TOWARD TECHNICAL UNDERSTANDING

Part 3. Advanced Levels

J.M. Haile
Clemson University • Clemson, SC 29634-0909

The papers in this series* stalk the question of what we mean by an understanding of technical material. We have asserted that *to understand* has multiple meanings, and we organized those meanings into a hierarchy of seven levels: (1) Making conversation; (2) Identifying elements; (3) Recognizing patterns; (4) Solving problems; (5) Posing problems; (6) Making connections; (7) Creating extensions. In the second paper of this series, we discussed understanding at Levels 1 through 4, which we refer to as elementary understandings. To progress beyond problem solving at Level 4, we must realize that solving a problem is not the same as knowing *how* to solve it. This realization marks the beginning of the transition to the more advanced levels addressed in this paper. The discussions here rely on the descriptions of brain structure and function that were summarized in the first paper of the series.

Transition:
Level 4 (Solving Problems)
to
Level 5 (Posing Problems)

Motivation: Solving a problem is not the same as knowing *how* to solve the problem.

Reformulation: The initial solution procedure is refined by rehearsal and the problem plus its solution are explored by exercising variations on a theme.

LEVEL 5: POSING PROBLEMS

We practice problem solving not to obtain an answer, but to learn how to solve problems. That is, implementing a procedure to obtain an answer occupies a lower level of understanding than does devising the procedure. To develop skills for solving problems, we must confront new problems, solve them, and then solve them again and again. Repetition allows us to shift our attention from obtaining an answer to learning a procedure. Repetition also promotes creation of long-term memories, which we need for reusing a procedure in the future. The connections between repetition and memory will be discussed first, then we will make connections between repetition and problem solving.

Posing Problems to Create Memories

Creating memories serves as one hedge against future needs. In particular, long-term memories (certain long-lasting neural networks and combinations of networks) enable us to reuse problem-solving strategies that we have found successful in the past. At the subconscious level, we don't know how the brain selects what ideas are to be remembered. That is, in spite of popular wisdom, the brain does not lay down a memory for every mental state nor for every sensory experience; on the contrary, most pass through short-term memory and are lost. But we do know that we can consciously select what ideas are to be remembered and we can consciously create those desired memories; the operative mechanism is repetition—repeatedly thinking about the ideas.

Repetition causes the cortex to repeatedly fire the same pattern of neurons; such repeated activations appear to strengthen synaptic junctions and perhaps develop new junctions. Thus, by repeated use, a track through a wilderness becomes a path, then a walk, and finally a highway. Moreover, besides strengthening connections, repetition also seems to refine the neural network that reproduces the desired

***J.M. Haile**, a professor of chemical engineering at Clemson University, is the author of* Molecular Simulation, *published by John Wiley & Sons in 1992.*

---

* *Part 1, "Brain Structure and Function" was published in the summer 1997 issue of CEE (Vol. 31, No. 3) and Part 2, "Elementary Levels," appeared in the fall 1997 issue (Vol. 31. No. 4).*

firing patters—it makes the network more efficient. When the mind contrives a pattern for a new idea, it seems to arise "on the fly"—the new thought is hastily thrown up as a permutation of an existing pattern. If a new idea seems interesting or important, and therefore worth remembering, then repeatedly thinking about it may create a new network that is largely separated from the parent network but connected to other networks that represent related ideas.

Perhaps a helpful metaphor here would be scaffolding. A new, hastily constructed, network is a fragile thing, momentarily stabilized by a scaffolding of neural connections that allow us to examine the new idea. If the idea is judged worthy, then we use repetition to strengthen important synaptic junctions in the network and remove the scaffolding. Often, the scaffolding is produced by studying the intermediate details that appear in any logical development, such as those that connect a conclusion to a hypothesis, those that relate an effect to a cause, and those that connect an answer to a problem statement. Without intermediate details, scaffolding is sparse or nonexistent, and student understanding remain poorly developed. Repetitions make such connections at first plausible, then acceptable, and finally obvious—these correspond to stages in removing the scaffolding.

. . . we must realize that solving a problem is not the same as knowing how to solve it. This realization marks the beginning of the transition to the more advanced levels addressed in this paper.

Repetition also serves to distinguish procedural memories from episodic memories. *Procedural* memories are created by conscious practice, while *episodic* memories are apparently created from a single experience. How might episodic memories be formed? Deep in the brain, forming part of the limbic system, is the hippocampus—a pair of structures whose shapes each resemble that of a sea horse. If a hippocampus is damaged or removed, we lose the ability to form new long-term memories; old memories remain, but new ones do not form. Thus, the hippocampus plays some crucial role in forming long-term memories. Further, it communicates with the cortex through two bundles of axons, one apparently for input and another for output. This suggests that the hippocampus may act, in effect, as a buffer between short-term and long-term memories.[1]

Perhaps when many networks in the cortex are busy—attacking a hard problem—the cortex is too preoccupied to continue the structural changes that produce long-term memories. Perhaps, instead, networks in the hippocampus are activated, loading the buffer. Later, when networks become available in the cortex (perhaps during rest, or sleep), the hippocampus "replays" important patterns in the cortex, thereby creating long-term memories in the cortex.[2] If such a scenario is true, then all memories are formed by repetition; the difference between procedural and episodic memories is merely that procedural memories are created by conscious repetition, while episodic memories are created below consciousness via repetitions instigated by the hippocampus.

Posing Problems by Repetition

We invest time and effort in learning so as to realize future benefits; this implies that we intend to remember what we learn. Problem posing is the level of understanding at which we use repetition for learning how to solve problems and for creating memories of the solution procedure. We identify two kinds of repetition: *rehearsal,* in which we repeatedly pose and solve the same problem, and *variational,* in which we pose and solve new problems that are closely related to the original problem.

Rehearsal • Having solved a problem, we rehearse the procedure to learn how we solved it. Since we know that the procedure leads to the solution, our minds during rehearsal are free to consider (1) why each step is important and how it contributes to the solution, (2) whether alternative steps may be more economical, and (3) whether the steps and intermediate results can be connected to other things we know, thereby attaching additional meanings to the procedure, the solution, and the problem.

Thus, part of our activity during rehearsal is to probe and verify the logic of the algorithm; such activity conforms to Poincaré's statement that in a chain of logic, the order of the elements is more important than the elements themselves.* Another part of rehearsal is the search for a better algorithm. That is, problems are interesting and instructive to the extent that they can be solved in more than one way. Problems can themselves be viewed as patterns with their multiple meanings reflected in the various ways by which they can be solved. By repeatedly posing the same problem to ourselves, we create opportunities for finding alternative solution procedures and therefore for finding additional meanings.

A powerful motivation for rehearsal occurs when we intend to present the solution to others—perhaps as a lecture or as a written document. Such presentations are most effective when the chain of logic is economical, with every element moving the development in an obvious way toward the goal. Such presentations are developed by rehearsing, wherein we systematically try to reformulate a logical development into

---

**See Level 3 in Paper 2,* Chem. Eng. Ed., *31(4), 1997).*

a sequence that is not only economical but also rich in meaning. Minsky[3] has emphasized that reformulation is the central act of creativity. For example, in spite of the common attitude that rehearsal is merely mechanical repetition, the rehearsal involved in preparing lectures and writing textbooks provides opportunities for high levels of creativity and originality.[3]

Variational • Besides repeatedly posing the same problem, we should also pose and solve other problems that we create by systematically changing the original problem. Thus, we enhance problem-solving skills by posing variations on a theme. This activity is analogous to a practice technique used by musicians. Consider the passage from Chopin's third *Prelude* (Opus 28) for piano, shown in Figure 1. This one measure is scored as a phrase—a musical pattern of sixteen notes. The third *Prelude* is marked *vivace,* which means a lively allegro, and corresponds to a speed of about a measure per second. In fact, a measure per second would be a little slow; five measures in four seconds would be more nearly correct. Thus, each of the sixteen notes should be sounded at a uniform interval of about 5/100 of a second.

How is such skill developed? *Not* simply by repeatedly playing the measure as written, but rather by practicing rhythmic variations, such as are also shown in Figure 1. Each variation shifts the emphasis to a different note, hence a different finger; additional variations would be used to shift the emphasis among different groups of notes. The figure shows only three variations, but in practice, the musician routinely works through 40 or 50 variations of the same phrase. And those are just the rhythmic variations; one also works through variations in tempo and in dynamics (loudness). It may seem paradoxical that to achieve what the composer has written, one practices something other than what is written, but such practice proves to be an efficient way to attain absolute control over the material; to embed a metaphor within a metaphor, a chain is made stronger by systematically and repeatedly strengthening one link at a time.

Likewise, we can improve our grasp of and control over technical material by posing variations on the theme inherent in any problem. Say the original problem requires us to obtain the volume V occupied by one mole of nitrogen at P = 2 bar and T = 50°C. Having obtained the answer, we can systematically vary that problem to create somewhat different, but related, problems to solve. For example: (1) What would V be if T were 100°C instead of 50°C? (2) What would V be if P were 3 bar at 50°C instead of 2 bar? (3) What would V be if we had 5 moles instead of one at 50°C, 2 bar? (4) Can we generalize what we've learned from these four calculations? (5) What if we knew N, V, T and needed to find P? (6) What if we knew N, V, P and needed T? (7) If the gas were a binary mixture of nitrogen and oxygen, what would change in all these calculations? (8) What if the gas were a twenty-component mixture? (9) Presumably, we have used the ideal-gas law in these calculations, so by what criteria do we decide that the ideal-gas law no longer applies? (10) When the ideal-gas law doesn't apply, what should we use instead?

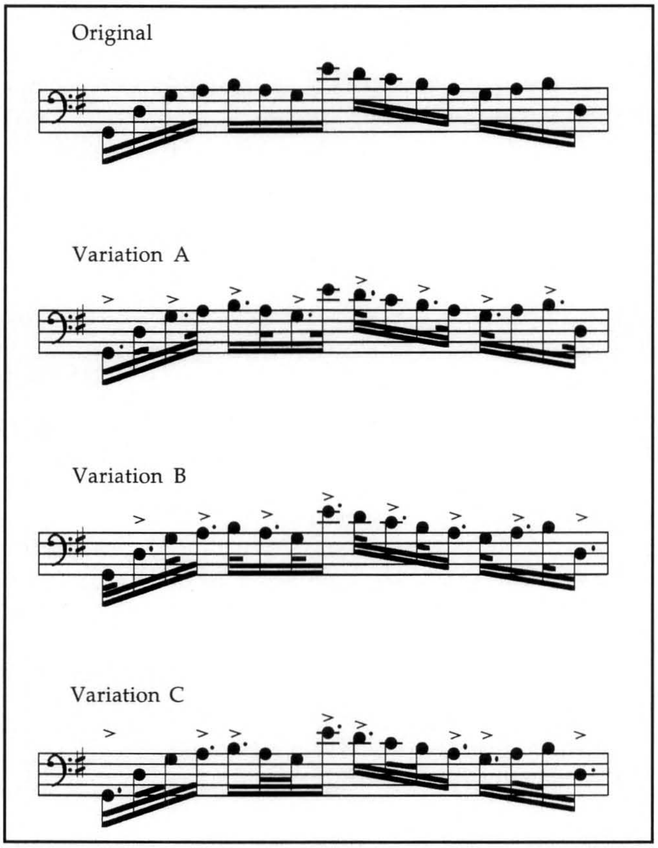

Figure 1. *Three rhythmic variations on the first measure of Chopin's* Prelude *for piano, Opus 28, No. 3. On each staff, horizontal lines and spaces between them represent keys on the keyboard; notes indicate keys to be struck. On each of the four staffs, the same keys are to be struck; thus, each staff contains the same pattern of notes. But the variations differ from the original and from one another in that they require keys to be struck with different amounts of force and held for different amounts of time. In an analogous manner, engineering students can exercise their understanding of technical material by repeatedly using the same pattern of information, but emphasizing different aspects of the pattern; that is, they can pose and solve several variations on a problem originally assigned by their instructor.*

Note that the original problem has led us to devise ten variations—effectively, ten new problems. This process is most effective if students are merely shown the strategy and they create their own variations. Hopefully, they eventually create problems that they don't know how to solve, then they initiate a dialog with the instructor. This process is systematic and can be applied to any problem; in fact, rather than solving 100 different problems, students seem to gain more by solving ten problems plus ten variations of each. More on variational problem posing can be found in a book by Brown and Walter.[4]

Earlier we noted that reformulation is a central aspect of creativity; this observation can now be pushed farther by noting that devising variations on a theme is itself a reformulation. Hence, as Hofstadter has discussed,[5] variations on a theme is the crux of creativity. Any new object, process, or idea is created by modifying, to a greater or lesser extent, existing objects, processes, and ideas. (There is, after all, nothing new under the sun.) This aspect of creativity undoubtedly reflects the way minds work—not by spontaneously creating a completely new neural net, but rather by continually modifying existing assemblies of neurons. But the lesson here is that in practicing variational repetition on solved problems students practice creating new things. And even though their first attempts are mundane and uninteresting, the habit, once acquired, can eventually serve them well.

Transition:

Level 5 (Posing Problems)
to
Level 6 (Making Connections)

Motivation: Having learned to solve a problem, we should then ask whether that knowledge can be applied to other problems within the same domain and to analogous problems in other domains.

Reformulation: Pattern, problem context, and solution are generalized to other domains.

LEVEL 6: MAKING CONNECTIONS

The understandings gained at Level 5 can require substantial effort and labor because they often require us to make substantial modifications to dendritic trees and neural networks. So once such modifications are made, we try to increase their usefulness by connecting them to other networks that represent other patterns and problem contexts. That is, we try to project our newly acquired understandings into other domains of knowledge. Sometimes ideas for cross-domain connections can be evoked by posing a simple heuristic: Having solved the immediate problem, can we now solve a similar problem or an analogous problem? But more often, we must employ cross-domain devices to help us find ideas that transcend domains. Cross-domain devices are relations, patterns, or procedures that are invariant under changes of context; thus, they can be extracted from one context and inserted into another. Such devices provide powerful ways to increase understandings, and therefore it is probably not surprising that relatively few of them are known. We are always seeking to add new cross-domain devices to our repertoire, for every such device gives us another way to learn. Five common cross-domain devices follow.

(1) Our most powerful cross-domain device is *mathematics*. This statement often surprises students, for they tend to view mathematics as a tool for computation. But the real value of mathematics is that its rules for reasoning are independent of context: mathematics is powerful because it is abstract. As a simple example, consider the exponential growth law

$$y = y_0 e^{\alpha x}$$

where α may be positive or negative. This one equation applies to certain processes in a number of very different and *unrelated* contexts. For example, it describes the decay of radioactive isotopes, the variation of density with altitude in a stagnant isothermal atmosphere, the growth of a population in a limitless environment, the cooling of a warm body in cooler surroundings, and the growth of capital in an interest-bearing investment.

(2) A second device for extending ideas across domains is provided by *scaling laws*. These devices exploit the extent to which certain behaviors are universal—independent of context—when variables describing phenomena are scaled appropriately. Thus, we have the many dimensionless groups that correlate fluid flow and heat transfer in transport phenomena, we have corresponding states ideas for correlating thermodynamic properties, and we have scaling laws for describing the behavior of materials near critical points.

More generally, we now have numerous disparate phenomena, referred to collectively as fractals, that are invariant under changes of scale. For example, the Brownian motion, first described by Robert Brown in 1828, originally referred to a microscopic scale; when viewed through a microscope, a minute particle displays random movements caused by collisions with molecules of the surrounding medium. But such movements are also observed on macroscopic scales in colloidal suspensions and on galactic scales in the motion of stars in open clusters, such as the Pleiades.

(3) Another effective way to cross domains is by using an *analogy:* the presumption that if two things have certain similarities, then they also have other similarities. Analogies can be structural or functional, and it is wise to keep clear which you intend in a particular case; the common pitfall is to assume that structural similarities imply functional similarities. Examples of fruitful analogies include those among the linear transport laws of Newton, Fick, Fourier, and Ohm. In thermodynamics, certain phase diagrams for vapor-liquid equilibria are structurally analogous to diagrams for liquid-solid equilibria. And in process control, artificial neural networks bear certain functional analogies to biological neural networks.

(4) Still another device is the *metaphor,* which we use to describe an unfamiliar thing in terms of some more familiar thing. Unlike an analogy, a metaphor typically attempts to relate two things that have neither structural nor functional similarities. Minsky has emphasized that we typically use spatial forms and concrete objects as metaphors for abstract ideas and concepts.[3] For example, we talk about an idea being solid, firm, fluid, or off-the-wall. More generally, we have family trees, the tree of life, the tree of knowledge, dendritic trees, and logic trees; we have roots of a family, the root of an idea, the root of the matter, the root of a problem, and the root of all evil; we have bridges between domains of knowledge, the bridge of time, a bridge over troubled water, and a bridge (no less) to the 21st century. One of the most compelling metaphors of recent years has been that of the desktop for manipulating operating systems on personal computers.

(5) The last cross-domain devices we mention here are various graphing templates for representing relations. The most common is a simple x-y plot, which shows how an effect is correlated with its cause. To have something less familiar, we show in Figure 2 examples of an interaction square,[3] which shows how two causes either reinforce or compete in contributing to a single effect. In the first and third quadrants of the square, the two causes act together to either amplify (quadrant I) or suppress (quadrant III) the effect. The interesting behavior occurs in the second and fourth quadrants, in which the two causes compete. If we have a mathematical relation for how the two causes contribute to the effect, then we can usually solve for the locus along which the two causes exactly compensate for one another. This locus transverses the second and fourth quadrants. A particular example appears in the bottom of Figure 2, which shows how temperature and flow rate contribute to a particular value of the Reynolds number for fluid flow.

In the discussion of Level 3 (Paper 2 of this series), we observed that organizing knowledge into patterns provides a mechanism for improving the efficiency of education. Understanding at Level 6 provides a similar opportunity for efficiency. At this level our intention is to find existing neural structures created in one context and apply them to problems in other contexts. When this can be done, we avoid much of the laborious effort required at Level 5 in making major structural changes to old networks.

Transition:

Level 6 (Making Connections)
to
Level 7 (Creating Extensions)

Motivation: Having learned to recognize and solve analogous problems in various domains, we should ask what problems can still *not* be solved, but which might be solved if we could extend, modify, or reformulate what we have learned.

Reformulation: Generalizations are modified to attack other problems.

LEVEL 7: CREATING EXTENSIONS

At Level 6, our understanding is sufficient for us to realize that a certain pattern, problem, or procedure, devised in one context, can be useful when transplanted *in toto* to another context. At Level 7 we realize that a complete transplant will not be useful, but if the pattern, problem, or procedure is modified, then the transplant will bear fruit. In some situations, the necessary modification can be generated by merely devising a variation on a theme, but more likely, we need a reformulation that is more elaborate than a simple variation.

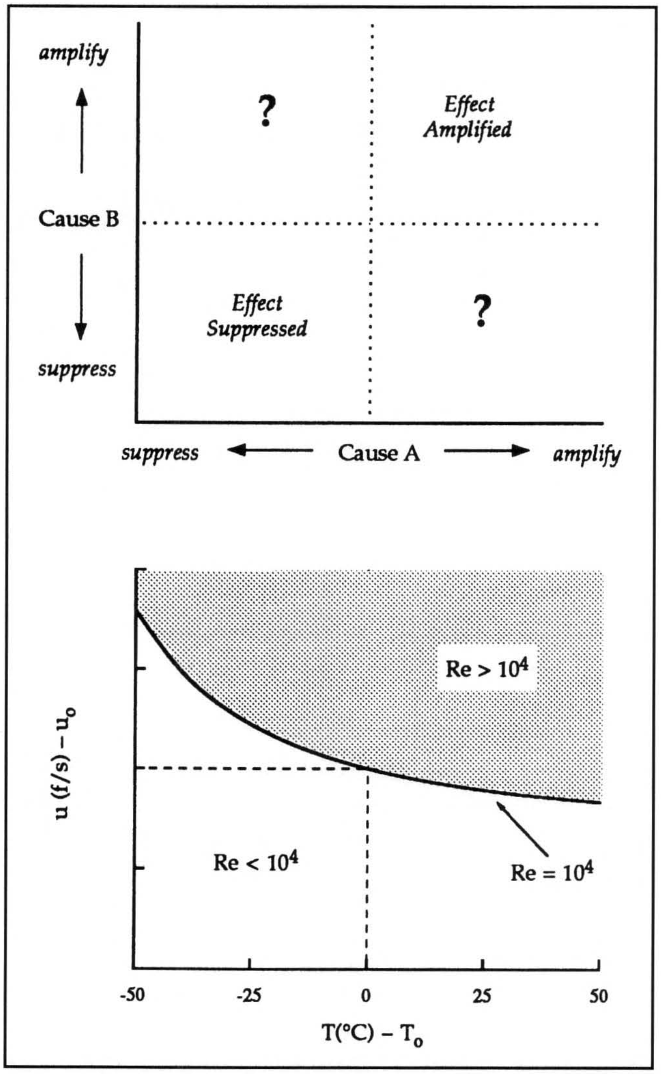

Figure 2. (Top) *Generic template for an interaction square that shows how two causes, A and B, contribute to one effect.* **(Bottom)** *A particular example, showing how temperature and flow rate combine to maintain the Reynolds number at 10^4 for water flowing through a 2.54-cm pipe. If the water temperature increases from the nominal conditions of $T_o = 50°C$ at $u_o = 0.72$ f/s, pushing the operating point into the shaded region, then the desired Re can be regained by adjusting a supply value to decrease the flow. Inversely, if the temperature decreases from T_o.*

That is, we are seeking a homomorphic projection across domains—a projection that identifies the essential features and that suppresses the inessential details.

An example is Maxwell's development of his theory for electromagnetic fields, which grew out of an analogy with vortices created in rotating incompressible fluids, as described by Helmholtz and Thomson. Here is Maxwell reviewing some of Thomson's papers on electrostatics and magnetism:[6]

> . . . illustrations of magnetic force . . . are not put forward as *explanations* of magnetic force. . . . They belong more properly to that remarkable *extension* of the science of hydrokinetics . . .

(The first italics is Maxwell's; the second is mine.)

Creating extensions is a first step in the more general topic of pattern posing and as such it links the study of established patterns to the research involved in creating new patterns. A principal strategy for posing new patterns is to shift, remove, or otherwise violate boundaries. By boundaries, we mean the assumptions and preconceptions that are inherent in any established pattern, concept, or procedure. Even experimental work involves assumptions; that is, we design an experimental protocol involving certain pieces of equipment under the preconceptions that certain phenomena will be observed and not others. But bounds serve as barriers that limit our thinking. So when a problem does not yield to attacks using established patterns and procedures, then we should test the bounds—examine our assumptions and preconceptions. As Root-Bernstein has noted,[7] in such situations it's not the problem that causes our lack of comprehension; rather, the impasse arises from assumptions that we take for granted.

Bounds are a product of negative thinking. Up to now, this paper has focused on positive thinking—on identifying ways to promote firing of useful patterns of neurons. But the brain has both inhibitory and excitatory synapses, so not only can we learn productive ways to think, but we can also learn to *avoid* unproductive ways to think. By imposing bounds on positive thinking,[3] negative thinking helps us be more effective because it helps us avoid wasting time on unproductive and counterproductive trains of thought. But we don't want the bounds produced by negative thinking to be too rigid because creative extensions can sometimes be found by shifting those bounds or by recognizing that some bounds have been misinterpreted or are inappropriate. Achieving a balance between positive and negative learning requires a delicate hand on the part of the instructor, for overemphasis on negative thinking can easily suppress creative impulses in students.

Lastly, note that violating bounds—juxtapositioning the incongruous—is a principal attribute of intellectual humor. Indulgence in intellectual humor exercises the mind in violating bounds and produces combinations of thoughts that might otherwise remain unconnected. It is a conceit of mine that such exercise preserves some flexibility in neural networks, and it might—just might—represent some lowly practice at creating extensions.

CONCLUSIONS

In this series of papers, we have presented a strategy for studying technical material; the strategy is organized into a hierarchy of seven levels. We enter the hierarchy at Level 1 when our attention is drawn to a topic and we begin to pose questions about it. We leave the hierarchy, as it applies to a particular topic, at Level 7 when we begin to consider how the topic's objects and concepts can be modified so that they can be applied to other topics. Note that problem solving, at Level 4, occupies the central level in the hierarchy, but problem solving is neither the goal nor terminal point of the hierarchy.

An overriding theme of these papers has been that anything interesting or useful has multiple meanings, and understandings of those meanings arise out of connections: connections among objects and concepts to form meaningful patterns, connections between patterns and a problem context, connections among different problems and their contexts, and connections among different domains of knowledge. The hierarchy of understanding provides a scheme for systematically making connections. The hierarchy can be used by instructors to help organize how material is presented to students and to help assess student understanding. Similarly, it can be used by students to help organize their study of a topic, to assess their comprehension, and to identify what should be done to move to the next level.

We have devoted considerable effort in trying to find meanings for the word *understanding*. Perhaps some additional insight can be gained by inverting the issue and identifying things that are *not* understanding:

1) *Verbal fluency is not understanding*—people can engage in conversations about a topic without being able to answer questions about the topic or to explain the topic to others;
2) *Experience is not understanding*—people routinely use automobiles and computers without understanding how such things work;
3) *Solving a problem is not understanding*—people can solve a problem without realizing how they solved it and without being able to explain their procedure;
4) *Making predictions is not understanding*—before 500 B.C., the ancient Babylonians had correlated sufficient observations so that they could predict lunar eclipses,[8] but they could not explain the geometry that causes an eclipse;[9]
5) *Accumulated knowledge is not understanding*—the Nobel laureate Albert Szent-Györgyi once remarked

Continued on page 39.

Toward Technical Understanding

Continued from page 35.

> that during his study of muscular action he came to realize that the more he learned, the less he understood, and so he became fearful of finally learning everything, but understanding nothing.[10]

The discussions here raise many questions that would seem to serve as starting points for further, more detailed investigations. Here is a list of some of the more obvious ones.

1. If the pattern can indeed serve as the fundamental unit of understanding, then what are those patterns that distinguish one topic from another? For example, what patterns distinguish transport from thermodynamics and thermodynamics from reaction kinetics? Then, by extension, what patterns distinguish chemical engineering from chemistry and from other engineering disciplines?
2. Repetition is necessary to solidify certain kinds of understandings, and therefore some amount of redundancy needs to be incorporated into a curriculum. But efficiency in education can be attained by appealing to patterns and other devices that cross subject domains. To what extent can a curriculum be made more effective by organizing it around patterns rather than topics?
3. What are appropriate cues that will activate, in student brains, proper patterns and homomorphic projections needed to address particular problem situations? Are there minimum numbers of cues that are sufficient?
4. Can we contrive a complete list of devices for making connections across subject domains? Is there a minimum number of such devices that a student should be able to use? What are the most effective ways for students to develop facility with cross-domain devices?
5. Can we devise systematic procedures for identifying and testing default assumptions and probing tacitly assumed boundaries?
6. Are there ways to gauge the importance and impact of negative thinking relative to positive thinking?
7. What indicators can we devise for determining when students successfully make a transition from one level of understanding to another?
8. Presumably, we do not expect all students to achieve the same levels of understanding. What levels are appropriate for BS students? For MS students? For PhD students?
9. Traditional descriptions of brain function use time to identify two kinds of memories: short-term (you look up a phone number and remember it only long enough to dial it) and long-term (you still remember your name). But recent evidence suggests a third: intermediate-term memory, in which a buffer (perhaps the hippocampus) is loaded while structural changes are made in the brain to lay down the corresponding long-term memory. Thus, students who cram before a test often do not retain the crammed information because they are only loading a buffer, not creating long-term memories. This suggests that simple linear progression through material over a semester may not be as effective as some cyclic procedure in which important patterns are revisited at intervals. Revisiting amounts to repetition, which stimulates creation and solidification of long-term memories and pares away superfluous scaffolding. If this conjecture were confirmed, what kinds of cyclic presentations should be used? What are the optimum times between re-exposure to the same patterns?
10. Finally, note that throughout these papers we have emphasized *what* rather than *how*. So, how do we help students progress through a hierarchy of understanding?

> *Understanding never ends.*
> Minsky[3]

ACKNOWLEDGMENTS

Many of the ideas presented in this series were tested and clarified by continually referring to Marvin Minsky's book,[3] *The Society of Mind;* without that book, these papers would have taken a very different form. Over the years of my struggle to understand understanding, I have learned much from discussions with my colleagues R.W. Rice (Clemson) and J.P. O'Connell (Virginia); my thanks for their forbearance.

REFERENCES

1. Rolls, E.T., and A. Treves, "Neural Networks in the Brain Involved in Memory and Recall," *Progress in Brain Res.,* **102**, 335 (1994)
2. Calvin, W.H., *The Cerebral Code,* MIT Press, Cambridge, MA (1996)
3. Minsky, M., *The Society of Mind,* Simon and Schuster, New York, NY (1986)
4. Brown, S.I., and M.I. Walter, *The Art of Problem Posing,* Erlbaum, Hillsdale, NJ (1983)
5. Hofstadter, D.R., "Variations of a Theme as the Crux of Creativity," in *Metamagical Themas,* Basic Books, Inc., New York, NY (1985)
6. Maxwell, J.C., "Review of Reprint of Papers on Electrostatics and Magnetism by Sir W. Thomson," in *The Scientific Papers of James Clerk Maxwell,* Vol 2., W.D. Niven, ed., Cambridge University Press, Cambridge (1890); reprinted by Dover Publications, New York, NY, p. 301(1965)
7. Root-Bernstein, R.S., *Discovering,* Harvard University Press, Cambridge, MA, p. 296 (1989)
8. Neugebauer, O., *The Exact Sciences in Antiquity,* reprint of 2nd ed., Harper & Brothers, New York, NY (1962); lst ed., Princeton University Press (1952); 2nd ed., Brown University Press (1957)
9. Toulmin, S., *Foresight and Understanding,* Indiana University Press, Bloomington, IN (1961); cited in D.A. Crosby and R.G. Williams, "Creative Problem-Solving in Physics, Philosophy, and Painting: Three Case Studies," in *Creativity and the Imagination,* M. Amsler, ed., University of Delaware Press, Newark, DE (1987)
10. Szent-Györgyi, A., "In Search of Simplicity and Generalizations (50 Years Poaching in Science)," in *Current Aspects of Biochemical Energetics,* N.O. Kaplan and E.P. Kennedy, eds., Academic Press, New York, NY (1966) ❐