7/9/01

Thoughts on Adaptive Expertise

John Bransford

At the risk of being a “Johnny come lately”, let me try out some thoughts about the recent discussion that Tom Harris started about the concept of adaptive expertise.

1. Why should be care about the concept of adaptive expertise?

We’ve talked a lot about the idea of “working backwards’ as we think about courses and programs (e.g., Wiggins & McTighe’s book on “Understanding by Design, 1997, which is highly compatible with HPL; Sean’s excellent description of this process).

There are at least two important sources of information that can guide the working backwards process. One is knowledge of our particular domains (e.g. Biomechanics, Optics) and what we want students to know and be able to do that is relevant to that domain. Another is knowledge of the world “out there” -- after graduation. The concept of adaptive expertise becomes especially important when we explore the world “out there”.

We have some great interviews with people like Peter Vaill on “whitewater worlds” and the kinds of people it takes to navigate them. Another interview is with Tony White (CEO of Celera Genomics) who talks about the need to create environments were “his” people use more of their brains (than is typically the case) to invent and innovate. Larry Howard (not yet on tape but we hope to capture him soon) has noted how he has dealt with many experts in engineering and technology who are used to “applying technical algorithms” but not used to being adventuresome in their thinking. All of you probably know lots of good examples of contrasting cases of “by the book” people versus those who are willing ---AND EQUIPPED--to jump in and try new things. [We should get more descriptions of these kinds of cases in interview form.]

The ability to change and continually innovate is where the concept of equipping students to be adaptive experts comes into play. But as Tom Harris emphasizes, we need to be much clearer about the meaning of this construct. I’ll try to explain below why I think that a focus on adaptive expertise will involve a focus on what we have defined as “core competencies” PLUS. The “PLUS” is the extra needed for adaptive rather than simply routine expertise. I’ll also try to sketch some of the cognitive processes that seem to underlie adaptive expertise. I think it is crucial to try to understand and study these.

2. Research on Adaptive Expertise: Hatano and other researchers have differentiated between “routine” and “adaptive” experts. As an illustration, HPL (p 45 of the expended edition) discusses Miller’s studies of information systems designers who work with clients to design computer systems that allow them to efficiently store and access relevant information . Routine experts (“artisans”) try to identify the functions that their clients want automated. They tend to accept the problem and its limits as stated by the clients. Their approach to these tasks is primarily to find things that they have done before that can be applied to the new situation. They attempt to “get the problem solved’ as efficiently as possible and then move on to the next task.

In contrast to the artisans (routine experts), the adaptive experts (virtuosos) listen to the clients’ statement of the problem but only as a point of departure for further discussions. They realize that how one defines one’s problem is half the battle, and they know that clients are often too close to a situation to see their problems from a variety of different perspectives.

Adaptive information designers also look forward to the opportunity to expand their thinking and increase their existing solution strategies. They treat their clients’ problems as opportunities for new learning rather than simply as a job to do. (At the risk of making him blush, I think that Larry Howard is a fantastic model of an adaptive expert in the computer systems world. To emulate him, we need to get all our students to drink really strong coffee at least 6 times a day.)

In general:

- **Routine experts** have learned a set of routines that can be very complex and sophisticated, and the experts become very skilled at applying them. These routines can involve expert communication skills, design analysis skills , data gathering and analysis skills, and so forth. Many of the competencies we have been discussing can easily fit “routine expertise” . This doesn’t mean they aren’t important. But it does mean that --in order to define adaptive expertise--we need to explore the issue of “routine competencies” PLUS.....

Routine experts continue to learn throughout their lifetimes, but the learning tends to be one of becoming increasingly efficient at doing what they have also been doing, and perhaps of adding a few new tricks along the way. Studies of cigar rollers in a cigar company showed that they kept getting faster and faster over time.

Being a routine expert is great if one’s world stays stable. About 40 years ago, adolescents could learn how to fix cars from their Dad or Mom and turn this into a lifelong area of employment. Today, good car mechanics have to undergo rigorous training about every 6 months because there is so much change.

As humans, we all need to have routines that we can count on. As William James said, “Habit is the flywheel of society”. I don’t want to be an adaptive expert with respect to typing. I just want a good keyboard and the ability to increase my efficiency over time. On the other hand, some people with Carpal Tunnel syndrome have faced the need to learn a whole new, more ergonomic keyboard--and have done so with excellent long-term results. Others have failed to change and hence remain in pain.

_ **Adaptive Experts:** Compared to routine experts, adaptive experts are more likely to relish challenges that require them to “stretch” their knowledge and abilities. They tolerate ambiguity, at least for a while, and they think of themselves as people who know a lot, yet still know little compared to all that is knowable. They are particularly aware of the “assumptive nature of knowing” (e.g., how their current beliefs and knowledge affect their “fish is fish” constructions), and they are able to “let go” of these assumptions without feeling overly threatened. They also actively try to make their tacit assumptions explicit and test them against various criteria. As the Philosopher of Science Toulmin put it:

> A person demonstrates his (or her) rationality, not by a commitment to a fixed set, stereotyped procedures, or immutable concepts, but by the manner in which and the occasions on which, he (she) changes those ideas and procedures. (p. v).

The ability to be an adaptive expert requires that people deal with emotions (hot cognition) as well as skills and knowledge. The physicist David Boehm points out the emotional turmoil that is often involved in changing one’s thinking. His description refers to a scientist (in this case a male scientist) being confronted by conflicting opinions:

> His first reaction is often of violent disturbance, as views that are very dear are questioned or thrown to the ground. Nevertheless, if he will “stay with it” rather than escape into anger and unjustified rejection of contrary ideas, he will discover that this disturbance is very beneficial. For now he becomes aware of the assumptive character of a
> the conscious criticism of one’s own metaphysics, leading to changes where appropriate and ultimately, to the continual creation of new and different kinds.

We probably all know people who set up barriers to protect their comfort zones and try to avoid upsetting their cherished procedures and beliefs. Others are willing to explore new possibilities despite the initial turmoil involved. But it’s important to note that even adaptive experts have to pick their battles carefully. There’s not enough time in life to continually rethink everything we believe.

3. Developing Adaptive Experts:

I think the core conjecture for us to explore is that the development of adaptive expertise is not something that simply happens AFTER people develop routine expertise. You don’t develop it in a “capstone course” at the end of students’ senior year. Instead, the path toward adaptive expertise is probably different from the path toward routine expertise.

Adaptive expertise involves habits of mind, attitudes, and ways of thinking and organizing one’s knowledge that are different from routine expertise and that take time to develop. I don’t mean to imply that “you can’t teach an old routine expert new tricks”. But it’s probably harder to do this than to start people down an “adaptive expertise” path to begin with--at least for most people.

- **Helping students learn about themselves as thinkers and problem solvers**

I think that a key to developing adaptive experts is to help students understand their own processes of knowing and problem solving, plus help them develop an identity as a lifelong learner rather than as an expert who is supposed to know all the answers. There are issues here that are much deeper than the standard cliches (e.g., “being a lifelong learner is good”)

Helping students develop adaptive expertise requires a **metacognitive** approach to teaching. Students need to understand how they think, and how what they currently know can be both a blessing and a curse. They need to understand that their (usually tacit) ideasof “what it means to be a competent professional” (e.g. knowing most of the answers versus really being a learner) has major effects on how they think and act, and how comfortable they feel about taking a chance. They need to pay attention to the processes they use to solve problems. [There are some powerful examples in HPL that we should pull out and highlight of ways that metacognitive additions to instruction have increased achievement in courses in physics and other areas.]

One thing our students need to understand is that their problem solving is always affected by their current knowledge and assumptions (just like Fish is Fish). Having prior knowledge and beliefs is both necessary and nice -- we can’t function in a mental vacuum. But our current assumptions can also trap us in a box that confines us to a “problem space” that is much narrower than it should be.

Consider the following problem:

Two men played 5 games of checkers. Each won three games. How is this possible?

Most people who try to solve this problem first assume (reasonably) that the two men are playing each other. Given that problem space, the problem is impossible to solve. After a few seconds of confusion, most people typically let go of the assumption that the two men were playing each other. This helps create a larger problem space where the answer becomes trivial. If the two men are not playing one another, it is easy for each to win three games.

The preceding is one of those trick problems, of course. But nature is also full of trick problems (or maybe it’s more accurate to say that we inadvertently play tricks on ourselves when we attempt to solve problems). The trick we play on ourselves is that we tend to jump into solutions based on tacit, restricted definitions of a problem--hence we operate in a restricted problem space. This is similar to the routine information system designer (see above) who simply takes the clients’ definition of the problem as a given and doesn’t try to reframe the problem to reveal alternative problem spaces that can be explored.

One of my favorite examples of restricted problem spaces comes from a book by Adams called Conceptual Blockbusting. He discussed the bruised tomato problem that a group of engineers tried to solve about 40 years ago. The problem as stated by the client was : “we can’t afford to have humans pick tomatoes--we need automated tomato pickers. But the current ones are bruising the tomatoes. We need you engineers to design us a tomato picker that is less likely to bruise tomatoes".

The engineers worked for about 6 months and made mild headway but no major improvements. They padded the picking arms, slowed down the machine a little, etc. But the improvements were pedestrian.

Eventually, some biologists were brought into the picture. They redefined the problem and hence opened up a how new problem space that contained new solutions. Instead of trying to design a automated tomato picker that was less likely to bruise tomatoes, they set out to design a tomato that was less likely to be bruised. And they succeeded (their new tomato had thicker skin and grew further out on the vine).

Interestingly if this bruised tomato problem were posed in this day and age, traditional biologists and bioengineers would probably have different views of how to frame the "reinvent the tomato" problem. One would look at selected breeding, the other at genetic engineering. A major reason for getting collaborative, "distributed expertise" teams involved in problem solving is to generate alternative problem definitions and increase the problem spaces that can be explored.

4. Should we think about a systematic framework for thinking about and doing problem solving?

It might be useful for us to develop a systematic approach to problem solving that helps our students think more carefully about their own assumptions and processes as they solve problems. One possible framework comes from The Ideal Problem Solver (1993) that my colleague Barry Stein and I wrote. I have NO investment in getting us to use this framework versus some other--I use it simply as an illustration. For those who have read the book (there are a total of 5 in the world, I think), you'll notice that I've modified the IDEAL acronym somewhat to better fit our purposes. Here's a quick overview of what IDEAL means.

I = Invest the time to treat problems as opportunities for new learning. We would need to help students understand the many ways that this is important. For example, previously solved problems can serve as example cases that provide the bases for analogous solutions to new problems. Concrete examples become even more valuable if seen as instances of more general issues (see HPL). For example, if a previously solved problem can be framed and represented in memory at a more abstract level, it is more likely to be accessed when needed later rather than remain inert. So it is valuable to invest the time to relate particular problems to more general issues. When students are helped to focus on the PROCESSES of solving problems, they can also learn more that helps transfer (e.g. see the physics learning studies by Mestre et al. in Ch 6 of HPL).

D = Develop an understanding of the problem you are trying to solve and the processes you are using to solve it. This means that students must be helped to define **learning goals** for obtaining more information rather than simply jumping to solutions based on superficial understandings of problems. And they must understand WHY this is important. **This has implications for how we design challenges.**

Consider a simple challenge we just built for Emergency Medicine instruction. A college student has been brought to the emergency room because “she won’t wake up”. If novices are asked,“what do you think might be wrong with her” they’ll generate lots of possibilities based on everyday experiences like “too much to drink”, “drug overdose”, etc. But this is essentially “guessing” based on existing knowledge. To develop adaptive experts, we need to help students do more than this.

The Emergency Medicine Challenge we built asks three questions:

(1) What might be wrong with this patient?

(2) What additional information would you want in order to feel confident making a diagnosis and deciding on a treatment--and why would you want it?

(3) What are some things that you would want to learn in order to develop more expertise in this area?

Students revisit this challenge as they proceed with their course.

Questions (2) and (3) don’t have to be in every challenge, but they are particularly important to include in some challenges in order to develop adaptive expertise. Question (2) gets at the idea that problems don’t typically come prepackaged with all the data for solving them--we need to know what to seek in order to be confident and competent. Question (3) gets students in the habit of defining learning goals. For example, as I thought about Q 3 I realized that I don’t really know what it means to be unconscious, and I wondered if there are a number of different unconscious states (e.g., being knocked unconscious, being unconscious due to a drug overdose, etc.). I want to learn about unconsciousness and it’s relationship to sleep, fainting, hypnosis, etc. I also want to know if and how different reasons for unconsciousness can be detected by different sets of vital signs, etc, and whether different sets of signs suggest different methods for “awakening” patients (e.g., smelling salts?, Letting patients awake on their own?, etc.). My guess is that generating these kinds of questions first will create a “time for telling” (Schwartz & Bransford) that helps students better learn new information that is relevant).

Similar to the Emergency medical challenge (above), consider the“Port wine stain” problem. If posed simply as a problem to be solved, many people might simply start to generate solution strategies like “try to replace the skin with a skin graft”, etc. We need to help students develop the habit of mind of thinking thoughts like “I really need to understand what is causing this. Does it just involve the skin, or are other things like blood vessels involved”, etc?). The tendency to jump right to possible strategies without trying to first understand and conceptualize problems is a tendency found in all fields. And it is a dangerous tendency because it builds habits of overassimilation to existing knowledge and beliefs.

Kay Burgess and Xiaodong Lin on the LTC have documented the effects of an over-reliance on assimilation (cf Bransford & Schwartz, 1999) They created catalogs of items that one might purchase for an Eagle recovery plan. Some of the items were bogus and would not work. Fifth and sixth graders were given the catalogs plus resources they could use to help make their catalog selections. Most were over-confident in their competence for this task. They simply used their everyday intuitions as human beings and made decisions (which tended to be wrong) without consulting any resources. For example, they chose the option of hand-raising baby eagles because the babies would "feel like orphans" if they put them in an isolated incubator. (The problem with hand-raising eagles is that the babies would imprint on humans). In contrast, a science expert who had no knowledge of eagle recovery took a totally different approach. She knew that she needed more information and used the contrasting catalog items to formulate learning goals that guided her search through the resources that were available for the study. She exhibited adaptive expertise (Hatano & Ingaki, 1986).

Wineberg (in press) provides an additional example of the importance of overcoming the tendency to simply assimilate. He studied a historian who was asked to analyze a set of history documents that focused on a topic that was outside his area of specialization. At first, the historian resolved puzzling contradictions in the documents by using his existing knowledge of present day culture. Eventually he came to the conclusion that he did not have enough historical knowledge about the situation to make an informed judgement, so he devised learning goals and carried them out. After opportunities to learn, the historian did as well at analyzing the history documents as an expert who specialized in that area. In contrast, college students presented with the same documents tended to use their intuitive everyday knowledge and generated erroneous conclusions (Wineburg & Fournier, 1994). They failed to question their existing assumptions and, ultimately, failed to take advantage of new opportunities to learn.

Overall, we need to create challenges that give our students lots of opportunities to ask for more information and define learning goals. Often, the problems that we use develop habits of mind like those of the routine expert in systems design (see above) who always accepts the client's problem statement as a given rather than simply as a point of departure for further exploration. In medical education, Barrows (the father of "problem based learning") does a great job of consistently emphasizing the formulation of learning goals as students proceed through medical school.

E = Explore at least two different ways to frame (define) the problem you are working with. This is a very simple but powerful strategy that moves one's thinking one level of abstraction about most attempts to solve problems. For example, a bird flew into the window of a seminar I was holding and started to panic and dive at students--they were scared. I tried to catch the bird and then remembered to ask myself (What's my tacit problem of the current problem?). It was: How do I catch the bird before it hurts someone or itself? The problem space for this definition of the problem has lots of obvious strategies, but none of them worked.

It then hit me to explore a different definition of the problem; namely, “How can I get the bird to voluntarily leave the room”. For this, I turned out all the lights and hoped it would head back out the lighted window. It did. This solution required knowledge (e.g. Of phototropism), and it was lucky that the room could get dark and there was light outside. But I would never have tried it if I had not consciously used the IDEAL strategy of exploring **at least two** ways to define the problem to be solved.

Interestingly, just last year I encountered a similar problem at a conference session in Virginia---a bird flew in the room and started darting at people and also flying into the screen--which was lit. Since I had seen this problem before, it was a snap to know what to do (this is one of the advantages of “case based” representations of previously solved problems --provided the cases are coded so people can retrieve them when needed). People were amazed at my “quick thinking”. Actually, it wasn’t quick thinking at all. I simply accessed a familiar case that I had used previously. I suspect that most “quick” problem solving involves access to well known cases (or more general schemas) like this. Once again, however, these cases need to be represented at the appropriate level of abstraction or they will remain inert when the problem arises again (e.g., see Bransford et. al’s “Because Wisdom Can’t be Told”).

A = ACT on the strategy(s) that seems most promising --initially through thought experiments or pilot studies and eventually in real world contexts.

L = Look at your whole process of problem solving and **learn** from the experience so you can improve next time around. [This includes **Investing** the time to explore one’s solution strategies and consider possible alternatives, and to try to see how the specific problem one has solved relates to more general issues in a domain (e.g., I might be exploring a cave and find life that does not need oxygen and be helped to realize that this is highly relevant to astrobiologists who are trying to find other places that support life.)

- **Beyond Memorizing IDEAL:** Obviously, It doesn’t do much good to simply memorize the IDEAL framework--it needs to be used to structure students activities of problem representation, problem solving and reflection. There are probably other frameworks that are even better than IDEAL (The IDEALer Framework?) But I thought I’d throw IDEAL out as an example of how we might take a systematic, metacognitive approach. .

- **Mini Studies:**I see the potential for some very powerful mini studies we can and should do to show NSF that there are really “teeth” in pursuing the goal of developing adaptive expertise. We should also look at research on Problem Based learning in Medicine by Cindy Hmelo (she got her Ph D at Peabody) as a guidelines for these studies. She has strong results).

5. An “Adaptive Expertise” Challenge for All of Us

Imagine that we are part of the show "To Tell the Truth" and are introduced to two seniors in Bioengineering who each claim to be adaptive experts. But in reality, one is a routine expert. What would we do in order to decide which is which? If we simply ask "Are you an adaptive expert?", both will say yes (or perhaps "huh?")

Both should show evidence of basic competencies like core knowledge of Bioengineering, the ability to solve familiar (routine) problems that are similar to ones solved before, the ability to communicate, etc. What's the "PLUS" for the adaptive expert?

Here are some sacrificial thoughts about the two most important things:

- A systematic understanding of themselves as learners and problem solvers -- including their strengths and weaknesses, plus their understanding of the importance of building "distributed expertise" teams.

- An ability to LEARN to solve novel problems (not simply solve routine problems based on existing knowledge). This includes the ability to generate learning goals, define multiple perspectives on a problems, deal with criticism and contradictory data, etc.

Assessing the latter is relevant to new ways of thinking about the issue of learning and transfer (Bransford & Schwartz, 1999). A striking feature of most transfer studies is that they all use a final transfer task that involves what Bransford & Schwartz (1999) call "sequestered problem solving" (SPS). Just as juries are often sequestered in order to protect them from possible exposure to "contaminating" information, subjects in experiments are sequestered during tests of transfer. There are no opportunities for them to demonstrate their abilities to learn to solve new problems by seeking help from other resources such as texts or colleagues, or by trying things out, receiving feedback and getting opportunities to revise. Accompanying the SPS paradigm is a theory that characterizes transfer as the ability to directly apply one's previous learning to a new setting or problem (we call this the Direct Application (DA) theory of transfer). Bransford & Schwartz's thesis is that the SPS methodology and the accompanying DA theory of transfer is responsible for much of the pessimism about evidence for transfer.

An alternative to SPS methodology and DA theory is a view that acknowledges the validity of these perspectives but also broadens the conception of transfer by including an emphasis on peoples' "preparation for future learning" (PFL). Here, the focus shifts to assessments of people's abilities to learn in knowledge-rich environments. When organizations hire new employees they don't expect them to have learned everything they need for successful adaptation. They want people who can learn, and they expect them to make use of resources (e.g., texts, computer programs, colleagues) to facilitate this learning. The better prepared they are for future learning, the greater the transfer (in terms of speed and/or quality of new learning).

Examples of how a PFL View of Transfer helps us Rethink the Quality of Various Learning Experiences.

The PFL perspective helps us notice evidence of positive transfer that is often invisible in the traditional SPS paradigm. This, in turn, helps reveal the value of various learning experiences that often remain hidden when assessed by traditional methods.. Prevailing theories and methods of measuring transfer work well for studying full-blown expertise, but they represent too blunt an instrument for studying the smaller changes in learning that lead to the development of expertise. New theories and measures of transfer are required.

As a simple illustration of a PFL perspective on transfer, consider a set of studies conducted by Kay Burgess & Sean Brophy . In one study they asked fifth graders and college students to create a state-wide recovery plan to protect Bald Eagles from the threat of extinction. The goal was to investigate the degree to which their general educational experiences prepared them for this novel task; none of the students had explicitly studied Eagle recovery plans.

The plans generated by both groups missed the mark widely. The college students' writing and spelling skills were better than the fifth graders, but none of the college students mentioned the need to worry about baby eagles imprinting on the humans who fed them, about creating tall hacking towers so that fledgling eagles would imprint on the territory that they would eventually call home, and about a host of other important variables. In short, none of the students--college or fifth graders-- generated a recovery plan that was even close to being adequate. Based on these findings, one might claim that the students' general educational experiences did not prepare them adequately for transfer.

However, by another measure of transfer, the differences between the age groups were striking. We asked the students to generate questions about important issues they would research in order to design effective recovery plans for eagles . The fifth graders tended to focus on features of individual eagles like how much do they eat? How big are they? Where do they live?. In contrast, the college students were much more likely to focus on issues of interdependence between the eagles and their habitats. They asked questions such as "What type of ecosystem supports Eagles," (reflecting an appreciation of interdependence); "What about predators of Eagles and Eagle babies?" (also reflecting interdependence); "Are today's threats like the initial threats to eagles?" (reflecting an appreciation of history and change); "What different kinds of specialists are needed for different recovery areas?" (reflecting an appreciation for a possible need for multiple solutions). Because they had not studied eagles directly, the college students were presumably generating questions that were framed by other aspects of biology that they had learned. So, by this alternative form of transfer test, it would appear that the college students had learned general considerations that would presumably help shape their future learning if they chose to pursue this topic (Scardamalia & Bereiter, 1992). In this regard, one would call their prior learning experiences a success.

Overall, the traditional SPS approach to assessment in the eagle study revealed how far the fifth grade and college students were from developing an adequate Eagle recovery plan, and it invited the inference that the students' K-12 experiences had not prepared them for this kind of transfer. From the PFL perspective, one looks for evidence of initial learning trajectories. So, rather than evaluate whether people can generate a finished product, the focus shifts to whether they are prepared to learn to solve new problems. For example, one determinant of the course of future learning is the questions people ask about a topic, because these questions shape their learning goals (e.g., see Barrows,1985; Bereiter & Scardamalia, 1989; Hmelo, 1994). For the eagle experiment, the PFL perspective yielded a deeper appreciation of how the college students' K-12 experiences had prepared them to learn.

Clarifying the value of various teaching strategies: The PFL perspective can also help clarify the advantages of various teaching techniques (e.g., challenge-based inquiry) that look inefficient from other perspectives. For example, consider efforts to compare the benefits of (a) beginning lessons by first having students generate their own, perhaps incorrect, thoughts about phenomena versus (b) simply telling students the correct answers. Examples might include attempts to have students begin an instructional sequence by first generating their own experiment to test some idea (Bransford et al., 1990) or creating their own formula for capturing the variance of statistical distribution, (Schwartz & Moore, in press). Since novices will often generate ideas that are incorrect, they must eventually be guided toward more fruitful ways of thinking. Why not "cut to the chase" and present the correct ideas right from the start?

The PFL perspective suggests a number of reasons for first having students generate their own ideas about phenomena. The most important is that it provides an opportunity for students to contrast their own thinking with that of others, including experts in an area. This sets the stage for appreciating the critical features of the new information that is presented to them--analogous to the perceptual examples from Garner (1974) discussed earlier. For example, students who first generated their own thoughts on how to design an experiment to test a particular idea expressed appreciation about the elegance of the experiments discussed in an article that was then assigned to them (Bransford et al., 1987). In contrast, students who were simply assigned the article did not have the advantage of experiencing how the article helped clarify their own thinking. As a result, they treated the article simply as a set of facts to be learned.

Schwartz and Moore (in press) illustrate a similar example in the domain of statistics. The idea is that students are better prepared to appreciate the formula for standard deviation if they are first given opportunities to differentiate the elements of "spread" that the formula has to account for. To help differentiate these elements, students are shown an initial pair of distributions, say {2, 4, 6, 8, 10} and {4, 5, 6, 7, 8}. The experimenters point out that the two sets have a similarity, and they ask the students to notice that there is a single number for each set that helps determine this similarity--the average. This single number is easier to keep in mind and communicate than the total distribution.

The experimenters then ask students to come up with a method for determining a single number for each set that could capture what is different (i.e., the variance). After they invent their own methods (often a range formula) they receive a new pair of distributions, say {2, 4, 6, 8, 10} and {4, 8} and determine whether their formula works for this set as well. If it does not they should fix it. This continues for several cycles where students generate a formula and then try to apply it to new distributions that highlight new quantitative properties (like sample size). At the end of these exercises, students may be shown the formula for variance used by experts. The question of interest is: How do these exploratory activities prepare students to understand the variance formula in ways that go beyond teaching the formula from the start?

Initial results from the Schwartz and Moore studies suggest that even though the students generated faulty formulas, these experiences helped the students become aware of the quantitative properties of distributions that a formula should take into account. This set the stage for noticing critical features of experts' formulas; for example, that it yields a smaller number for smaller variances (many of the students' self-generated formulas had done the opposite); that it elegantly solves the problem of set size, and so forth. As a consequence, students in the "generate first" group were much better able to appreciate the strengths and weaknesses of different non-standard formulas for capturing variance (e.g., a formula that summed the deviations from the median instead of the mean). In contrast, students who had been directly taught the standard formula (with no previous attempts to generate their own thoughts) simply declared that the non-standard formulas were "wrong". They were not as prepared as the other students to learn about the expert formula. In Broudy's terms, they had a less well-differentiated "field" for "knowing with".

The instructional procedures noted above can also help accomplish another goal; namely, to let people experience how seemingly "intuitive" or "obvious" ideas that they initially generate can look suspect when subjected to closer scrutiny. This is important because adapting to new situations (transfer) often involves "letting go" of previously held ideas and behaviors. This is very different from assuming that transfer represents "the degree to which a behavior will be repeated in a new situation" (Detterman, 1993, p. 4). In many cases, repeating an old behavior in a new setting produces what his been called "negative transfer". Luchin's (1942) classic studies of filling water jars illustrate this point nicely. When given a transfer task, participants in these experiments repeated a complex set of water-pouring strategies despite the fact that the task permitted a simple, efficient response. Land, inventor of the Polaroid Land camera, coined a colorful definition of "insight" that highlights the importance of "letting go" of previous assumptions and strategies rather than simply repeating them. He defined insight as "the sudden cessation of stupidity" (Land, 19). It's a great definition. And it relates especially to the "Explore" part of the IDEAL framework discussed above (I.e. Exploring alternate ways to frame a problem--which can reveal assumptions that can be changed or relaxed).

6. A Final Thought

First, I hope this is helpful.

Second, it occurs to me that the Wiggins and McTighe approach to working backwards might be rephrased to read:

"Define what we want our students to <u>experience</u> during our courses and to know and be able to do at the end". They need to experience changes in their own thinking (and thinking about thinking)over time. And we need to help them learn from these experiences. I think that doing so creates a highly metacognitive environment where content knowledge (including knowledge about domain knowledge AND the nature of adaptive expertise) and processes of learning and problem solving are combined to develop adaptive expertise.

Cheers.