluralistic phenomenon. What counts as wise is going to vary from time to time and place to place. However, the self-organizing and self-defining processes of relevance realization that comprise wisdom are universal, and therefore, there can be a science of how wisdom is realized within cognitive processes. This means that in a deep sense there cannot be a universal and complete product theory of wisdom.

The self-organizing nature of relevance realization helps to explain both its adaptive nature and how it is realized in the brain. Also, this theoretical feature makes important contributions to a theory of foolishness. Perkins (2002) has argued that a specific form of self-organizing processing is responsible both for making the brain adaptive, and importantly, how this can be a source of foolishness. Following the seminal work of Bak, Tang and Wiesenfeld (Bak, Tang & Wiesenfeld, 1987), he argues that the human brain demonstrates *self-organizing criticality*. Consider a pile of sand that has a stream of sand falling onto its apex. The causal constraints of gravity and friction enable the sand pile to organize itself into a cone shape. This cone shape channels the sand in such a

manner to reinforce the cone. So we see a feedback loop in which the shape of the pile directs the sand, which helps to further shape the pile. For a period the feedback loop maintains a stable system. However at some point the sand pile becomes too tall for its base, and there is an avalanche. There is a period of instability and loss of integrity for the system. However, this instability makes possible a wider base for the sand that then starts to constrain the placement of the sand, and the whole system has now re-organized into a newer, more stable one.

Perkins argues that much of cognition runs in this manner. The brain is a dynamical system with multiple feedback loops that create periods of stable organization followed by transition periods of instability in response to inputs to the system. These periods make possible new emergent structures that re-organize and re-stabilize the system. There is now considerable evidence that the brain works this way at many levels and scales of analysis (Beggs and Plenz 2003; Beggs and Plenz 2004; Bassett et. al. 2006; Bullmore and Sporns 2009). Perkins agues that this is how most decisions are made in what he calls "emergent activity switching." The brain is like the pile of sand receiving input from the world. Like the pile receiving more sand, the brain configures itself into a stable feedback loop of activity. But as the input to the brain changes, the brain goes through a transition period of instability that affords a new structure of activity, and this is how the brain constantly redesigns its behaviour in order to adjust to changes in the environment. The brain switches activities in an emergent, self-organizing manner. This helps to explain how the brain is constantly making "decisions" about when and how to switch activities and modes of processing without any direct deliberative decision-making. It also helps to explain how the brain switches activities under circumstances that are too rapid for deliberation (e.g. in the middle of playing sports), yet still in a highly sophisticated and intelligent manner.

Irving, Vervaeke, and Ferraro (2010) have recently argued that the self-organizing criticality found in the brain is a powerful way that the brain implements the self-organizing process of relevance realization, and, moreover, that this is the basic ability that makes us intelligent. This is evident in that relevance realization has a central role at every level of cognition; likewise, self-organizing criticality is found in the brain at all scales and levels of analysis. Likewise, this helps to explain how the self-organizing criticality of emergent activity switching results in sophisticated and intelligent behaviour. With its self-organizing criticality the brain engages in a kind of on-going opponent processing between integration and differentiation of information processing. This means that the brain is constantly complexifying (simultaneously integrating as a system while differentiates its component parts) its processing as a way of continually adapting to a dynamically complex environment. One important property of such self-complexification in any system is that it results in emergent functions and abilities for that system. Complex systems can do more while retaining their integrity as systems. The brain is thus constantly transcending itself in its ability to realize relevant information. It is constantly evolving its cognitive fitness to its environment. So the

connection between relevance realizing self-organizing criticality and intelligence helps to explain how brains can be dynamically intelligent in a dynamically complex world. However, this connection has yet more promise, offering significant insights into a process theory of wisdom via a process theory of foolishness.

Perkins argues that the emergent activity switching that makes one so adaptive is also the source of human folly, which he defines as maladaptive behaviour that is not explained by ignorance (lack of factual knowledge) nor by lack of intelligence. His idea is that the emergent activity switching very much has a life of its own: the feedback loops that constitute it are self-reinforcing and self-maintaining, robust and resilient. They can become very complex and stable structures of behaviour. In a manner analogous to a computer virus, the dynamical processes of the brain can get caught up in self-destructive feedback loops that are very compulsive and highly resistant to change because they have hijacked the self-organizing criticality of the brain. Foolishness is *parasitic processing*. Figure 4 provides a schematic example of such parasitic processing. In it we can see how various construals, biases, and factors of cognitive processing such as encoding specificity can all powerfully reinforce one another such that a complex, compulsive (because of positive feedback), and resilient/resistant self- organizing system can take shape within one's cognition. It is a system that warps our sense of reality and robs life from us. Ironically, the very thing that enables our intelligence and makes us adaptive, i.e. complex self-organization, is also what makes us vulnerable to foolishness.

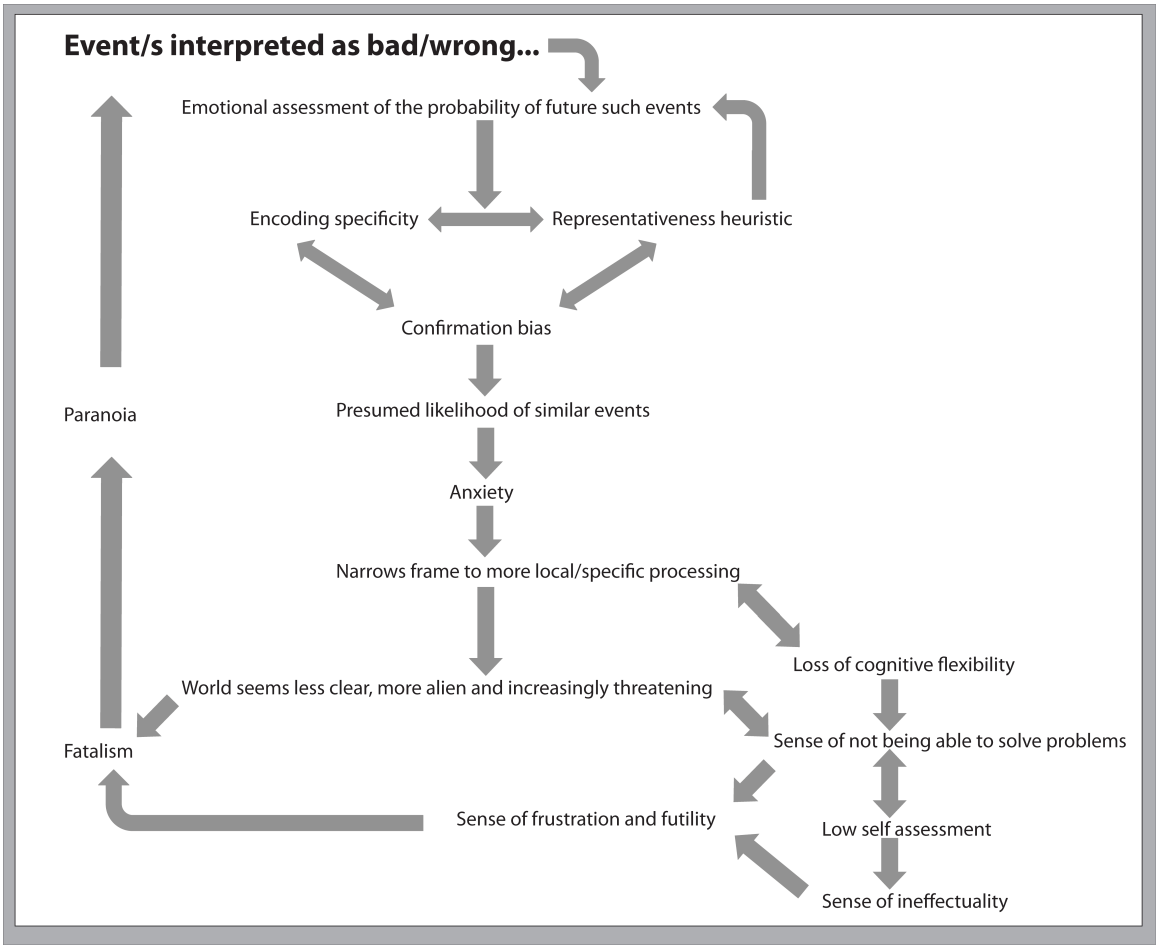

This proposal by Perkins leads us to make an important suggestion: perhaps wisdom involves the purposeful cultivation of self-organizing processes within cognition that enhance depth insights and the rationality of framing. Many schools of thought suggest that wisdom involves just such a self-organizing nature. Examples include: Taoism emphasizes that wisdom involves becoming an empty vessel for the Tao; Buddhism emphasizes that the eightfold path should be understood as an eight spoked wheel in which all the components feed into each other and "roll" along; the Platonic tradition emphasizes the importance of a dialogue between people that has a life of its own; the Neo-platonic tradition emphasizes the advent of the One being something that happens to someone seeking enlightenment; Czikszentmihalyi and Rathmunde (1990) emphasize the importance of "flow" to wisdom, etc. In short, it is commonly believed that the pursuit of wisdom involves the cultivation of the conditions for wisdom, but that wisdom itself takes shape within individuals and very much has a life of its own. In this sense, wisdom is the obverse of the parasitic processing we see at work in foolishness. It is plausible that the only way to deal with the dynamic self-organizing nature of foolishness is with a counteractive self-organizing system within cognition and behaviour.

This common element of the basic natures of both foolishness and wisdom suggests another important theoretical insight. Thus far, the terms 'wisdom' and 'foolishness' have been treated as categorical concepts. However, given their common machinery, a machinery that pervades human cognition in general, it makes sense to consider wisdom and foolishness as opposed points on a single continuum.

As it stands Perkins' theory has much going for it in that it helps dissolve the paradox of foolishness, i.e., how can knowledgeable intelligent people perform such stupid actions? There is no paradox in that this is not a matter of not possessing factual knowledge nor does it involve mysterious processing, for the very same processing that makes us intelligent makes us vulnerable to foolishness. Foolishness is, in this sense, inevitable.

However, there are two important and related lacunae in Perkins account. The first is that Perkins really has not explained the self-deceptive nature of foolishness, only its self-destructive nature. If we return to Figure 4, we can see that an important element of the self-reinforcing nature of depression is the attributional bias that shapes one's perception of one's circumstances. This damaged construal of situations is one of the sorts of illusions that wisdom allows one to penetrate. If wisdom is going to ameliorate the self-destructive nature of foolishness it is going to do this through cognitive processes that see through the illusions of self-deception. Second, while Perkins names top down processes that tune and manage emergent activity switching, their overall place in the cognitive ontology is left unclear. Perkins does not explain how a process other than our intelligence can make use of that intelligence to redesign how that intelligence is being applied. We argue, following Stanovich (Stanovich and West 2000; Stanovich 2002; Stanovich 2006; Stanovich 2009), that these processes are processes of *rationality*, which is derived from but also measurably distinct from intelligence. However, we argue that Stanovich's account of self-deception, while valuable, is inadequate for a theory of wisdom.

Stanovich (Stanovich and West 2000; Stanovich 2002; Stanovich 2006; Stanovich 2009) argues that what he calls dysrationalia is largely due to a failure of reflectively cultivating and applying a cognitive style that protects propositional processing from interference by procedural processing. Very often, intelligent and knowledgeable people allow themselves to use procedural processing when propositional processing is needed. Stanovich (2002; 2009) has good evidence that this failure is not due to a lack of intelligence and that intelligence is largely not predictive of the ability to behave rationally [cite things, plus book). Instead, what is missing is a *cognitive style* that can be cultivated and applied through reflective awareness (i.e., a metacognitive, phenomenological awareness, not a theoretical one) of both our propositional and procedural processing so that their relationship is properly managed. A cognitive style is a configuration of domain general abilities of directing attention, valuation, and motivation that produces a particular salience landscape within

which one undertakes one's tasks. We need a cognitive style that directs attention, valuation, and motivation such that propositional information should be processed by formal inferential procedures, i.e. what we need is a cognitive style that protects and promotes the *rationality of computation*. This is a reflective monitoring and managing of cognition that insures that propositional information is being properly encoded and manipulated by formal inferential procedures such as logic and probability theory.

There are key ideas within Stanovich's theory with which we strongly agree. The first is that wisdom (the rationality for overcoming foolishness) is a matter of reflection, i.e. it crucially involves directing attention to the medium of cognitive processing rather than the content of the cognitive product, such as a new belief or decision to act. Rationality, as a cognitive style, is not directed primarily at the content of one's beliefs, but rather the process by which beliefs are generated and validated. Being rational means, in part, that one's cognitive medium of computational processing is salient. The second insight with which we agree is that self-deception involves conflict and misalignment between different cognitive processes. Finally, the third is that rationality is not a matter of simple intelligence but rather how intelligence is reflectively applied to its own operations through the cultivation of cognitive styles. Such cognitive styles (Kozhevnikov 2007; Sternberg 200?) involve learned skills and sensitivities of *learning to learn* (Mercado 2008; Harlow 1949).

Intelligence can be directed not only to learning about and interacting with the world to produce a particular cognitive product e.g. some knowledge or behaviour, but it can also be directed to learning about the cognitive processes of learning itself, i.e. it can become learning to learn. A habitual way of learning to learn results in a cognitive style for an organism. Such learning to learn means that an organism notices important patterns in how it is processing information and intervenes in those patterns, i.e., restructures them, in order to improve how the information is being processed. Rationality presupposes such learning to learn, for without the ability to intelligently intervene in its own cognition a creature cannot be held responsible for its cognition nor deemed subject to rational standards. Such intervention also means that the creature is *using intelligence to improve how intelligence is being used and developed*. This leads to an interesting suggestion: if rationality is the reflective application of intelligence to the use and development of intelligence, then perhaps, in a similar manner, wisdom is the reflective application of rationality to the use and development of rationality. Perhaps wisdom is *using rationality to improve how rationality is being used and developed.* However, we should note that such learning to learn is often initially associated with insight since insight involves restructuring what patterns ones finds relevant. It involves altering how co-relevance is being realized so that to better facilitate the finding and using of information that is important to the organism. Learning to learn would also be needed by the depth

insights associated with wisdom, since these involve self-reflective insights, i.e., insights into one's own cognitive processing.

We argue that in addition to the rationality of computation advocated by Stanovich that we also need a rationality of framing. As previously argued, this framing is always simultaneously a matter of attention, motivation and valuation. We will call this a rationality of construal, however, this should not be conflated with our previously articulated notion of perspectival rationality, what we will later articulate in more depth as our rationality of communion. We argue that the self-deception of foolishness is largely about being locked into a feedback loop that maintains and reinforces an inappropriate construal of problems, i.e. a mis-framing of problems and situations, that cause one to misjudge the co-relevance and importance of the information presented [fig 1]. This is why our foolish behaviour persists – it leverages the same adaptive self-organization that underlies all of our intelligent cognition. It possesses the same persistence, resilience and growth that is the hallmark of human adaptivity. In short, self-deception is a self-imposed lack of the cognitive flexibility needed for depth insights, so that we trap ourselves in illusion. This entrapment follows from the self-organizing nature of relevance realization; this means that a rationality of construal will have to be one that can counteract the formation and persistence of such parasitic processing. We will also argue, following the literature on insight problem solving (Weisberg and Alba 1981; Adams et. al. 1988; Lockhart et al. 1988; Needham and Begg; 1991; Schooler et al. 1993) that factual knowledge is largely impotent to overcome such mis-framing and parasitic processing. Instead, highly related skills of construal, self-awareness, and self-regulation are needed to overcome the mis-framing at the heart of foolishness. These skills need to be deployed in a self-organizing system in order to counteract parasitic processing. Computation alone is largely inadequate to the task of realizing the depth insights needed to reframe the problems of life and escape foolishness. Moreover, computation is also largely inadequate on its own for manifesting the self-regulation needed to change patterns of behaviour. Instead, success or failure in self-regulation is largely dependent upon the skills and sensitivities of construal. (Ayduk and Mischel 2002; Myrseth and Fishbach 2009).

So, in addition to a reflective cognitive style that protects and promotes a rationality of computation, the wise person needs to develop a reflective cognitive style that protects and promotes the rationality of construal, i.e., how we are framing events so that we may interact with the world free from parasitic processing. The wise person would also need to have skills and sensitivities for how to co-ordinate these two cognitive styles, both within themselves and also in the context of a larger community. They would need to know how to govern the relationships between the rationality of computation and the rationality of construal and balance these within a rationality of communion (both inter- and intrapersonal communion). The manner in which these cognitive styles are coordinated within an individual is an internalization of how we perspectivally manage our belonging to other people and participating in a culture. One learns how to get the two cognitive

styles to belong together by internalizing knowing what it is like to belong to a social group, i.e. knowing what it is like for different takes on the world to fit together. In a similar manner, one learns how to commune with one's envisioned future self leading a good life, and thereby belong to one's future self, by means of internalizing how one currently communes with others. In this way the rationality of communion involves perspectivally appropriating one's development in order to transform oneself into a better future self. This involves internalizing the perspectives of appropriate role models. One powerful way to enhance learning to learn it to look at one's processing from the perspective of someone who is performing better. From these role models the person cultivating wisdom would through such learning to learn come to emmulate a higher order cognitive style of perspectival rationality for the application and co-ordination of the lower order computational and construal cognitive styles of rationality. This higher order process would realize a cognitive meta-style for rationally improving how rationality is being applied and developed. Wisdom is such a rationally self-transcending rationality.

Stanovich has considerable work supporting the conclusion that Baron's *active open mindedness* (AOM) (Baron 1994) is an optimal cognitive style for a rationality of computation (Stanovich and West 2000; Stanovich 2002; Stanovich 2006; Stanovich 2009). AOM means actively looking for how bias is warping and thwarting our problem solving. A bias is when a heuristic of procedural processing is being used on a problem for which computational processing is better for achieving the goals of the problem. A bias is a misplaced heuristic that interferes with computational processing. AOM requires learning about cognitive biases and actively searching for them in one's processing in order to actively counteract their effect by applying more formal computational procedures (Baron 1994). It should be noted that AOM is applied to deduction, induction, and abduction (plausibility reasoning) as well as to their successful integration within reasoning. AOM enhances our sense of how information is relevant to and within our propositional grasp of facts. It enhances our sense of *reasonableness*. If affords that good vision of factual reality that Aristotle called *sophia* (Aristotle, Nichomachaen Ethics).

However, a rationality of construal is also needed. Langer (1989; 1997; 2000) **[footnote**] and current independent work within clinical psychology (Teasdale 1999; Baer 2003; Hayes et. al. 2004) and neuroscience (Farb et. al 2007) point to the cognitive style of mindfulness as being very important for comprehensively transforming and improving the framing of situations so as to facilitate not being trapped in self-defeating construals of situations and problems. It sees through parasitic processing. Mindfulness involves paying attention not to one's propositional encoding or inferences. Instead it involves paying attention to how one is paying attention to situations, i.e. to how one is distinguishing relevance from irrelevance. This involves training skills of the shifting of the direction, aspect and scale of attention and thereby enhancing the cognitive flexibility of construal. Mindfulness enhances our situational insight, foresight (insight into potentials within

situations) and mindsight (insight into others' cognition (Siegel 2007; 2010), and the interaction between them. These skills are trained in mindfulness cultivation practices such as meditation and contemplation. Mindfulness is an optimal cognitive style for the rationality of construal. Mindfulness enhances our sense or how information is relevant to and within our procedural grasp of events. It enhances our sense of *fluency*. It affords that efficacy of interaction with situations that Aristotle called *phronesis* (Aristotle, Nichomachean Ethics).

Wisdom relies upon a meta-style of rationally self-transcending rationality that governs the rational cognitive styles of mindfulness and AOM in a manner that enhances the developmental complexification of information processing that is constitutive of basic general intelligence. Wisdom involves simultaneously differentially developing each cognitive style of rationality and integrating them together in a self-organizing system which thus enhances relevance realization overall. We become more fluently reasonable and more reasonable in our fluency. However, while this framework may provide a theory of how foolishness is overcome, and this may be the central therapeutic aspect of wisdom, it is questionable that this will provide for a sufficient account of wisdom. It is very plausible that there is much more to leading a good life than merely avoiding foolishness. As we've argued, a person cultivating wisdom needs to appropriate their own development by learning how to commune with their future flourishing self. A theory of wisdom needs a theory of the cognitive processes that afford such communing and flourishing. There has been a lot of important work by psychologists (Diener et. al. 1999; Diener 2000; Ryan and Deci 2000; Deci and Ryan 2000; Ryan and Deci 2003) and philosophers (Kekes 1986;1995: 2000; 2002:2006; Frankfurt 2004;2006; Russon 2009; Wolf 2010) about the cognitive processes that create the conditions for the possibility of a good life. All these theorists realize that wisdom can at most create conditions of possibility for flourishing. The world must co-operate for an individual to be flourishing.

By drawing on all of this work it seems plausible that there are three central dimensions that must be in place in order for an individual to be leading a good life. The individual must have subjective well-being (Diener et. al. 1999; Diener 2000). The individual must judge and experience their life as satisfying. Second they must judge and experience that their lives are morally respectable, both by themselves and by others whose moral judgment they respect (Kekes 1986;1995: 2000; 2002:2006). Finally, and more recently, Susan Wolf (2009; 2010) has argued for the independent dimension of *meaning in life*. Meaning in life is to judge and experience oneself as connected appropriately to something that has an important value independent from one's valuing of it. More loosely, this can be put as that one judges and experiences that one is connected to something "bigger" than oneself or as Wolf puts its "subjective attraction meets objective attractiveness" (Wolf 2010). To use our terminology meaning in life is to judge that one has

significant transcendence. So overall the conditions that makes one's' life good are that one judges and experiences that ones life is satisfying, virtuous (morally good), and deeply connected.

A theory of the cognitive conditions that afford achieving such ends would be a theory of the flourishing that is necessary to wisdom. What might those cognitive conditions be? One helpful suggestion, that draws these three dimensions together, is that a wise person must have a constellation of abilities focused on the making and protection of selves (identities, persons) and communities of selves. It is plausible that conditions of satisfaction (subjective well being) mark out conditions for the agency necessary to being a self. This is a subjective sense of well-being based on the sense that one has the autonomy and competence to have succeeded and to continue to succeed reliably (not perfectly) in one's projects (Ryan and Deci 2000; Deci and Ryan 2000; Ryan and Deci 2003). It is also plausible that meaning in life marks out those conditions that create the kinds of connections that are constitutive of selves such as connections between mind and body, mind and mind (both intra and inter personal connections) and between mind and world (conditions for action and interaction). It is also plausible that virtue marks out the conditions that protect, as inherently valuable, persons and their communities. The wise person would have depth insights into the processes by which selves/identities/persons are constituted, protected, and promoted.

These depth insights mean that the wise person is very focused on the transcendence dimension of relevance realization. The wise person is not only aware of how pieces of information are relevant to each other (co-relevance) or how they are relevant to the wise person himself (importance), they are also aware of and concerned for how he/she is relevant to others. As we have argued, the wise person is not only concerned with realizing what is relevant they are also concerned with saying and doing relevant things. Agency, and meaning in life, virtue come together in actions that connect one to others in ways that value those relations for their own sake

This concern for transcendence means that the self-knowledge involved is not primarily autobiographical knowledge nor is it narcissism. It is the perspectival knowledge of what it is like to be a self and the linked procedural and propositional abilities to cultivate selves. This knowledge results in *knowing what is good for selves and how to bring it about*. It means knowing how to make and share those identity-making connections between mind and body, mind and mind (both intrapersonal and interpersonal connections) and between mind and world that support a sense of successful agency.

The self-knowledge involved also means knowing how to care for the dynamic and reciprocal connections of identification (and therefore identity creation) between agency and meaning in life, i.e. skillfully understanding that we want meaning in life that fosters agency and agency that fosters meaning in life, and that we have a process of dynamic reflective equilibrium

between them. We sometimes alter which aspects of our actions we identify with (agency transformation) under the normative guidance of our sense of meaning in life, i.e., we evaluate what kind of agent we are in terms of how it is effecting our meaning in life. For example, we sometimes may wonder if our actions are messing up our ability to find and feel at home/at peace in the world. We sometimes realize that our attempts to find satisfaction are undermining our sense of connection. We also sometimes alter which aspects of our meaning in life that we identify with (meaning in life transformation) under the normative guidance of our sense of agency, i.e., we evaluate what kind of home in the universe we are making with ourselves and with others in terms of the kinds of agency we are fostering, For example, we may wonder if our meaning in life is causing the existential confusion or foolishness that undermine agency. We sometimes realize that our attempts to connect are undermining our ability to find satisfaction.

Finally, the self-knowledge involved is eudaemonisitc in nature. It involves knowing how to create and implement those rules of attitude and conduct that protect selves, their projects, and their meaning in life. Such self-knowledge involves knowing how to be virtuous and how to promote virtue and how to protect selves, their projects, and meaning in life from vice. This means that self-knowledge involves the perspectival awareness of how to align actions to a self-transcending vision of a good life to which the person cultivating wisdom is committed and to which they belong.

Wise people have the knowledge of the nature of selves such that they can promote and protect: agency, meaning in life, the agency-meaning in life identity formation relation, and the virtues needed to protect selves and communities of selves and to afford participation in this way of life.. Although these four areas of knowledge are analytically distinguishable, the four greatly interpenetrate and interact in reality. Knowing how to cultivate this dynamic system of knowledge constitutes knowing how to flourish, and how to promote and protect flourishing.

How then does one personally enact the cultivation of wisdom in one's own life? We've argued that wisdom requires the cultivation of the cognitive styles of AOM and mindfulness as well as the cognitive meta-style that affords their mutual complexification. We've already discussed practices such as meditation and contemplation for the cultivation of mindfulness, and practices such as bias identification and counteraction for the cultivation of AOM. These practices have already been examined in considerable detail elsewhere. However, the enactment and training of our proposed cognitive meta-style requires explication. Implicated in this cognitive meta-style is the management of perspectival knowledge under the normativity of the rationality of communion. This can be accomplished by a process we are going call "internalizing the sage". By taking the perspective of the sage on one's cognition one can enhance learning to learn until one comes to emulate the perspective of the sage. This a common strategy used by many wisdom traditions in different times and places. For example, St. Paul emphasizes that it "not I who live but Christ who lives in me" (Galatians 2:20).

Within Buddhism one is to realize one's own Buddha nature. Stoicism has been described as the process of becoming like Socrates (Long 2004). The followers of Epicurus explicitly practiced imagining how Epicurus would reflect upon their actions. A plausible explanation of Plato's use of Socratic dialogues is to create a "spiritual exercise" (Hadot 1995; 2004) for internalizing Socrates and his method of enlenkos. The Tao Te Ching continually talks about the Sage and how the sage sees the world and acts within it. More recently, Baltes and Staudinger have produced empirical evidence that simply imaging talking to another person improves ones performance on tasks related to wisdom (Baltes and Staudinger 2000). This work indicates that simply viewing one's cognition form the perspective of another does facilitate becoming more insightful. It is reasonable that Vygotsky's (1978) concept of internalization within proximal development helps to explain why internalizing someone with more wisdom could even further enhance one's insight. Why would taking the perspective of the sage be good for overcoming foolishness and affording flourishing? One suggestion is that by taking the perspective of the sage one comes to have a salience landscape that is similar to that of a wise person. One tends to find salient those things that a sage would. This is a method for procedurally implementing the factual knowledge we have concerning the good life and revealing that gap between who we are and who we want to be. All of the above wisdom traditions placed a high priority on turning theory into practice. Recently, Kosslyn and Moulton (2009) have described a similar process in athletic training in which mental practice is used to enhance performance. The athlete first stores motion images from the coach. The athlete then tunes the images to their own body, but from the third person perspective. The athlete then compares the imagined action to real actions and shifts from the third person perspective to first person perspective. Finally the athlete trains to make this response habitual. So it is plausible that internalizing the sage involves both taking the perspective of the sage on one's cognition, a kind of heightened third person perspective on one's cognitive processing that enhances learning to learn, and then taking the perspective of the sage on the world as a kind of heightened first person perspective that transforms one's salience landscape so as to implement the insights gained through the learning to learn. Finally, this orientation is trained until it becomes a habit of mind and way of life.

Internalizing the sage also engages processes of transforming one's salience landscape, and such salience transformations are key to self-regulation (Metcalfe and Mischel 1999; Ayduk and Mischel 2002; Myrseth and Fishbach 2009) required to appropriate the development of one's character. Finally, internalizing the sage develops the narrative skills for excellence in perspectival knowing. Following the ancient Greek wisdom tradition, having the salience landscape of the sage such that one naturally self-regulates so as to realize the good life was called *sophrosyne* (Schmid 1998; McGhee 2000) As such, we propose calling the cognitive meta-style upon which wisdom relies sophrosyne. In sum, the pursuit of wisdom requires cultivating sophrosyne by internalizing the sage,

and using sophrosyne to govern the complexification of AOM and mindfulness. One can cultivate AOM by becoming aware of inferential bias and actively counteracting it. One can cultivate mindfulness through meditative and contemplative practices and break up parasitic processing. In this way one cultivates a self-organizing system that enhances and develops the relevance realization central to cognition (See Figure 5). One improves one's ability to see through foolishness and into the good life.

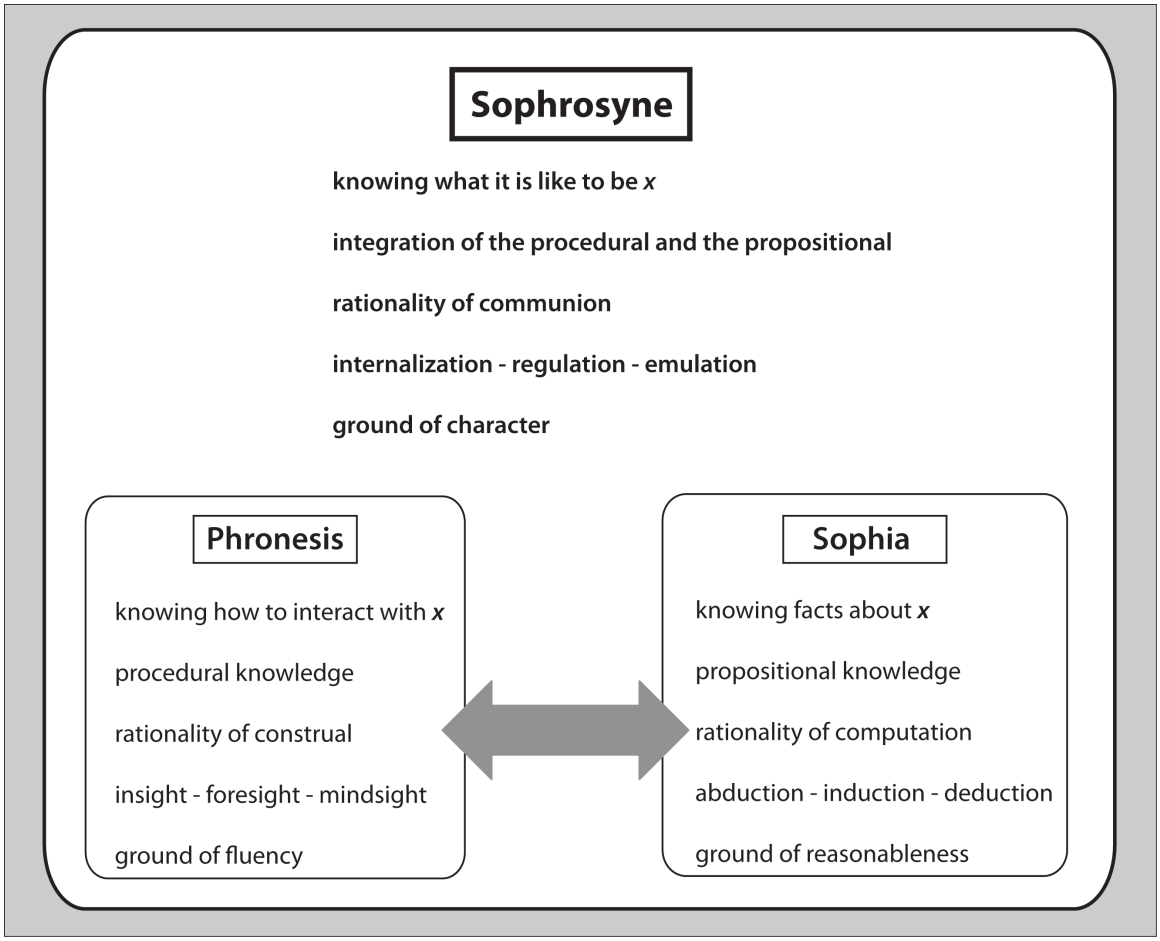

Acknowledgement: We would like to thank Greg Katsoras for his help with the idea of perspectival knowing. We would like to thank Najam Tirimizi for his help with the idea of rationally self-transcending rationality. Finally, we would like to thank David Kim for his help in discussions of the role of wisdom in resolving internal conflict.