

# Structural Complexity and Effective Informing

*T. Grandon Gill*

*University of South Florida, Tampa, Florida, USA*

[ggill@coba.usf.edu](mailto:ggill@coba.usf.edu)

## Abstract

A task that is unfamiliar and non-routine is sometimes characterized as having high structural complexity. Through increasing the task performer's knowledge, informing processes can have a significant impact on structural complexity—as well as being impacted by it. The present paper examines the structural complexity construct, which can be further broken down into state complexity, operator complexity, and goal complexity, with the objective of better understanding how the qualitative nature of structural complexity impacts the informing process. The analysis finds that the gap between conceptual expertise and practical expertise is particularly difficult to overcome. Informing techniques favored by the conceptual expert, which emphasize top-down informing using symbolic concepts, tend to be very much at odds with those favored by the practicing expert, who likely favors a bottom-up approach based upon examples. Techniques for overcoming these barriers are discussed, such as the use of vertical slicing and stories.

**Keywords:** Structural complexity, task complexity, problem space, goal space, informing systems, learning, knowledge compilation, problem solving.

## Introduction

The idealized informing system consists of a sender supplying information to a client through a channel in support of a task (Cohen, 1999). For such informing to be effective, the information being supplied must impact the mental model of the client. We refer to this model as a problem space (Gill & Hicks, 2006). Given this general framework, it makes sense to speculate that a qualitative understanding of the sender's and client's respective problem spaces could lead to insights regarding both the nature of the information that will be most effective in informing the client and its presentation. Such insights, in turn, could then be translated into rules that could be applied towards achieving more effective informing—an objective that is at the very core of the informing sciences.

It has been proposed that task complexity may be an important contributor to the informing process (Gill & Hicks, 2006). Of particular interest is structural complexity, which defines task complexity in terms of the nature of the task specific knowledge in the problem space. As such, it

---

Material published as part of this publication, either on-line or in print, is copyrighted by the Informing Science Institute. Permission to make digital or paper copy of part or all of these works for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage AND that copies 1) bear this notice in full and 2) give the full citation on the first page. It is permissible to abstract these works so long as credit is given. To copy in all other cases or to republish or to post on a server or to redistribute to lists requires specific permission and payment of a fee. Contact [Publisher@InformingScience.org](mailto:Publisher@InformingScience.org) to request redistribution permission.

changes as the performer learns more about the task. This learning aspect of structural complexity makes it a good candidate for giving us insights into the process of informing. The present paper adapts a structural complexity model, first proposed for use in matching task characteristics to expert system tools (Gill, 1991), to the general informing problem. It then analyzes how different levels of problem space structural com-

plexity virtually demand different qualitative types of informing.

We begin by presenting some definitions and then introducing a model of structural complexity, reviewing relevant literature as the model is developed. We next consider what types of informing are most likely to succeed at each level of structural complexity. Finally, recommendations deriving from the hypothesized relationship between informing and structural complexity are presented.

### Definitions

Prior to developing the structural complexity model to be used in this paper, it is useful to present a number of definitions. We begin with how we will be using the term task:

*Definition:* a *task* is a set of assigned a) goals to be achieved, b) instructions to be performed, or c) a mix of the two. (Hackman, 1969)

We next define the term problem space as follows:

*Definition:* A *problem space* is a representation of the cognitive system that will be used to perform a task “described in terms of (1) a set of *states of knowledge*, (2) *operators* for changing one state into another, (3) constraints on applying operators, and (4) *control knowledge* for deciding what knowledge to apply next.” (Card, Moran, & Newell, 1983, p. 87)

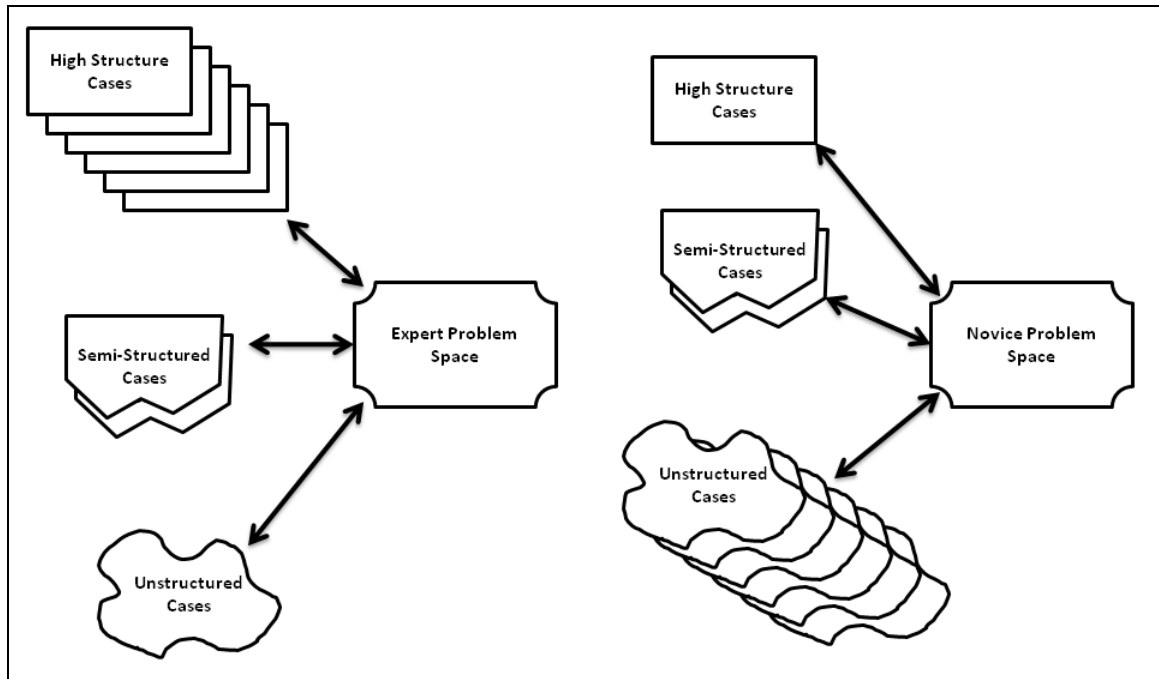
To organize our problem space, we will describe the set of knowledge states as the *state space* and the set of operators, constraints, and control knowledge as the *operator space*. Because tasks without goals—item (b) in our task definition—tend to be highly structured (almost as a matter of definition), we will also assume the existence of a *goal space*, containing one or more operators that allow the task performer to assess the fitness of the state space with respect to task goals and to determine relative progress towards these goals.

*Structural task complexity*, or *structural complexity*, refers to the degree to which a task is performed using task specific (as opposed to general purpose) knowledge, operators, and goals. Under a structural complexity definition, low structure is more complex than high structure, i.e., unfamiliar tasks are more structurally complex than routine tasks. To avoid confusion with the many conflicting definitions of task complexity (see the complete review in Gill & Hicks, 2006), it should be noted that structural complexity differs from other widely used task complexity definitions in a number of important respects, including:

- Unlike objective task complexity (Wood, 1986)—which is defined to be a function of the number of task components, their degree of inter-relationships, and the degree to which these relationships are dynamic—structural complexity is not a property of the task itself but of the problem space used to perform the task. Thus, structural complexity nearly always declines as task performer expertise grows (e.g., through practice or through being informed as to how better to perform the task).
- Unlike experienced or information processing complexity (Gill & Hicks, 2006), structural complexity is not necessarily a good predictor of task performance metrics such as bytes of information processed, arousal, or perceived difficulty.
- Unlike other problem space complexity metrics, such as number of paths through the problem space (e.g., Campbell, 1988; McCabe, 1976), structural complexity is less concerned with the potential size and precise construction of the problem space than with the qualitative nature of the cognitive structures likely to be used during task performance.

The final definition to be presented is that of a *task case*. Whereas the problem space describes the cognitive elements used to perform a task, a task case is an individual instance of a task to be performed. For example, the problem space of a physician would include a vast collection of medical knowledge; a task case might involve an individual patient coming in to visit the physician with a specific complaint to be diagnosed.

It is important to recognize that structural complexity can vary across task cases. As suggested by Figure 1, for an expert's problem space, most commonly performed task cases are likely to be highly structured. We'd also expect some semi-structured cases—where some task-specific knowledge is available but not enough to complete the task (e.g., the underlying cause of a patient's symptoms were the result of a rare local disease contracted during previous travels to a remote region)—and even a small number of unstructured cases—where existing knowledge in the problem space does not apply to the task. The typical problem space of a novice, in contrast, would likely include a vastly larger percentage of task cases that would require unstructured problem solving.



**Figure 1: Expert vs. Novice Problem Spaces**

Figure 1 highlights the point that structural complexity is really an interaction between the performer's problem space and the particular task case being performed. When we consider the nature of informing, we will generally be looking at the problem solving activities invoked for a particular task-case. Sometimes, however, it will also be useful to speak generally about complexity across a collection or representative sample of task cases. When doing so, the term *overall structural complexity* (or *overall structure*, when referring to the opposite of structural complexity) is used. It can be loosely defined as an average or weighted average (by frequency of occurrence) of structure levels across the set of task cases being considered. In the Figure 1 example, we could therefore say that the overall structural complexity of the novice problem space is substantially *higher* than that of the expert because its overall structure is *lower*.

## Structural Complexity Model

The qualitative framework used to describe structural complexity, modified from an earlier model developed for mapping tasks into expert system tools (Gill, 1991), characterizes problem space structure on three dimensions: state, operator, and goal. In each dimension, there are two scales: one reflecting the qualitative nature of the dimension and one reflecting the degree of uncertainty or ambiguity present. Each of the three dimensions is now described in detail using a number of example tasks—e.g., playing chess, solving a physics problem, writing an essay—to illustrate the different levels of structure.

### **State Space Structure**

The state space component of the problem space consists of the knowledge schemata that can be applied to a particular task case or, more broadly, to a set of task cases. We need to be a bit vague regarding the precise nature of schemata for two reasons. First, it makes sense to avoid unnecessarily limiting the generalizability of the proposed model. Second, and more to the point, although a large and sometimes impressive body of research has accumulated addressing the nature of individual schema—in fields that include cognitive psychology, cognitive science, computer science, neuroscience, and even philosophy—a great deal more research will be necessary before we will be ready to assert their structure in detail. Certainly included would be symbolic representations of task information, mental images (Kosslyn, 1980), cognitive representations of other sensory and emotional inputs (e.g., Gilbert, 2007), and semantic linkages between these elements, both of a hierarchical (e.g., inheritance) and relational (e.g., semantic networks; Woods, 1975) nature.

Approximate levels of structural complexity, as they relate to the state space, are presented in ascending order in Table 1. A particularly interesting aspect of this framework—later repeated for the state and goal problem spaces—is that our intuitive notion of an expert (e.g., a master of theory, such as a professor) is almost precisely the opposite of the way typical experts most commonly reason. The most widely cited evidence for this assertion comes from the study of chess, which has been referred to as the *drosophila* (fruit fly) of the cognitive sciences owing to its frequent use in the study of problem solving (Charness, 1991, p. 39). An obvious candidate for the most important skill required of a chess master would be the ability to look more moves ahead than a novice. Studies suggest, however, that once a player reaches an intermediate level of play a relatively stable level of move look-ahead is reached. What does change, on the other hand, is the number of board positions (state spaces representations) than can be recalled by the player. Among the evidence supporting this view is the astonishing ability of chess masters to recall board positions after a short inspection (an ability that declines rapidly as pieces are placed in random positions) and the ability to recall the sequence of moves in past games. As expertise continues to grow, grand masters develop enhanced ability to overlook moves that are unsuitable without conscious thought (Charness, 1991). A similar novice-expert distinction has been observed in many domains (e.g., Larkin, 1981).

The human cognitive system is actually built so as to encourage such movement towards reasoning through unconscious processes. The process of chunking (Miller, 1967), for example, allows scarce working memory—where most conscious processing typically takes place—to be conserved by taking schemata and composing them into ever-larger chunks. Automization (Shiffrin & Dumais, 1981) has a similar impact on sequences of operators applied during task performance. Thus, higher levels of state space knowledge (e.g., abstract models) are only required for those task cases where existing task related schemata are inadequate. It is the low structure cases—not the high structure cases—that require theoretical reasoning. If an expert's job consists entirely of routine task cases, such theory may never be required. Collectively, the processes by which expert knowledge is transformed—sometimes referred to as knowledge compilation—

contribute to a wide range of observed phenomena, such as the learning curve (Neves & Anderson, 1981).

**Table 1: State Space Complexity Levels, with lower levels being associated with more practiced tasks and greater expertise in task performance.**

<i>State Complexity Level</i>	<i>Name</i>	<i>Description</i>
1	Full Schema	A fully structured state space, with all the schema and relationships required to perform a particular task. A physician, for example, would enter into a diagnostic task with a large set of schema that might be knowledge for representing different conditions, tests that could be applied, mental image representations for symptoms observed in the past, knowledge of common and uncommon patterns of symptoms and measures, and so forth.
2	Schema Templates	A collection of knowledge structures defining information needed for task performance and interrelationships within such information. As an example, a blank set of tax forms and instructions identifies the information that is required for preparing a tax return. The layout of a chessboard and knowledge regarding how pieces move prepares the way for playing chess.
3	Abstract Models	Knowledge that helps the task performer identify what knowledge is and is not relevant to the task at hand, the general nature of the components of the task, and the general nature of relationships expected to be present in performing the task. A physician, for example, might have mental models of the functioning of different body systems, the nature and impact of various bacteria and viruses, the conceptual basis for various treatment options, and so forth.
4	Unstructured	Little or no task specific knowledge is present. In addition, it is likely to be difficult or impossible to determine what information is and is not relevant to the task. While such a situation will certainly make task performance more complex, it will not necessarily preclude task performance since general purpose techniques (most notably, analogy; Rumelhart & Norman, 1981) are available for mapping knowledge and relationships from one domain into another in cases where task-specific structure is unavailable.

Along with substantial body of literature regarding the knowledge compilation process, a number of compelling examples of the phenomenon have accumulated in the literature and popular press. One that illustrates the compiled nature of specialized expertise involved the identification of a forged artwork at the Getty Museum, where a number of experts expressed the sentiment that a particular statue didn't "feel" right despite its relatively solid provenance and the fact that they could not articulate what bothered them (Gladwell, 2005). In an example of everyday expertise, an experiment involving choice in the context of shopping investigators found that buyers confronted with a complex furniture choice were happier a week later if they decided based upon their first impression as opposed to being asked to think about it (Dijksterhuis, Bos, Nordgren, & van Baaren, 2006). In a similar vein, one of the most common multiple choice test taking strategies is to go with your first guess unless strong disconfirming evidence is later encountered (Es-trada, 2008).

For each of the three spaces, a second set of levels reflecting uncertainty or ambiguity is also specified as a part of overall structure. For the state space, these specifically relate to uncertainty and are presented in Table 2. These differ from the state space complexity levels in three important ways. First, they are properties of individual values within the state space, rather than of the entire state space being applied to a particular task case. Thus, a given task case may involve schemata at a variety of uncertainty levels. Second, uncertainty levels may change over the course of task

performance as a result of task performer action. Third, a given uncertainty level may be an inherent property of the task—the oven being used in a cooking task has an inaccurate thermometer—rather than being a function of the particular problem space. When such is the case, expertise is unlikely to change that particular uncertainty level. Otherwise, enhancements that come with expertise (e.g., more acute observational power) will tend to drive values towards lower uncertainty as expertise increases.

**Table 2: State Space Uncertainty Levels, often a function of the task itself**

<i>State Uncertainty Level</i>	<i>Name</i>	<i>Description</i>
1	Knowable	Values in the state space are directly observed, established by the task, fully determined by other values, or set by the task performer.
2	Measured	Values may not be observed directly, but may be measured with some quantifiable range of uncertainty.
3	Uncertain	Values that have yet to be specified, measured, or observed, values that may change between task performer observations, and values that are tacit in nature (as is often the case for moderating variables in the social sciences) all fall into this category. Tacit values, in particular, generally need to be deduced or confirmed indirectly from other values in the state space.
4	Unknown	Values for which little or no information is available, either with respect to the actual value or the range of possible values.

## **Operator Space Structure**

The operator space component of the problem space consists of operators (i.e., actions) that the task performer can apply to change the state space, constraints on when those operators can be applied, and control knowledge that determines what operators should be selected at each point during task performance. Such changes may involve direct changes to the external task environment (e.g., moving a chess piece) or may take place entirely within the task performer's internal representation (e.g., mentally generating the set of the opponent's possible responses to a possible move being considered by the task performer). Without any particular loss of generality, we can combine the constraint and operator elements by assuming that a condition-action format may be specified (IF constraint-condition THEN action) that prevents illegal actions from taking place. We refer to operators in this form as rules.

The levels of operator space structure, presented in Table 3, refer specifically to how operators are controlled during task performance. They differ from constraints in that the former determine whether or not an operator can be legally applied whereas the latter determine what to do when two or more operators that meet their individual constraints are available, referred to as the conflict resolution problem. This problem is nearly ubiquitous except in highly structured tasks. It is also present when the opportunity to apply operators in parallel is available (e.g., a physician may choose to order multiple tests from the same blood sample). As was the case for state complexity, the operator space also offers a path towards lower complexity. If, after repeated task performance, the same sequence of operators or rules is repeatedly invoked as a result of conflict resolution, they may be combined into a single program. Such a program synthesizes the entire collection into a structure resembling a logical flow chart, thereby eliminating the need for conflict resolution while it is active. The program may then, in some cases, interact with the state space such that all the values that are needed to determine the specific path to be taken within the program are established *prior* to starting the program. Doing so allows the program to be reconstructed as a set of separate sequences. By removing control knowledge (i.e. tests) from the op-

erator sequence, we meet the prerequisites for automatization (Shiffrin & Dumais, 1981) of the sequence into a single compiled operator with continuing practice.

**Table 3: Operator Space Complexity Levels, with lower levels being associated with more practiced tasks and greater expertise in task performance.**

<i>Operator Complexity Level</i>	<i>Name</i>	<i>Description</i>
1	Sequence	One or more operators that are applied without any intervening control or conditions
2	Program	A collection of operators that may contain different branches—regions of operators that can be invoked only if specified criteria are met, also referred to as control structures. A program may be interrupted as a result of the state space values but is normally followed to completion.
3	Task-specific conflict resolution	Choice of operators is informed with task-specific knowledge that may include experience with past choices in similar situations, considerations of the cost of operator invocation, and probabilistic knowledge of likely outcomes.
4	Weak methods	General methods of operator choice that are applicable to many tasks, such as search strategies (Newell, 1991) and a variety of heuristics, such as elimination by aspects (Payne, Bettman, & Johnson, 1993).

Some types of task may not be amenable to progress towards automatization. The huge number of possible board configurations in chess, for example, dramatically limits the degree to which operator space structure can be established (indeed, intermediate players appear to exhibit patterns of search that are qualitatively similar to those of experts, only not as effective; Charness, 1991, p. 44). That will not prevent some aspects of the task being automated, however. For example, chess players often employ very similar openings from game-to-game, implying that the opening portion of the game can become increasingly programmed as familiarity with the range of openings grows. Mathematically oriented tasks, on the other hand, may become so automated that intervening steps are totally omitted or forgotten. Tools and specialized knowledge may provide similar shortcuts. For example, math programs, formulas, and integral tables may allow a task performer to move directly from a math problem to its solution without any attention being paid to the underlying derivation.

The uncertainty-related component to complexity in the operator space results from the uncertainty associated with an operator's impact on the state space. The levels are summarized in Table 4. Completely analogous to state space uncertainty, some of these uncertainties may be inherent to the task (e.g., your opponent's moves in chess will always be uncertain at the time you move your piece) whereas others may become more structured with experience. For example, when experience shows that open ended responses to the final item of a questionnaire all fall into a series of categories, the "interpret the response" operator for that question could change from an uncertain state set to a determined state set.

**Table 4: Operator Uncertainty Levels, often a function of the task itself**

<i>Operator Uncertainty Level</i>	<i>Name</i>	<i>Description</i>
1	Deterministic	Application of an operator leads to a completely predictable impact on the state space. For example, a manager may choose to set the price for a given product as part of the product management task.
2	Determined State Set	Application of an operator leads to a state space that is an element of a known set of state spaces. For example, when you move a piece in chess, the resulting state space will be a member of the set of possible board positions that can be reached after your opponent's next move. Similarly, when a nurse takes a patient's temperature, the resulting set of possible state spaces consists of states that are identical with the exception of having differing temperature values.
3	Uncertain State Set	The operator affects a known set of state space values, but there is no way to bound the set of possible states that result. For example, if an individual is assigned the task of writing an essay, there is no obvious way to establish the set of all possible essays that could be written.
4	Indeterminate	The impact of the operator on the state space cannot be predicted prior to being invoked.

### **Goal Space Structure**

The goal space component of the problem space consists of an evaluator—which we will call the *goal fitness operator* or, for short, the *goal evaluator*—that allows the relative fitness of the state space to be assessed as the task is being performed. In essence, it evaluates the degree to which progress is being made towards the task goal and helps identify what actions will lead to further progress. In the latter context, the goal evaluator can be accessed by control systems within the operator space as part of the process of determining what operator to apply. It can be vital, for example, in pruning large search trees.

Consider the following examples. In chess, if the goal evaluator strongly rates a particular move as unpromising, there is no need to mentally consider the series of opponent responses and counter-moves that might result from making the move; the task performer simply moves on to considering the next move. In medical diagnosis, if the goal evaluator perceives that symptoms determined so far are a good fit with a particular disease, further diagnostic efforts can then be directed towards verifying that conclusion. When confronted with a large number of alternatives (e.g., buying a used car), the goal evaluator can also be used to eliminate alternatives from consideration (e.g., Payne et al., 1993). The availability and quality of such a goal evaluator can lead to dramatic differences in task performance. For example, it has been demonstrated that chess grandmasters search less intensively than master-level players (Charness, 1991, p. 44), implying they are better able to abandon moves without considering them further. This would be a natural consequence of a superior goal fitness operator.

While it would be possible to place such a goal fitness operator in the operator space (as was done by Card et al., 1983), there are a number of benefits to treating goals separately. Specifically, as mentioned earlier, goals will play a more prominent role than instructions to the performer for nearly all tasks that are not highly structured. A separate goal space facilitates mapping such a task to the problem space. Second, a number of highly influential models suggest that the goal is the principal pathway for establishing motivation (Bandura, 1991, p. 264; Locke, 1978). By explicitly incorporating goals into our model, we therefore incorporate a motivational element to the



problem space. As will be discussed later, this motivational aspect of task performance provides insights into the receptivity of the task performer to different types of informing.

The goal evaluator closely corresponds to the utility function, widely used in economics and decision theory. Expected utility is the conceptual source of decision-maker preference; if you choose one alternative over another, the presumption is that the chosen alternative had higher expected utility than the other. Extending this concept to all actions, not just economic choices, you have the goal evaluator.

A learning model of utility has been proposed that maps well to the goal space levels (Gill, 2008a). Under the learning model, as individual familiarity with a choice grows, utility migrates along the following path:

- *Generic utility*: General purpose goal feature detectors—analogueous to the specialized detectors often proposed for other cognitive systems, such as word recognition (McClelland & Rumelhart, 1981)—identify characteristics in possible alternatives that are widely applicable to different settings, such as gains, losses, uncertainty, and consumption opportunities. These send signals to accumulators where they are summed and a final utility level is established.
- *Task-specific goal utility*: Decision-specific goals and goal feature detectors are acquired that send signals based upon the degree to which an intended choice (or action) appears to satisfy active goals related to the specific task case.
- *State-based utility*: As decisions become automatic, specific reference to individual goals becomes no longer necessary to make choices. Instead, direct linkages from state values to utility are established. Here, it is specifically hypothesized that making progress towards goals becomes the source of utility (Gill, 2008a).

For our present purposes, only one refinement to the model is required to make these levels consistent with the other problem space components (state and operator spaces). From a complexity perspective, it makes sense to divide task-specific goal utility into two levels. The more structurally complex level exists where multiple, interrelated goals are present. The less complex level exists where a single goal or fully decomposable individual goals (i.e., goals that don't interact with each other) are present.

The justification for separating the two regions is that the associated goal spaces will tend to be, qualitatively, very different in structure. This draws upon the notion from evolutionary biology of a fitness landscape (Kauffman, 1993), which maps entity characteristics (such as genes or animal traits) to a value that signifies the survivability or desirability of the combination (Gill, 2008b). Where a single goal exists, or where individual goals can be pursued independently, the utility function will have a single peak. From a task performance perspective, this means that by choosing operators that increase expected utility, you will nearly always find yourself moving towards the desired goal. Where multiple goals that interact are present, on the other hand, you are far more likely to get a rugged fitness landscape (Kauffman, 1993). This means that choosing operators by optimizing utility locally only guarantees a local peak. Very often such local peaks will not exhibit very good fitness. The hypothesized goal space levels are presented in Table 5.

To make the ruggedness example more concrete, consider the game of chess. In the very early stages of learning the game, novices are taught to associate point values with pieces (e.g., pawns, 1; bishops and knights, 3; rooks, 5, queen, 10). This immediately leads to the creation of a task-specific goal—maximize points relative to your opponent. While this goal is generally sound, it is only one of many heuristic goals that are used to guide chess play (e.g., maintain control of the

center of the board is another). Often, during real games these goals will conflict. If you only worry about the maximizing the point count goal, you are unlikely to succeed against a strong player. Indeed, once your opponent recognizes that you are focusing exclusively upon that goal, he or she will continually tempt you with sacrifices that—if taken—will ensure your defeat.

**Table 5: Goal Space Complexity Levels, with lower levels being associated with more practiced tasks and greater expertise in task performance.**

<i>Goal Complexity Level</i>	<i>Name</i>	<i>Description</i>
1	Progress	Identifies the degree to which the current state space represents a completed task. Goal fitness is derived directly from the state space.
2	Decomposable	Goal fitness is a function of one or more separable goal frames within the state space which may be additively combined. Goal fitness directs the task performer towards a single peak.
3	Rugged	Goal fitness is determined by activation signals sent by individual competing goal frame agents. It can be modeled by a non-decomposable utility function that is heavily influenced by interrelationships between state values, with multiple local fitness peaks being present.
4	Generic	Goal fitness is a function of generic feature detectors that look for attributes of alternatives, such as gains, losses, and uncertainty that are widely applicable in many circumstances. As such, suitability for a specific task may be weak.

With increasing expertise, the task performer will acquire the ability to weigh alternative goals as appropriate for different situations. This requires two types of knowledge: knowing where the local peaks are and knowing how to weigh them appropriately under different circumstances. At this point, the goal space comes to more closely resemble a single function.

The process through which goal space structure emerges through practice closely mirrors that proposed for the state and operator spaces. Initially, unstructured task performance would utilize generic goal detectors along with any task-specific goals available to the task performer as a result of specific goal information incorporated into the task itself and prior informing with respect to specific task performance or theory. Over time, particular patterns of agent activation leading to successful task performance would be chunked into more sophisticated task-specific goals, applying set weights to the contributing goals and sending suppressing signals to those goals found not to be relevant. The composite goal frames would require vastly less attention than the generic-goal based approach and would gradually emerge and tune themselves to all routine task cases. Finally, task performers would begin to map these goal evaluators directly to state-space values—a process likely to be facilitated by parallel development of a large library of chunked schema patterns in the state space, as already discussed—eliminating the need to consider goals consciously except for the purpose of monitoring task progress.

In the goal space, the analog to state and operator uncertainty levels is goal dynamics levels, presented in Table 6. The idea here is that during the course of performing a task, priorities (Level 2) or goals (Level 3) may change in ways that significantly impact task performance. As was the case for the other uncertainty levels, this may be a consequence of the nature of the task itself and not just the problem space.

Goal dynamics are likely to present the greatest challenge in tasks that take place over an extended period of time. For example, during a multi-semester-long redesign of an introductory programming course (Gill & Holton, 2006), the initial goal criteria included designing a course that could: 1) accommodate diverse student backgrounds, 2) allow a single terminally-qualified

instructor to preside over many students (an accreditation-related criterion), 3) accommodate the unpredictable growth in the department's MIS major (which had been explosive in prior years), and, to a lesser extent, 4) act as a gatekeeper for subsequent courses in the MIS major. Over the duration of the two year redesign, however, severe contraction in the program's enrollments (which went from 1100 MIS majors to around 200) caused criteria (2)-(4) to be abandoned and replaced with (5) encouraging retention.

**Table 6: Goal Dynamics Levels, representing the degree to which goals change over time.**

<i>Goal Dynamics Level</i>	<i>Name</i>	<i>Description</i>
1	Static	Goals are unlikely to change during task performance.
2	Fixed Goal Set	Goals may change during task performance, but the nature of all possible goals is known at the outset of the task, meaning changes can be modeled in terms of priorities.
3	Uncertain Goal Set	Goals may change in some meaningful way during task performance but the precise nature of possible goals cannot be ascertained.
4	Indeterminate	It is unknown whether or not goals can change and the nature of such change cannot be predicted.

The principal approaches to reducing goal dynamics are to reduce the time taken to perform the task—thereby reducing the window of opportunity during which goals can be altered—and to focus on attaining goal peaks that are anchored to less dynamic criteria. For example, prioritizing task goals that maximize contribution to an organization's strategy (which, hopefully, changes less rapidly than its tactics) may reduce the volatility experienced during task performance.

### ***Integrated Model***

The concepts presented can be synthesized into an overall problem space model, such as that presented in Figure 2. There are, essentially, three dimensions that can be used to categorize task-related problem space knowledge:

1. The structural level (the Y-axis of the diagram), with lower levels generally corresponding to higher experience-based expertise (theory based expertise is considered in the next section.) Level 4 is omitted since it is not task specific.
2. The specific subspace (the Z-axis of the diagram)
3. The applicability of the knowledge, with some knowledge being broadly applicable across task cases (core knowledge), while other knowledge is more task case specific. (For the sake of illustrative clarity, the diagram somewhat oversimplifies this characteristic, since some knowledge may be too specific to be considered core yet may map to more than one task case.)

The relative importance of task case knowledge versus core task knowledge is likely to be a function of the ruggedness of the task fitness landscape. High levels of ruggedness, generally resulting from large numbers of interacting task elements, will produce many alternative solutions to a task that can vary widely from case to case. In such landscapes we would expect to see knowledge of specific task cases dominate core knowledge (Gill, 2008b). Where the task is highly decomposable, on the other hand, knowledge of task components will apply across the task landscape, making core knowledge much more generalizable across the problem space.

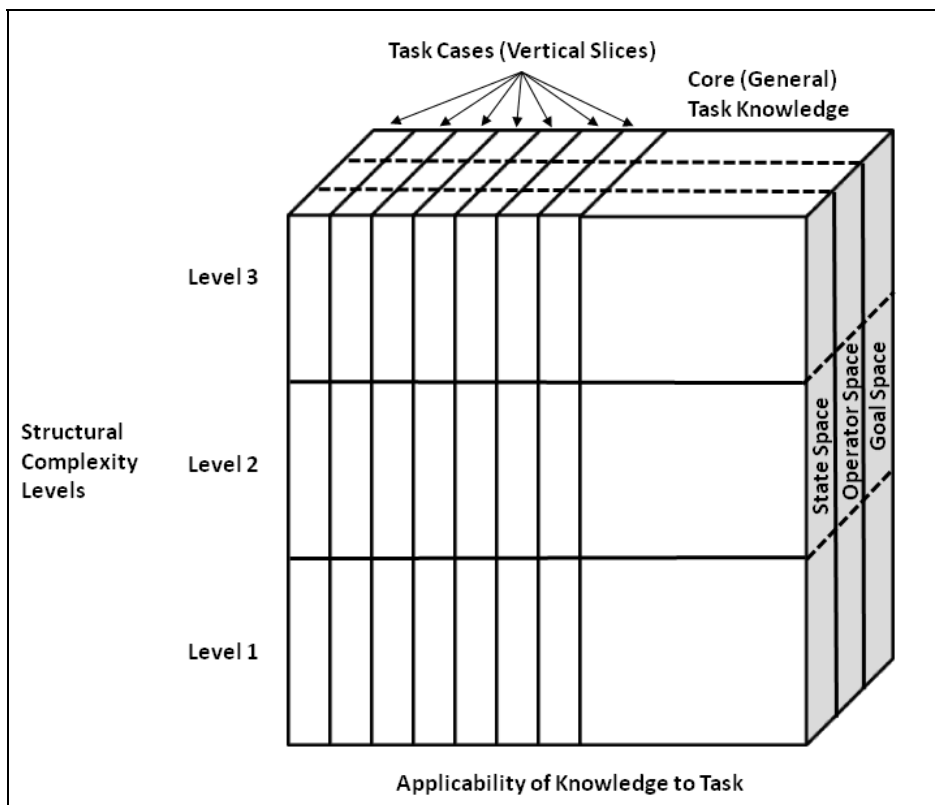


Figure 2: Problem Space Structural Complexity Model

## Issues in Informing and Structural Complexity

The remainder of the paper addresses how structural complexity impacts informing processes. Because such informing involves the interaction of sender, task, and client in a context that involves adaptation (i.e., changes to the client problem space), the goal space will not be easily described or predicted without detailed study of individual task cases and performers. The nature of the problem does suggest, however, a certain number general informing issues. Three of these involve: 1) awareness of prior client knowledge, 2) understanding the particular bottlenecks to performance of the task being conveyed, and 3) understanding the nature of the client's motivation for being informed. These are discussed briefly in the current section. In the subsequent two sections, we expand upon the discussion by considering some heuristics that appear to apply generally to complex informing situations, then address the particularly challenging problem of informing across very dissimilar problem spaces.

### ***Prior Client Knowledge***

If we are a sender seeking to inform a particular client, understanding that client's initial problem space (i.e., the problem space we are seeking to change) is bound to be important. Obviously, the informing process can be streamlined if the sender knows what knowledge the client already has. In a curricular setting, for example, this may be accomplished by having prerequisites to a course. Such knowledge may not be limited to task-specific elements, however. For example, differences in vocabulary or conflicting prior knowledge (Bain, 2004) may inhibit informing. We return to this topic later, when we consider the general challenges of informing across dissimilar problem spaces.

## **Location of Task Performance Bottlenecks**

For many tasks, bottlenecks to task performance tend to occur in specific regions of the problem space. Consider some of the examples presented in the previous sections:

- *Physics problems*: Understanding of elementary physics will nearly always involve growing the state space through the acquisition of theories (Level 3) and specific schemata structures (Level 2) that identify when the specific constraint of an operator (e.g., a formula) has been met. Operator control and goal fitness tend to be far less important.
- *Medical diagnosis*: For most medical diagnosis problems, the problem space will generally be dominated by state schemata (i.e., the ability to recognize patient symptoms) and control knowledge (i.e., the ability to determine what questions to ask and tests to run in an environment where both time and costs are likely to be highly constrained).
- *Chess*: Chess effectiveness is particularly dependent upon the sophistication of the goal space, which allows us to eliminate moves from consideration in order to avoid the combinatorial explosion of possible moves as we look ahead. Such fitness is likely to be greatly enhanced by the accumulation of state space schemata holding patterns of moves (e.g., past games, games described in magazines).

As a general rule, we may speculate that the most effective informing is going to occur with respect to the task performer's perceived bottlenecks.

## **Client Motivation**

The client's motivation towards being informed is also likely to be a critical determinant of what types of informing will be effective. Using our model, three general types of task performance improvements might be the goal of informing, either individually or in combination:

1. *Reducing structural complexity*: This involves transforming a particular task case to a more structured form or reducing the overall structural complexity of the task (i.e., increasing the average level of structure across all task cases).
2. *Increasing task outcome quality*: This involves coming up with solutions having higher goal fitness. For example, in a complex goal task, being informed about a peak with higher fitness than the client's existing peak could lead to an improvement in outcome if the informed client then targets the new peak.
3. *Increasing performance efficiency*: This involves improvements to performance such as reducing costs (both economic and cognitive) and increasing speed of task performance.

Client motivation extends well beyond the desire to achieve pure task performance improvements. In fact, the psychology and education literatures find that client motivation for informing can vary considerably across two dimensions (Elliott & Harackiewicz, 1996): mastery vs. performance and approach vs. avoidance. For mastery goals, motivation is largely intrinsic whereas performance goals are often motivated by external factors (e.g., more pay, better grade) as well as by intrinsic factors. We might speculate that mastery goals would make informing across task cases particularly attractive, whereas performance goals would tend to be focused towards achieving better performance on those specific cases being performed or anticipated. For approach goals, the motivation is directed towards enhancing performance whereas avoidance goals are motivated by preventing or ameliorating performance failures. We might therefore anticipate

that the performance-avoidance combination (far more common than the mastery-avoidance combination) would be particularly motivated by informing related to specific task cases.

## Heuristic “Laws” of Informing

The existence of intrinsic motivation for being informed has significant implications with respect to informing processes and structural complexity. Some of these implications are captured in three empirically-derived principles related to informing that are presented in this section. The first, the *Law of Abandoned Expertise*, relates generally to human informing activities. The remaining two, the *Law of Limited Visibility* and the *Law of Low Hanging Fruit*, specifically relate to informing situations involving tasks exhibiting a rugged fitness landscape (Level 3 Goal Space).

The particular significance of rugged fitness landscapes is examined in considerable detail in another paper (Gill, 2008b). One of the key conclusions of that paper is that rugged landscapes might be quite common in informing systems. The reasoning behind this conclusion was as follows:

- Many elements appear to have the potential to impact informing effectiveness (both sender and client motivation, client problem space contents, sender awareness of client problem space contents, task characteristics, delivery system characteristics, and so forth). Thus, the fitness function is likely to have many arguments.
- Most of the arguments to the fitness function would appear to require fit with other arguments. For example, if a client’s learning style leans towards visual presentation, we would expect more effective informing to occur if the system employs graphic display. Thus, we predict many interrelationships between fitness variables.
- Particularly where technology and globalization trends impact the system, the nature of the fitness functions will be continually changing.

These three attributes happen to represent the principal prerequisites for a rugged fitness landscape. Thus, such landscapes should be the rule, rather than the exception, when the fitness of an informing system is considered in conceptual terms.

### **Law of Abandoned Expertise**

*Hypothesis:* Clients will resist any task-related informing activities that require relinquishing existing expertise in their problem space.

Throughout the development of the problem space complexity model, expertise has been presented as an accumulative process of moving from higher structural complexity towards lower structural complexity. For such a process to continue, it would necessarily have an implicit motivational element. The likely corollary, then, would be that moving towards from practiced expertise towards higher structural complexity would be demotivating in character.

The reason for stating this law explicitly is that the oft cited “resistance to change” is overly general and does not reflect the real world very well. In knowledge diffusion contexts, some changes are observed to be welcomed. Rogers (2003), for example, described the remarkable pace of diffusion for the drug tetracycline in the 1950s. Key to this rate of adoption was the fact that the drug was nearly identical to existing antibiotics in its application, differing only in its significantly reduced side effects. What the hypothesis states, then, is that informing activities requiring individuals to “give up” their previously acquired expertise are those most likely to experience resistance.

Perhaps the most compelling evidence for the hypothesis comes from Kuhn's (1970) seminal research into the nature of scientific revolution, where it is noted:

...the act of judgment that leads scientists to reject a previously accepted theory is always based upon more than a comparison of that theory with the world. The decision to reject one paradigm is always simultaneously the decision to accept another, and the judgment leading to that decision involves the comparison of both paradigms with nature *and* with each other. (p. 77)

Wholesale acceptance of informing that requires the acceptance of another paradigm—a form of schema restructuring (Rumelhart & Norman, 1981)—is also likely to entail abandoning a substantial fraction of compiled knowledge (e.g., chunks, automatized processes), suggesting that the task performance benefits of being informed would need to lead to dramatic improvements when compared with the initial problem space.

Evidence for the Law of Abandoned Expertise also appears at the novice level. For example, faculty members teaching an introductory physics course were disturbed by exam results that suggested their students were holding fast to their existing (Aristotelian) model of motion, despite their having been taught the more accurate Newtonian perspective (Halhoun & Hestenes, 1985, as cited in Bain, 2004). To counter this, the faculty conducted individual interviews with students, devised experiments to demonstrate the inadequacies of the student-held models, and conducted experiments that refuted Aristotelian notions in front of the students. They then asked the student to explain the results.

What they heard astonished them: many of the students still refused to give up their mistaken ideas about motion. Instead they argued that the experiment that they had just witnessed did not exactly apply to the law of motion in question; it was a special case, or it didn't quite fit the mistaken theory of law that they held as true... The students performed all kinds of mental gymnastics to avoid confronting and revising the fundamental principles that guides their understanding of the physical universe. Perhaps most disturbing, some of these students had received high grades in the class. (Bain, 2004, p. 23)

The implication here is that even novice task performers may perceive themselves as having some task-specific knowledge in their problem spaces that they are reluctant to relinquish.

### ***Law of Limited Visibility***

*Hypothesis:* In the absence of concrete negative performance feedback or external pressures, an individual will gradually come to view the entire goal space in terms of the peak that he or she has reached or is presently climbing. The phenomenon will be particularly pronounced in a goal space where multiple peaks exist, but only one provides feedback.

This hypothesis, which addresses the issue of loss of perspective, builds upon the same basic assumption as the previous hypothesis, that the acquisition of expertise is intrinsically motivating. The underlying reasoning is as follows. As you acquire expertise in a rugged goal space, the process will naturally begin by choosing a peak to focus on (Level 2) then, ultimately, getting to the point where you don't explicitly think about the goal at all but only the progress you are making towards completing the task (Level 1). The presence of feedback that suggests you are not doing well, or external pressures to increase your exploration of the goal space, will motivate you to continue thinking about the goal space as a whole. In the absence of such pressures, however, you will naturally come to believe that your way of viewing the task is the "correct" way (a single peak perspective).

The predicted impact of the availability of feedback for only one peak is easily explained. In low structure situations, feedback regarding what constitutes a "good" strategy improves performance

(Payne et al., 1993). As a consequence, expertise can develop more rapidly in such contexts. As we are motivated by the acquisition of expertise, it therefore makes sense that we would gravitate towards solution fitness peaks providing such feedback.

There are a number of additional observed cognitive mechanisms that would tend to cement our loss of general perspective once the process is set in motion. The *availability bias* results in our being excessively influenced by examples that we can readily recall—which would all tend to arise from our activities relating to the individual peak (Reber, 2004). The *false consensus effect* (Ross, Greene, & House, 1977) means that we will tend to believe others see the world the way we do, meaning we will perceive that our peers view the task in the same single-peak terms that we do unless they explicitly inform us to the contrary. The *confirmatory bias* means that we tend to maximize the weight that we place upon incoming evidence that confirms our view of the task and ignore or minimize evidence to the contrary (Oswald & Grosjean, 2004). The *representativeness bias* (Teigen, 2004) means that we will tend to view information that partially supports our view as being completely supportive of our view. These may, of course, also be taken as supporting the built-in resistance to loss of expertise previously hypothesized.

On the positive side, the Law of Limited Visibility offers the promise of task performance benefits, at least in the short run. All the processes associated with knowledge compilation result in increasing speed combined with reduced cognitive demands. Thus, focusing on a single peak and acquiring expertise as rapidly as possible will promote task performance efficiency.

The potential risks involved with the single peak stem from our entire notion of evolutionary fitness. When entities become so specialized that they cease to adapt, they are extremely vulnerable to changes in the environment. An interesting example of this vulnerability involves two companies that, in the mid-1980s, were widely touted as being the future of business: Mrs. Fields' Cookies and Batterymarch Financial Management (Gill, 1995). In both cases, the companies had achieved acclaim through their innovative informing systems. In the case of Mrs. Fields, they implemented a centralized system that controlled cookie production, scheduling, record-keeping, and even facilitated hiring at all their store locations. Essentially, it eliminated the need for expertise at the store manager level—which was good, since the average tenure of store managers was less than a year. Using this system, they achieved spectacular growth in sales and units without needing to franchise. Batterymarch managed investments and, in similar fashion, developed a fully automated portfolio system—based on the founder's contrarian philosophy—and the first automated trading system, allowing the firm to eliminate costly human traders. By the early 1980s, the Batterymarch had experienced a huge increase in assets managed, had established an extraordinary track record of beating the market averages, and had a managed-assets-per-employee ratio on the order of ten times the industry average.

In both cases, then, the companies had profited dramatically by employing technology in place of human expertise, in support of their respective founder's particular view of the world (problem space). From an adaptive systems standpoint, however, this caused a problem. The types of individuals the systems replaced (experienced managers and franchise owners for Mrs. Fields; human traders for Batterymarch) were precisely the type of individuals whose role included environmental sensing and who would typically be the first to identify fundamental changes taking place within the competitive environment. Thus, when the stock market transformed from its lackluster days of the 1970s and early 1980s into the bull market that commenced in 1982, Batterymarch continued to try to fine tune its portfolio using its original contrarian model instead of restructuring its approach and seeking an alternative peak. In doing so, it achieved a track record that was as breathtakingly bad as its previous record was good. In the case of Mrs. Fields, changes in consumer taste, the role played by malls, and a bi-coastal recession that hit its prime areas in the late 1980s fundamentally changed the cookie business and their subsequent attempts to tune their existing models to the new environment proved ineffective. As a result, the company suffered huge



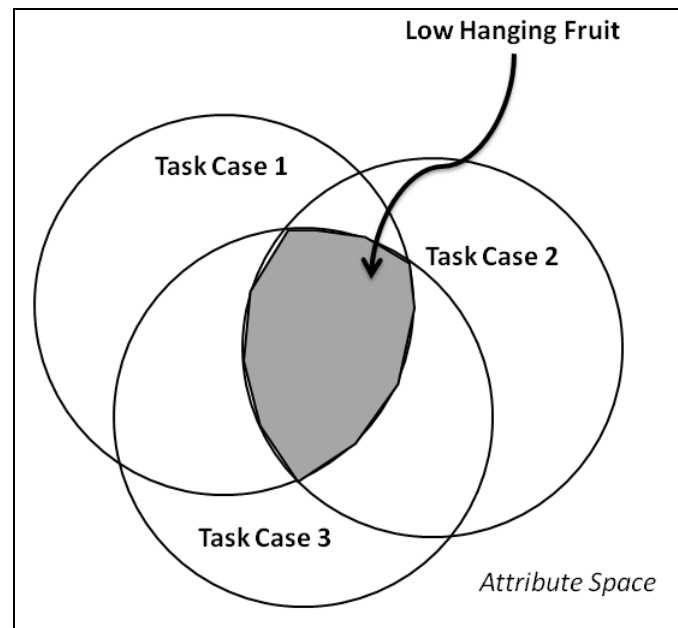
losses and the Fields lost control of the business that they had founded. The cautionary lesson provided by the two stories is an important one: the efficiencies gained from exclusive focus on a single goal in a rugged goal space come only at the expense of adaptability. Should the environment change, the original problem space will be brittle.

### **Law of Low Hanging Fruit**

*Hypothesis:* Within a rugged goal space, those problem space attributes that enhance or detract from goal fitness decomposably across nearly all fitness peaks will tend to obscure equally important contributors to fitness that only act upon certain specific peaks.

This hypothesis particularly relates to the challenges associated with acquiring expertise across a task domain that includes a rugged fitness landscape (i.e., multiple fitness peaks, Goal Level 3) for task performance, as opposed to domains with decomposable (Goal Level 2) fitness landscapes.

Within rugged task domains, there are likely to be some attributes that influence fitness in a decomposable manner, collectively referred to as *low hanging fruit* (Gill, 2008b, Gill & Sincich, 2008). Viewed from a task case perspective, these are the attributes that are common to all task cases, as illustrated in Figure 3. From a theory-building perspective, these attributes are very attractive. One of the principal criteria for judging the quality of a theory is its compactness. Prior to Copernicus, for example, it was perfectly feasible to predict the position of the planets in the sky—it simply required a very elaborate model based on the assumption that the earth was the center of the universe. The solar centered model, in contrast, was vastly simpler in its construction. Thus, reinforced by principles such as Occam’s razor, when attempting to construct a theory that applies to all task cases, we would prefer that *relatively few* task attributes actually determine fitness. Thus, we tend to *accept* state space values as making a contribution to goal fitness only on the basis of a large body of evidence, which may include both models and observations. When presented with a series of task cases, each having a set of attributes that appear to be significant contributors to goal fitness, the most attractive and generalizable theory would consist of the *intersection* of those attributes. Any theory encompassing additional attributes is going to involve qualifications and contingencies that undermine the compactness and elegance of the theory.



**Figure 3: Low Hanging Fruit in Attribute Space**

If, on the other hand, the goal is not to develop generalizable theory but rather to maximize task performance on a case-by-case basis, the opposite strategy may emerge. With each new case, additional attributes that impact task fitness will be identified and, rather than casting them out, we would tend to incrementally add them to our existing schemata. Effectively, then, we end up with a problem space that involves the *union* of those attributes. It has been observed, for example, that practitioners often develop heuristic models in their personal problem space that are unnecessarily convoluted. In fact, regression analysis frequently outperforms expert judgment on well structured tasks, an observation referred to as the “process-performance paradox in expert judgment” (Camerer & Johnson, 1991, p. 195). Moreover, the resulting scoring models often prove to be considerably more parsimonious than the approach employed by the original experts. Gladwell (2005, pp. 125-136) relates a story of how a researcher, through the study of numerous sample cases, developed a simple 3-factor decision tree for predicting whether or not a patient with chest pains was having a heart attack that was 95% accurate, compared with a 75-89% accuracy range for experts who used substantially more information.

### **Informing Between Dissimilar Problem Spaces**

The structural complexity model would predict that the process of informing novice clients—who have a very small task-specific problem space—will be very different from that of informing expert clients—who enter the process with a large existing problem space. In addition, the expert-to-expert informing process is likely to be heavily impacted by the qualitative differences between the sender and client problems spaces. We now consider how these differences could impact the informing process.

### **Informing Novice Client Problem Spaces**

A thorough treatment of the process of informing novices would, quite obviously, involve revisiting nearly the entire body literature generated by the field of education. Equally obvious, such a review is beyond the scope of the present paper. Our objective in this section, therefore, is the considerably more modest task of framing the novice client informing process in terms of structural complexity.

What makes informing novices different from informing experts is our ability to assume little or no pre-existing task-specific knowledge in the problem space. (Although, as our previously cited example relating to learning the Newtonian laws of motion suggests, even this assumption may be somewhat suspect in real world settings). Thus, the conceptual challenge of informing the novice revolves around building up the problem space from scratch.

There would seem to be three general approaches that could be employed to support development of the problem space:

1. *Top-down approach.* In this approach, structure is built from higher levels (e.g., theory, models, control knowledge), works down through intermediate levels (e.g., structures, programs), and, finally, clients are presented with real-world schemata and operators. This approach roughly corresponds to the traditional teaching approach that begins with lecturing and ends with examinations.
2. *Bottom-up approach:* In this approach, clients are presented with a set of rich task examples—consisting of realistic schemata and operators—and are guided through the process of inducing their own theory. This very closely corresponds to the *constructivist* model, currently very popular in education.

3. *Vertical-Slice approach*: In this approach, whose name is suggested by the task case knowledge presentation of Figure 2, complete problem spaces (all elements and levels) that permit the solution of a specific problem (task case) are provided in a predetermined sequence. These slices become successively more challenging, based upon the task cases being considered—a process sometimes referred to as scaffolding—thereby allowing the evolution of a more comprehensive problem space. Such a *problem-based learning* approach is fairly typical in the math and science domain.

As previously noted, the rugged goal space concept that is central to the present paper is that multiple fitness peaks are virtually inevitable in problems involving a particular sender informing a particular client about a particular complex task. In such situations, it would therefore be very likely that different strategies would prove to be optimal for different clients. Thus, rather than advocating a single approach, we will consider, from an abstract standpoint, the likely strengths and weaknesses of each approach. These are summarized in Table 7.

**Table 7: Strengths and Weaknesses of Approaches to Novice Informing**

<i>Strategy</i>	<i>Strengths</i>	<i>Weaknesses</i>
Top-Down	Allows the sender to maintain a high level of control over the client's problem space—particularly with respect to keeping it from becoming too well established (and therefore hard to change) prior to providing all information. Feedback is generally available, through testing.	Since substantial motivation may derive from successful task performance, clients may become demotivated over the long period it takes before they are ready to solve realistic problems. Senders are likely to have difficulty establishing a high level problem space that is suitable for novice clients.
Bottom-Up	Allows the client to immediately begin developing realistic task schema and, shortly thereafter, begin task performance. The requirement that the client induce his or her own abstract models over time reduces the likelihood that incorrect models that need to be changed will be frozen into the client problem space. Provides significant opportunities for reinforcement through peer-to-peer informing.	Extremely sensitive to the nature of the task examples presented. In task domains where a verifiably “correct” body of theory exists, there is no guarantee—or even likelihood—that the correct theory will be induced by the client. Senders must generally have deep knowledge of all the task cases presented. Feedback is generally ambiguous, often intentionally so to avoid driving the theory-creation process.
Vertical-Slice	Allows client to begin task performance very early in the informing process, providing a source of continuous motivation. The fact that information is continuously being transferred at all levels of structure may allow a wide range of client learning preferences to be accommodated. Feedback is generally available through exercises.	Clients may begin to freeze knowledge based upon “toy” problems. May be difficult to construct independent slices in complex task domains. Lack of feedback in task domains where goal criteria tend to be fuzzy or dynamic may make constructing and learning from vertical slice problem spaces difficult.

The client's motivation for seeking informing is also likely to have a significant bearing on the most appropriate technique. As previously noted, three distinct motivational profiles are typically encountered with populations of students (Elliott & Harackiewicz, 1996): mastery motivation (e.g., “I want to learn the material for its own sake, or so I can do things in the future”), performance-approach motivation (e.g., “What do I need to do in order to get an A?”), and performance-avoidance (e.g., “What do I need to do in order to get a passing grade?”). As a general rule, the two performance motivations may be less compatible with the ambiguous feedback provided by the bottom-up approach.

The nature of the sender's problem space can have a huge impact on appropriate approaches. Even when limiting ourselves to considering the case where the sender is an expert, a clear distinction needs to be made with respect to the form of expertise in the sender's problem space. Aristotle characterized knowledge as coming in two forms—*episteme* (theory) and *phronesis* (practical wisdom) (Kessels & Korthagen, 1996). In the structural complexity model, *episteme* would exist mainly at the upper levels. This type of knowledge is well communicated symbolically and is of a type prized by scholars. It would also be most commonly seen in individuals who know a great deal about a task, but who do not perform the task very often (frequent performance leading to the accumulation of more lower-level knowledge). Thus, we refer to individuals for whom *episteme* dominates as *academic-experts*. *Phronesis*, on the other hand, manifests itself most prominently in the *practitioner-expert*, which is to say an individual whose expertise has been acquired as a consequence of frequent performance of the task.

Where the sender is a practitioner-expert, the top-down approach will often be a poor fit. The problem here is that the problem-space of the sender is likely to be in precisely the wrong form for novice use. The knowledge engineering paradox states that the more expert we are at performing a given task, the more likely that we will have trouble explaining what it is that we are doing (e.g., Waterman, 1986). It is completely consistent with the path towards expertise described earlier, namely that repeated task performance moves us towards increasingly opaque schemas (chunks) and operators (automatized sequences) and what limited higher level models and control knowledge remain serve mainly to deal with unusual, non-routine cases. To make matters worse, experts are singularly bad at predicting how long it will take novices to acquire expertise, a finding consistent with the compiling of knowledge anticipated in the structural complexity model and referred to as the “curse of expertise” (Hinds, 1999). Completing the mismatch, where a rugged goal space is present, practitioner-experts are also likely to be intimately familiar with a few fitness peaks while, quite often, being unfamiliar with, or even unaware of, others (i.e., the Law of Limited Visibility). While these issues have less impact on bottom-up or vertical-slice strategies, the practitioner-expert sender is likely to require considerable re-education in order to be effectively employing a top-down strategy.

The academic-expert sender, on the other hand, is likely to have far less difficulty applying the top-down approach. Such experts will generally aim their personal learning efforts towards gaining a broader perspective of the entire problem space than would typically be required for the practitioner expert. The potential danger inherent in the broader perspective would be a natural tendency to overemphasize those low hanging fruit that generalize well across the entire problem space, thereby implicitly conveying (perhaps unintentionally) that the knowledge schemata embodied in these fruit represent the only critical elements needed for dealing with any task situation. On the other hand, the academic expert sender may find the depth and ad-hoc quality of knowledge required for bottom-up strategies to be somewhat intimidating, unlike practitioner-expert senders.

### **Informing Expert Client Problem Spaces**

Expert-to-expert informing is a very important process, being a critical element in the diffusion processes for most advanced knowledge—where word of mouth appears to be the dominant means of informing for all but the initial stages of most diffusion processes (Rogers, 2003). It is, however, an extremely difficult problem setting to address in abstract terms. Beginning with the situation where expert problem spaces are similarly structured (e.g., practitioner-expert informing practitioner-expert or academic-expert informing academic-expert), some general patterns—perhaps better described as low hanging fruit—suggested by our task structure model include:

- Informing will generally take place across levels where most expertise resides. Thus we would expect informing between practitioner-experts to be dominated by exchange of

schemata, whereas academic-expert informing is likely to take place at the model and task control knowledge level.

- In complex goal environments, informing will generally be more effective in situations where similar goal peaks are selected. For example, practitioner-experts pursuing very different peaks within the same domain are likely to have evolved very different schemata that inhibit communications. Academic-experts from different schools (e.g., scientists whose model is evolution versus those who believe in intelligent design) face formidable barriers to informing, since accepting premises from another school necessarily means abandoning existing beliefs (thereby violating the Law of Abandoned Expertise).
- Some material differences between similar expert problem spaces must exist, otherwise even extensive communication will likely result in minimal informing (which we narrowly define as communications leading to changes to the client’s problem space).

Summarizing these patterns, informing between comparable experts is probably a matter of achieving balance. It will be most effective where the problem spaces involved are similar enough to ensure effective communications but different enough so that such communications lead to informing. Not surprisingly, this conclusion closely corresponds to what is generally considered to be interesting (Davis, 1971).

Informing activities between academic-experts and academic-practitioners face formidable obstacles in both directions, a number of which may be predicted from the model. As suggested by the distribution of dots in Figure 4, representing the hypothetical presence of knowledge elements in the problem space, pure practicing experts and pure academic experts are likely to look quite different. Practitioners are likely to exhibit considerable depth with respect to the task cases that they perform routinely, plus some special purpose higher level knowledge that is used to address cases that they encounter infrequently. Academic experts, on the other hand, are motivated to acquire broad theoretical knowledge of the task domain, while never practicing enough to acquire the low-level structures necessary for efficient performance.

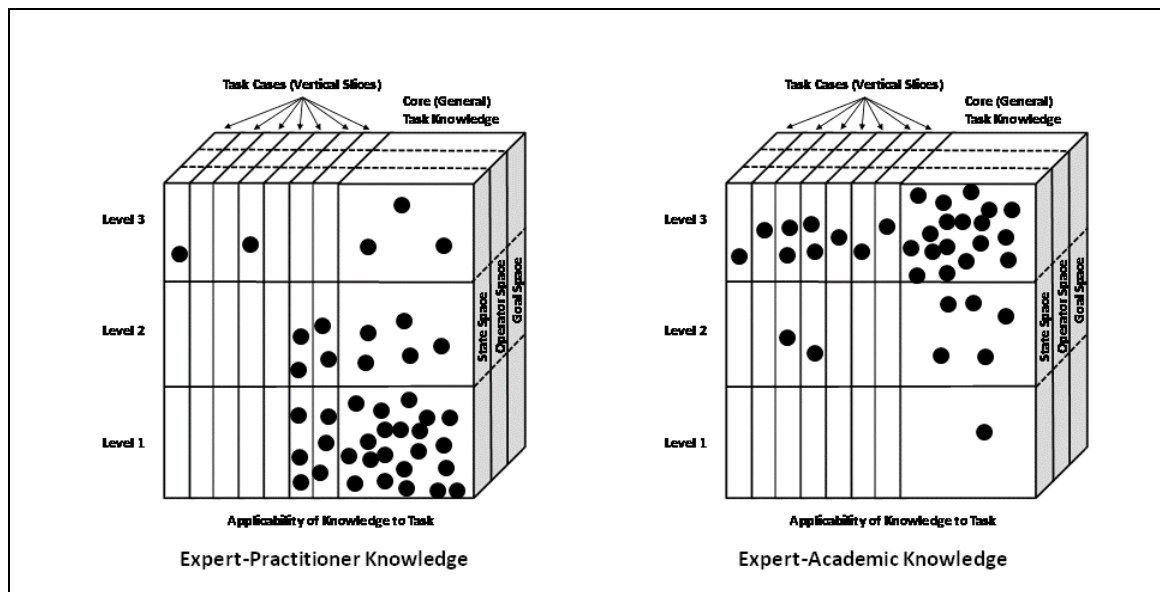


Figure 4: Illustrative Patterns of Practitioner and Academic Knowledge Distributions

## Structural Complexity

The worst barriers to communication between these spaces seem to present themselves where both client and sender, through the process of accumulating expertise, have come to view a rugged goal space (Level 3) as decomposable (Level 2). This would mean, for the practitioner-expert, concentrating on a single goal peak to the exclusion of all others. For the academic-expert, it would involve developing an extensive theory space that generalizes to a broad range of task cases through eliminating relationships not observed in all (or, at least, most) cases—i.e., dominated by low hanging fruit. When the practitioner-expert attempts to inform the academic-expert in this context, barriers from the academic expert client's perspective include:

- While some of the practitioner-expert's key contributors to fitness confirm what the academic-expert already knows, many will represent items that the academic-expert has already rejected as being spurious as a result of a rigorous analysis of the intersection of many task cases.
- The practitioner-expert will gravitate towards informing through the use of highly specialized schemata, none of which seem particularly useful in achieving the academic expert's goal of theory building. When asked to explain, the practitioner-expert will often be unable to do so (Waterman, 1986).
- Accommodating the practitioner-expert's perspective will often require the reformulation of theory or adding special cases that undermine the theory's compactness and elegance.

When informing moves from academic to practitioner, the barriers may prove to be even more formidable. From the practitioner-expert's perspective, the academic-expert's informing is all of the wrong type. For example:

- It is full of low hanging fruit that are of little value to the practitioner-expert because they are obvious.
- It is of the wrong form, since the practitioner-expert relies principally on schemata and operators for task performance and has little use for theory in reaching the particular peak that has been selected.
- It is likely to have at least some elements that are not relevant, or may even be wrong, when applied to the expert's particular task peak. The reason for this, as previously identified in the Law of Low Hanging Fruit, is that the appearance of significant values can be heavily influenced by the nature of the sample that is the basis of generalization. Thus, unless the practitioner-expert's particular task peak happened to be well represented in the observed sample, there may prove to be values in the academic-expert's model whose inclusion or anticipated direction of effect are entirely inconsistent with the practitioner's own knowledge.

These barriers are sufficiently great that we would expect effective academic-to-practitioner informing to be relatively uncommon. In fact, the diffusion literature seems to provide considerable support for this prediction. Consider the following generalizations made by Rogers (2003):

- *The complexity of an innovation, as perceived by members of a social system, is negatively related to its rate of adoption* (p. 257). Advanced academic ideas are frequently complex, especially when viewed from the practitioner's eyes, since they involve knowledge presented in an unfamiliar form.

- *The observability of an innovation, as perceived by members of a social system, is positively related to its rate of adoption* (p. 258). Theory-based informing frequently lacks the concreteness that makes for effective informing (Heath & Heath, 2007).
- *The compatibility of an innovation, as perceived by members of a social system, is positively related to its rate of adoption* (p. 249). Where the innovation is an idea, this compatibility will be with respect to both the idea's form and substance. Where both academic and practitioner perceive that a single peak dominates the landscape, unless it happens to be the same peak, their ideas may be fundamentally incompatible.

All of these factors will make the persuasion stage of adoption—during which “the individual forms a favorable or unfavorable attitude towards the innovation” (Rogers, 2003, p. 216)—particularly critical. The persuasion stage, however, tends to be most effective when conducted through local channels (Rogers, 2003, p. 207), meaning that other similar practitioners, rather than academics, are likely to be most effective in persuading a client practitioner to adopt an idea. This phenomenon also creates the likelihood of a time window after which the potential for academic-practitioner informing processes will be minimized. If the informing process starts in the very early stages of an idea, then cosmopolite channels—that is, channels where the sender and client can be quite different in terms of social network and knowledge—may be reasonably effective (Rogers, 2003, p. 207). If the diffusion of an idea is delayed (as will often be the case for academic ideas given the demands of rigorous reviews and characteristically long publication cycles) then the opportunity for efficient cosmopolite informing may well have passed.

The chances for academic-practitioner informing in the presence of a rugged goal space are likely to improve significantly where both parties accept the inevitability of multiple peaks. From the practitioner perspective, the range of the academic-expert's knowledge of the overall goal space provides the opportunity to consider alternative peaks that might prove to be more attractive than the peaks already well-understood. Such motivation will be particularly strong in where the practitioner-expert client has a strong mastery motivation. While such clients are probably the exception rather than the rule in a general population of practitioners, they are likely to exist—going under a variety of names such as strangers, innovators (Rogers, 2003, p. 42), and mavens (Gladwell, 2000). If these individuals within the practitioner community can be identified and informed, they can play an important role in subsequent diffusion.

The key to forming an informing relationship in which an innovator-expert becomes the client of an academic-expert is establishing motivation. One such motivation may be the fact that practitioner expertise often becomes unnecessarily specialized during movement towards compiled schemata based reasoning. As previously noted (e.g., Camerer & Johnson, 1991, p. 195; Gladwell, 2005, p. 125-136), expert performance can often improve through a systematic analysis of task cases aimed at making the expert's mental model more parsimonious. Similar improvements over expert accuracy are also frequently reported in expert systems performance (e.g., Gill, 1996). Practitioner-experts have room for improvement just like the rest of us—even in tasks that they routinely perform. The challenge is in identifying what knowledge will be most valuable to the expert and, of course, achieving effective informing.

When practitioner-experts communicate with academic-experts, the same general principles apply. The academic-expert client who perceives that the problem space has a rugged goal space is unlikely to have a hard and fast model that would be disrupted by the practitioner-expert's perceptions of the task. Indeed, the practitioner-expert's insights could well lead to insights on schemata that might prove applicable to other peaks, recalling that accumulating knowledge of task cases from a rugged goal space tends to expand, rather than reduce, the attributes considered in the goal fitness operator.

Perhaps the best approach to avoiding academic-practitioner communications challenges is to nurture the development of a hybrid expert, the academic-practitioner-expert (the academic who practices on the side or who had a long history of practice prior to acquiring academic training) or the practitioner-academic-expert (the practitioner who has completed most or all the training sequence that also produces academic researchers). In some fields, such individuals are widely distributed. In the sciences, for example, non-academic readership of research journals is common:

...there are many scientists who read but never write, in particular in non-university settings such as industry, government, and national laboratories. This is especially true in the medical field. Virtually every physician—even those not involved in research or teaching—must read medical journals to keep up with advances in patient care. In the life and physical sciences about 70% of all readings are done by nonacademicians, who write considerably less frequently than academic scientists. (Tenopir & King, 2001)

Similarly, a terminal degree in the sciences—as well as some social sciences, most notably psychology and economics—can often prove advantageous when looking for professional, industry, or government positions. In many professional fields—including medicine, law, engineering, and education—the same can be said.

The one major exception to this pattern of nurturing hybrid experts is business where, with the exception of finance (Pfeffer, 2007, p. 1335), there is little or no evidence that academic research credentials are an effective path to advancement in practice in the U.S.—in Germany, on the other hand, such credentials are more highly prized (Mintzberg, 2004)—or that considerable practical experience will facilitate entry into the research community. Thus, we would expect the academic-practitioner gap to be particularly pronounced in the business domain.

## Conclusions

Although the structural complexity model proposed here is applicable to virtually any informing activity between a human sender and a human client, the present paper has chosen to focus particularly on complex informing situations. There are two reasons for this. First, such informing is particularly prevalent in academic settings. Since the author is employed in such a setting—as is likely to be the case for most readers—it is an area of personal relevance. Second, there is probably no other domain where the transformation of client knowledge structures, both in competition with existing knowledge structures and where none previously existed, plays such an active role.

From our analysis of the structural complexity in the problem space, there are a number of useful general conclusions we can draw about the nature of problem spaces and informing:

1. One of key factors impacting informing is the ruggedness of the goal space presented by the task domain and internally represented by both sender and client. Rugged goal spaces tend to promote problem spaces where task case knowledge dominates core knowledge; decomposable goal spaces, in contrast, are accompanied by proportionately larger bodies of core knowledge that tends to be highly generalizable.
2. Academic-experts and practitioner-experts tend to have knowledge distributed in very different ways within their respective problem spaces. Academic-experts favor episteme, or conceptual knowledge, that tends to be generalizable and relatively easy to communicate symbolically. Practitioner-experts tend to favor phronesis, or practical wisdom, that is often situation specific and difficult to articulate.
3. There are a variety of generic strategies that may be used in informing. The top down informing approach, relying heavily on communicating conceptual knowledge, is a natural fit with academic-experts, but tends to be ill-suited to task domains with a rugged goal



space. The bottom up approach focuses on inducing concepts from very specific examples; it tends to be a good fit with practitioner-expert knowledge but presents significant communications barriers based upon the compiled nature of the knowledge to be conveyed, particularly where novice clients are involved. The vertical slice approach involves studying individual task cases at all knowledge levels; it can be well applied by either academic-experts or practitioner-experts but may not lead to consistent theory when applied to an overall task with a rugged goal space.

4. In the rugged goal space, the barriers to informing between experts-in-practice and experts-in-theory are formidable; if such informing is to take place at all, then meeting the prerequisites of effective informing will generally have to take precedence over the actual substance of the information being transmitted or exchanged.

## References

- Bain, K. (2004). *What the best college teachers do*. Cambridge, MA: Harvard University Press.
- Bandura, A. (1991). Social cognitive theory of self-regulation. *Organizational Behavior and Human Decision Processes*, 50, 248-287.
- Camerer, C. F., & Johnson, E. J. (1991). The process-performance paradox in expert judgment: How can experts know so much and predict so badly? In K. A. Ericsson & J. Smith (Eds.), *Towards a general theory of expertise: Prospects and limits* (pp. 195-217). Cambridge, UK: Cambridge University Press.
- Campbell, D. J. (1988). Task complexity: A review and analysis. *Academy of Management Review*, 13(1), 40-52.
- Card, S. K., Moran, T. P., & Newell, A. (1983). *The psychology of human-computer interaction*. Hillsdale NJ: Earlbaum.
- Charness, N. (1991). Knowledge and search in chess. In K. A. Ericsson & J. Smith (Eds.), *Towards a general theory of expertise: Prospects and limits*. Cambridge, U.K.: Cambridge University Press.
- Cohen, E. (1999). Reconceptualizing information systems as a field of the transdiscipline informing science: From ugly duckling to swan. *Journal of Computing and Information Technology*, 7(3), 213-219.
- Davis, M. (1971). That's interesting!: Towards a phenomenology of sociology and a sociology of phenomenology. *Philosophy of the Social Sciences*, 1, 309-344.
- Dijksterhuis, A., Bos, M., Nordgren, L., & van Baaren, R. (2006). On making the right choice: The deliberation-without-attention effect. *Science*, 311, 1005-1007.
- Elliot, A. J., & Harackiewicz, J. M. (1996). Approach and avoidance achievement goals and intrinsic motivation: A mediational analysis. *Journal of Personality and Social Psychology*, 70(3), 461-475.
- Estrada, R. (2008). *Common sense test taking*. Retrieved on 4/23/08 from <http://www.teachervision.fen.com/teacher-training/educational-testing/25091.html>
- Gilbert, D. (2007). *Stumbling on happiness*. New York: Knopf.
- Gill, T. G. (1991). *Expert systems: A mapping between symbol susceptible tasks and software tools*. Unpublished doctoral dissertation. Boston, MA: Harvard Business School.
- Gill, T. G. (1995). High-tech hidebound: Case studies of information technologies that inhibited organizational learning. *Accounting, Management & Information Technologies*, 5(1), 41-60.
- Gill, T. G. (1996). Expert systems: Task change and intrinsic motivation. *MIS Quarterly*, 20(3), 301-329.
- Gill, T. G. (2008a). A psychologically plausible goal-based utility function. *Informing Science: the International Journal of an Emerging Transdiscipline*, 11, 227-252. Retrieved from <http://inform.nu/Articles/Vol11/ISJv11p227-252Gill220.pdf>

## Structural Complexity

- Gill, T.G. (2008b). Reflections on researching the rugged fitness landscape. *Informing Science: the International Journal of an Emerging Transdiscipline*, 11, 165-196. Retrieved from <http://inform.nu/Articles/Vol11/ISJv11p165-196Gill219.pdf>
- Gill, T. G., & Hicks, R. (2006). Task complexity and informing science: A synthesis. *Informing Science: the International Journal of an Emerging Transdiscipline*, 9, 1-30. Retrieved from <http://inform.nu/Articles/Vol9/v9p001-030Gill46.pdf>
- Gill, T. G., & Holton, C. (2006). A self-paced introductory programming course. *Journal of Information Technology Education*, 5, 95-105. Retrieved from <http://jite.org/documents/Vol5/v5p095-105Gill114.pdf>
- Gill, T. G., & Sincich, T. (2008). Illusions of significance in a rugged landscape. *Informing Science: the International Journal of an Emerging Transdiscipline*, 11, 197-226. Retrieved from <http://inform.nu/Articles/Vol11/ISJv11p197-226GillIllusions.pdf>
- Gladwell, M. (2000). *The tipping point*. New York: Back Bay Books.
- Gladwell, M. (2005). *Blink*. New York: Back Bay Books.
- Hackman, J. R. (1969). Toward understanding the role of tasks in behavioral research. *Acta Psychologica*, 31, 97-128.
- Halhoun, I. A., & Hestenes, D. (1985). The initial knowledge state of college physics. *American Journal of Physics*, 53, 1043-1055.
- Heath, C., & Heath, D. (2007). *Made to stick: Why some ideas survive and others die*. New York, NY: Random House.
- Hinds, P. J. (1999). The curse of expertise: The effects of expertise and debiasing methods on predictions of novice performance. *Journal of Experimental Psychology: Applied*, 5(2), 205-221.
- Kauffman, S. A. (1993). *Origins of order: Self organization and selection in evolution*. Oxford, U.K.: Oxford University Press.
- Kessels, J. P. A., & Korthagen, F. A. J. (1996). The relationship between theory and practice: Back to classics. *Educational Researcher*, 25(3), 17-22.
- Kosslyn, S. M. (1980). *Image and mind*. Cambridge, MA: Harvard University Press.
- Kuhn, T. S. (1970). *The structure of scientific revolutions* (2nd ed.). Chicago, IL: University of Chicago.
- Larkin, J. H. (1981). Enriching formal knowledge: A model for learning to solve textbook physics problems. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 311-334). Hillsdale, NJ: Lawrence Erlbaum.
- Locke, E. (1978). The ubiquity of the technique of goal setting in theories of and approaches to employee motivation. *Academy of Management Review*, 3, 594-601.
- McCabe, T. J. (1976). A complexity measure. *IEEE Transactions on Software Engineering*, 2(4), 308-320.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, 88(5), 375-407.
- Miller, G. A. (1967). *The psychology of communication*. New York, NY: Basic Books.
- Mintzberg, H. (2004). *Managers, not MBAs*. San Francisco, CA: Berrett Koehler.
- Neves, D. M., & Anderson, J. R. (1981). Knowledge compilation: Mechanisms for the automatization of cognitive skills. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 57-84). Hillsdale, NJ: Lawrence Erlbaum.
- Newell, A. (1991). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Oswald, M. E., & Grosjean, S. (2004). Confirmation bias. In R. F. Pohl (Ed.), *Cognitive illusions* (pp. 79-96). Hove, U.K.: Psychology Press.

- Payne, J., Bettman, J., & Johnson, E. (1993). *The adaptive decision maker*. Cambridge, UK: Cambridge University Press.
- Pfeffer, J. (2007). A modest proposal: How we might change the process and product of managerial research. *Academy of Management Journal*, 50(6), 1334-1045.
- Reber, R. (2004). Availability. In R. F. Pohl (Ed.), *Cognitive illusions* (pp. 147-164). Hove, U.K.: Psychology Press.
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). New York: Free Press.
- Ross, L., Greene, D., & House, P. (1977). The false consensus effect: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, 13, 279-301.
- Rumelhart, D. E., & Norman, D. A. (1981). Analogical processes in learning. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 335-359). Hillsdale, NJ: Lawrence Earlbaum.
- Shiffrin, R. M., & Dumais, S. T. (1981). The development of automatism. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 111-140). Hillsdale, NJ: Lawrence Earlbaum.
- Teigen, K. H. (2004). Judgment by representativeness. In R. F. Pohl (Ed.), *Cognitive illusions* (pp. 165-182). Hove, U.K.: Psychology Press.
- Tenopir, C., & King, D. W. (2001). The use and value of scientific journals. *Serials*, 14(2), 113-120.
- Waterman, D. A. (1986). *A guide to expert systems*. Reading, MA: Addison Wesley.
- Wood, R. (1986). Task complexity: Definition of the construct. *Organizational Behavior and Human Decision Processes*, 37, 60-82.
- Woods, W. A. (1975). What's in a link: Foundations for semantic networks In D. G. Bobrow & A. M. Collins (Eds.), *Representation and understanding: Studies in cognitive science* (pp. 35-82). New York, NY: Academic Press.

## Biography



**Grandon Gill** is an Associate Professor in the Information Systems and Decision Sciences department at the University of South Florida. He holds a doctorate in Management Information Systems from Harvard Business School, where he also received his M.B.A. His principal research areas are the impacts of complexity on decision-making and IS education, and he has published many articles describing how technologies and innovative pedagogies can be combined to increase the effectiveness of teaching across a broad range of IS topics. Currently, he is an Editor of the *Journal of IT Education* and an Associate Editor for the *Decision Science Journal of Innovative Education*.