

Is Abstraction a Kind of Idea or How Conceptualization Works?

Commentary on a Special Issue of Cognitive Science Quarterly

William J. Clancey

NASA/Ames Research Center¹

In this commentary, I review papers by Ohlsson & Regan (O&R), van Oers, and Dreyfus, Hershkowitz, & Schwarz (DH&S). The papers are nominally about abstraction and learning, but emphasize different kinds of problems and levels of analysis. O&R focus on mathematical, “domain independent” characteristics of abstract thinking, claiming that experience in a domain is not the main determinant of scientific discovery. van Oers focuses on the development of abstraction within activities, especially as a sequence of nested domains of concern. DH&S emphasize how nested conceptualizations co-define and provide meaning for each other (a dialectic relation).

In this review of papers on the topic of “abstraction and learning,” I distinguish between *articulated mathematical theory* (e.g., terms and laws), *artifacts* (physical constructions), and *conceptualization* (mental processes). I suggest not polarizing abstractions and “domain-specific knowledge”—abstraction within the conceptual domain of mental processes is a relative term (Arnheim, 1969). There is no *experiential* sense in which an understanding is “domain independent” (rather, this is a defined, theoretical distinction within mathematical theory). In particular, educational research should use the term “abstraction” to include procedural and causal abstractions emphasized in (“knowledge-based”) models of scientific discovery and learning (Buchanan & Wilkins, 1993). Second, I heartily endorse “an activity-based, dynamic meaning of abstraction” (DH&S, p. 36), and suggest that we might study this as a neural, conceptualization coordination process. In particular, different modalities of conceptualization (verbal, visual, gestural, etc.) and perceptual-motor systems are being coordinated in ways that are sometimes visible in the experimental interactions considered by these papers.

¹ Address for Correspondence: William J. Clancey, Computational Sciences Division, MS 269-3, NASA/Ames Research Center, Moffett Field, CA 94035 (on leave from Institute for Human-Machine Cognition, UWF, Pensacola) bclancey@mail.arc.nasa.gov

Abstraction in Natural Sciences: Ohlsson and Regan

Ohlsson and Regan argue against the idea that “specific expertise is the main determinant of creative problem solving” (p. 31), as well as the more general claim that “specific knowledge is the main determinant of cognitive change” (p. 32). They are not saying “that specific knowledge is unhelpful” (p. 31). Rather, “to operate beyond the boundary of current knowledge, one has to be guided by something in addition to the facts and skills already acquired” (p. 32). Based on an experiment in which students solve a problem claimed to be analogous to the discovery of the molecular structure of DNA, Ohlsson and Regan conclude that in “the natural sciences... the learner’s repertoire of prior abstractions, not what he or she knows about the domain itself, might be a crucial cognitive resource” (p. 34). In short, they claim that abstract information is more helpful than specific information and is often the main determinant in scientific discovery.

I argue that Ohlsson and Regan have adopted an unnecessarily narrow definition of what constitutes an abstraction, and hence taken an unnecessarily harsh (and historically inaccurate) judgment about knowledge-based models of cognition (AKA expert systems). More importantly, from the perspective of situated cognition research, their example and analysis presents a limited view of scientific discovery (Dewey, 1938b) and an inadequate theory of the mechanism of conceptualisation (Clancey, 1999). Much stronger and far reaching claims about the nature and role of abstraction can be made, if only the work in expert knowledge modeling is more fairly appraised and the practice of scientific discovery is more carefully examined.

What does discovery in natural sciences look like?

To begin, let’s consider a current example of discovery in the natural sciences, namely the interpretation of surface features on Mars. Recently, scientists have discovered ravines and gullies on cliffs, which appear relatively fresh. The most obvious cause, flowing water, is not possible:

“The visual evidence, although compelling, challenges experts’ understanding of Mars. With temperatures well below freezing and an atmosphere 1000 times thinner than at sea level on Earth, how could liquid water be stable long enough to flow downhill?” (Vaugh, p. 12)

To resolve the dilemma, the scientists have adapted what they know about water to tell a new story about water flows on Mars:

“We’ve come up with a model to explain these features and why the water would flow down the gullies instead of just boiling off the surface,” explained Edgett. “When the water evaporates it cools the ground – that would cause the water behind the initial seepage site to freeze. This would result in pressure building up behind an ‘ice dam.’ Ultimately, the dam would break and send a flood down the gully.” (From “Questions and Answers: Factinos” p. 21)

Consider the understanding of causal processes that comes into play here: evaporation and freezing of liquids, pressure building, weakening of dams.

Reasoning by analogy from their general understanding, scientists have gone from the particular (observable features on Mars) to the abstract (theories about liquids) and back to the specific (a causal story about ice dams on Mars). This is abstraction at work. The scientists extended their knowledge of Mars by adapting principles they had developed in a different context (liquid behavior on Earth).

As a second example, consider the problem of explaining huge expanses of patterned ice and ground on Mars:

“Large pits, troughs, and flat mesas on the south polar cap do not seem to appear anywhere else on the planet, leading scientists to think the distinct landforms may have something to do with the amount of frozen carbon dioxide in the region.” (Vaughn, p. 13)

Here a different kind of abstraction is brought into play. Given that one area has features that other areas do not have, they ask, “What’s different here?” The north pole of Mars lacks such features. How is the south pole *otherwise known to be different* from the north pole? It has large amounts of frozen carbon dioxide. Perhaps CO₂ plays a causal role in the formation of the patterns. Causal reasoning, which is central to most discovery in the natural sciences, is replete with schemas like “What’s different here?” Indeed, one way of understanding student reasoning mistakes is to see what causal abstractions they are conceiving. For example, Feltovich, Spiro, & Coulson (1989, p. 142) show how medical students often inappropriately apply the length-tension relationship of elastic structures—a causal abstraction—to cardiac muscle.

Now, the examples of abstraction I have presented—which I claim are pervasive in knowledge-based models²—do not count as abstractions by the definition of Ohlsson and Regan. This could be why they believe abstraction has been “all but ignored” (p. 5) by contemporary cognitive science.

Unnecessarily narrow definition of abstraction

O&R’s experiment shows that formal constraints (shapes, symmetry, pairing, elongated chains, etc.) are “the main determinants” for solving topological problems. This is not very surprising, but it doesn’t say much about scientific discovery in the natural sciences or creative problem solving in engineering, medicine, business, and so on.

The problem of determining the molecular configuration of DNA is a kind of design layout problem. The discovery proper specifies the topological relations of known molecular components. It is not surprising that mathematical-spatial relations (the “abstract hints”) are essential. As stated, the problem is not to reason about causal processes (e.g., how the DNA duplication process is involved in cancer), but to specify a particular molecular structure. Using X-ray crystallography data, heuristic search applies con-

² I use the conventional term “knowledge-based” here, but a better name is “model-based.” Most of the early work in the 1970s into the 1980s used the term “model” only for qualitative simulation models, viewing classification models in programs like Mycin as “shallow” and hence not models at all (Clancey 1989). See also the elaborate comparison between human knowledge and computer models in Clancey (1997).

straints of “angles, distances and sizes” to lay out the components. Chemical facts, such as the nature and formation of molecular bonds, were necessary but “insufficient to constrain the search...the abstractions provided the additional guidance” (p. 15).

But is this a complete or even typical view of the nature of “abstract ideas in conceptual discovery and learning” (the title of O&R’s paper)? The experiment tells us about mathematical (geometrical and topological) abstractions and their role in formal puzzles, but what does it say about conceptualizing in general?

By O&R’s definition, apparently only mathematical relations count as abstractions: “A knowledge structure is abstract if it is represented in memory in such a way that it is independent of any particular domain of experience or knowledge” (p. 6). To them, an abstraction must be devoid of any domain content: “An abstraction encodes relational information while being silent about what is being related” (p. 6).

This notion of independence appears unnecessarily strong and psychologically suspect. Concepts fall more on a spectrum, *both specific and abstract*, not either-or. Abstraction, as it is naturally found in human thinking, should not be contrasted with concrete ideas (Arnheim, 1969).

The spectrum is obvious when we consider that not all abstract relations are mathematical. Other kinds of abstract relations are possible, involving causality, classification, temporality, etc. Abstraction, *per se*, is the relation of one concept to another, not a specific set of relations which are mathematical (e.g., spatial configuration relations).

Consider explaining the apparent gullies and ravines on cliffs of Mars. Clearly we must start with facts about Mars’ atmosphere, climate, and properties of water. But then we consider theories about water: its phases, how it flows, freezing. These are all abstractions *about liquids*. They are independent of specific puddles, raindrops, gullies, etc. But as concepts they are not “independent of any particular domain of experience.” Indeed, a great deal of causal knowledge is abstract—it relates specific experiences or events through classifications (e.g., “liquids”) and causal stories (e.g., evaporation and freezing) by which structural changes and behaviors can be understood (e.g., presence of liquids flowing in low atmospheric pressure).

Consider the “shoe store” abstraction in van Oers’ paper. Obviously, this concept is not independent of a domain, yet van Oers’ shows that the shoe store schema organizes the children’s behavior. If the shoe store schema doesn’t count as an abstract idea, then are O&R making general statements about discovery and learning or only telling us about mathematical-spatial configuration problems?

Besides the claim of domain “independence,” there is something strange about the idea of “encoding,” which is central to O&R’s definition of abstraction. Consider an abstraction programmers currently use in designing graphical user interfaces, the idea of a “wizard.” A wizard is a sequence of dialog boxes (panels), identical in size and placement on the screen, each having “previous, cancel, next, finish” buttons. (Titles and other design elements, such as icons, are also often similar.)

What does the wizard abstraction *encode*? Nothing. It is just a pattern, a schema, by which graphic user interface structures and processes may be conventionally organized. The notion of encoding by representations and knowledge has been incisively refuted by Bickhard and Terveen (1995). Indeed, the paradox of O&R's view is apparent when they tell us that "some branches of pure mathematics have not found any use within natural sciences" (p. 8). How can a concept encode relational information if *it has no instances*? What is it encoding?

O&R make no claims about the origin of abstractions (p. 9), and indeed their restrictive, math-centric definition is at odds with other theories of memory, such as MOPS (Schank, 1980), which make abstraction integral with all learning. Furthermore, theories of analogical reasoning (which appears in the notion of an "ice dam" on Mars) most definitely do not "dispense with abstractions" (p. 5) as O&R claim. Theories of analogical reasoning explain how we "move beyond our understanding" by relating structures or behaviors in one domain to those in another domain (Gentner and Stevens, 1983). Parallel relations in different domains can be "mapped," suggesting structural and causal inferences in a new domain³. Obviously, the idea of a dam is an abstraction, from which we can reason about water dams (made by people and beavers) and ice dams (formed by natural processes)—and thus explain gullies on Mars. One way we move "into unknown territory" (p. 5) is by such analogical reasoning.

O&R might argue that the examples of causal properties of liquids, dams, a wizard, and shoe stores are really "generalizations," and not abstractions (p. 9). Whatever we wish to call such general concepts, they obviously play a central role in scientific discovery and learning. Such concepts are not independent of specific domains in the way characterized by O&R, they do not encode anything, and as I will show, they have been essential in knowledge-based models of reasoning and discovery for at least 25 years.

Too harsh judgment about expert systems

The specificity hypothesis, which O&R argue against, is "that cognitive ability is primarily a function of expertise" (p. 4). In expert systems parlance, "knowledge is power."

O&R say that the "contemporary *Zeitgeist* all but ignores the role of abstraction in human cognition" (p. 5). But I believe they have misinterpreted the notion of domain-specific expertise. The knowledge-based approach never excluded abstraction, and indeed emphasized its role. "Expertise" includes whatever representations and methods are appropriate for solving problems in the domain—including abstractions. Look again at the cited claim from Lenat and Feigenbaum (1987, p. 1173): "specific knowledge... the concepts, facts, representations, methods, models, metaphors, and heuristics about its domain of endeavor." The literature gives many examples of abstract problem-solving and learning strategies, often called "meta-knowledge" in the 1970s.

³ But see Schön (Clancey, 1997, Chapter 9) for a critique of "mapping" theories. The relation of "seeing as," for example, involves non-descriptive relations, not mapping between propositions. See also the discussion of van Oers that follows.

As an example of a domain-specific abstraction, Davis' metarules modeled how an expert directed use of knowledge, according to circumstances, "If the current culture is not from a sterile source, don't bother trying to deduce the current organism identity from the identity of previous organisms" (Davis, 1982, p. 418).

Similarly, Lenat's AM (Automated Mathematician) included hundreds of domain-general abstractions for discovery: "After creating a new specialization S of concept G, explicitly look for ties between S and other known specializations of G" (Lenat, 1982, p. 180). Or consider: "A generalization of X is interesting if all the previously known boundary non-examples are now boundary examples of the concept."

Ironically, AM's discovery heuristics are couched in a vocabulary of abstractions that obviously fits O&R's mathematical definition (e.g., specialization, boundary example, definition)—"independent of any domain" (p. 6). Thus, the work in AM, which falls squarely within the knowledge-based approach (indeed, Lenat is cited by O&R) disproves the claim that the knowledge-based approach ignored abstraction. Indeed, Lenat's work is not only based on the centrality of abstraction, it is explicitly about discovery learning.⁴

Feigenbaum et al.'s claims were most strongly influenced by the experience in developing DENDRAL—a tool for scientific discovery, emphasizing the centrality of abstraction—in the very same domain of molecular configuration analyzed by O&R. DENDRAL determines molecular structure from mass spectrometry data. The reasons for DENDRAL's success are listed by Buchanan (1992):

- A canonical representation of the "answers" (chemical structures)
- A systematic generator of the search space
- A program that inferred plausible constraints for the generator from the input data; and
- A strong enough theory of the analytic instruments to make testable predictions.

Improving the rules for analyzing mass spectrometry data drastically reduced the number of structures that had to be generated before the answer was found. Thus, the tradeoff was between a) a general problem solver that used mathematical constraints and b) specific heuristics that incorporated information about the relation of mass spectra and molecular structure. The issue was efficiency for computing a solution, which meant practically whether an exhaustive search would produce results in our lifetimes:

"Generality is being questioned when we ask: how broad a universe of problems is the problem solver prepared to work on? Expertness is being questioned when we ask: how good are answers and were they arrived at with reasonable cost? Generality has great utility in some ways, but is not often associated with superior performance. The experts usually are specialists." (Feigenbaum, Buchanan, and Lederberg, 1971, p. 165)

⁴ In fact, Lenat cites Adams, 1974, *Conceptual Blockbusting*, which is a well-known proponent of the idea that abstractions are important for creativity.

“The designer of a problem-solving system...is inclined, at least initially, to try to match the generality of the physical process of chemistry with generality of the reasoning process. Yet he soon finds, paradoxically, that he cannot afford this match, that he must retreat and rework his analysis into more specialized forms if he is to be able to use his problem solver on real problems.” (p. 168).

“DENDRAL’s Planning Rule Generator is powerful for the supra-family of saturated, acyclic, monofunctional compounds, but is useless for all other classes of compounds. When we extend DENDRAL’s capability to families of cyclic molecules, we may have to write a new Planning Rule Generator.” (p. 189)

Feigenbaum, et al.’s emphasis on expert knowledge is a reaction to the early view that problem-solving was mainly a matter of canonical representations (how a problem was described) and broadly applicable search algorithms (especially means-ends analysis used by the General Problem Solver)—which they term “generality.” DENDRAL, viewed by some as the first expert system, showed otherwise. The early version, which attempted to generate all possible structures, without ordering them with regard to the spectral data, was too inefficient. Here expertise is any fact, representation, method, etc. that relates specific mass spectrometry data to a theory of spectral analysis: “By far the most powerful method of gaining effective control over the generator is to force its search to be relevant to the specific problem data given as input (the spectral data)” (p. 172).

The lesson was that a priori information about the molecule’s configuration was essential for producing solutions more quickly. O&R make the same conclusion for molecular structure configuration problems. However, Feigenbaum, et al. called these relations “specialized” even though they are mathematical-spatial; O&R call these relations abstractions, evidently *because* they are mathematical-spatial.

Many more examples, indeed hundreds, can be given from the literature of knowledge-based scientific discovery (e.g., Buchanan & Wilkins, 1993). This work, which O&R appear to misinterpret, supports their hypothesis that abstractions are essential for discovery learning. But more importantly, these computational models demonstrate that:

- different kinds of abstractions are relevant (conceptual, procedural, causal),
- the pattern holds for all knowledge-based problem solving (design, diagnosis, prediction, control), and
- such abstractions fall on a spectrum from domain-specific relations (such as Davis’ meta-rules) to domain-general relations (e.g., analogies with liquids on Mars and Earth) and mathematical (“domain independent”) relations (as the topological relations in DENDRAL).

Limited view of scientific discovery

Like the knowledge-based approach, O&R's analysis equates natural science (p. 34) with puzzle solving and understanding gained through instruction (p. 19, 31) with the *practice* of scientific discovery. Instead, I would argue that four phenomena that need to be distinguished:

- Human expert problem solving (e.g., medical diagnosis)
- Human scientific discovery in natural science (e.g., recent Mars analyses)
- Model-based inference in computational systems (e.g., MYCIN and DENDRAL)
- Learning by being given a goal state and representational artifacts to manipulate (more generally, a notation and operators) (e.g., the task in O&R's paper)

For discussion about the relation of human knowledge and computer representations pertinent to these distinctions see Clancey (1997). Here I will restrict my comments to some observations about the conceptual content of the experimental example, the practice of discovery, the process of conceptualization, and social considerations.

Content

O&R's experiment reformulates the DNA structure problem into a puzzle with given 2D shapes. In effect, O&R are arguing, as the list of insights makes explicit, that the essential problem of DNA structure elucidation was geometric and topological. (Though we might ask what other biochemical knowledge was required to discover these clues?).

Configuration problems, although important, are only one aspect of scientific discovery. The Mars example shows how causal reasoning may be central. Arguably, causal reasoning is what scientific discovery is all about, with structure elucidation playing the role of defining the objects and their parts that enter into causal stories. O&R's claims about the nature of abstractions (independence and encoding) are not warranted for natural science generally.

The puzzle form of O&R's experiment differs from natural scientific discovery in the following ways:

- *The students were not given irrelevant information; a major issue in scientific discovery is to identify important data and know what can be discarded.*
- *The problem is made up. By design it does not draw on background knowledge about archaeology, ancient cultures, priests, etc. And lo and behold, it shows that such background is not relevant (i.e., "specific expertise is not the main determinant in creative problem solving").*
- *The problem was topological in nature, so mathematical relations were instrumental (cf. p. 31). All one can say is that mathematics plays a role in science, not that it is "the main determinant."*

Practice

Just as important, the view of problem solving as occurring exclusively in the head, an assumption of the knowledge-based programs that O&R criticize, is explicitly adopted by them. They refer to conceptual creation as “the construction or discovery of new concepts, ideas, problem solutions and theories, *as opposed to artifacts*” (p. 5, my emphasis). Ironically, their own presentation makes a crucial point about Watson and Crick’s approach: “They manufactured metal pieces...physically manipulating and combining the metal pieces that represented the relevant molecules” (p. 11). Similarly, the O&R’s experimental materials included foam board shapes. Problem-solving, thinking, reasoning, whatever you want to call it, is a human *behavior*, and as such occurs as an experience in time. As the very example presented by O&R demonstrates, scientific discovery is very often carried out in and by way of physical interaction with artifacts. The interactive, behavioral aspect of reasoning is emphasized by *situated cognition* research (Clancey 1997; Greeno, 1994). In practice, “specific” *physical* representational tools for learning are crucial, such as Watson and Crick’s metal model (as well as mental models).

The process of conceptualization

O&R also adopt the knowledge-based view of conceptualization, which is essentially an manufacturing metaphor: assembly of instantiated templates, with much talk of extraction, parts, components, pieces, etc. (p. 7) This metaphor views concepts as static parts and new idea formation as physical configuration. Parts are supposedly stored in memory from which they are retrieved (copied? moved?) to another place (“working memory”) and then used to form new configurations. The template parts then are left unchanged to be used again some other day. (Lave [1988] describes this in terms of tools being left in a shed; see Clancey [1997, Chapter 3; 1999, Chapter 1].)

As O&R indicate, they don’t explain where abstractions come from. Any model of discovery learning that leaves this out must be somewhat suspect. For if a theory of discovery learning doesn’t explain the formation of abstractions, what will? What is missing in O&R’s theory is the idea that abstractions are not merely applied, they are *formed* and *changed* through reasoning itself. Indeed, the stronger claim is that all experiences, all remembering, and all thinking involves reconceptualization. Reasoning is not retrieving and reassembling, but reactivating and re-relating concepts to form generalized sequences and compositions (Clancey, 1999).

The paradox of “thinking new thoughts” (p. 5) partly stems from the manufacturing view. By the manufacturing assumption, discovery involves creating a new part, and “new” by definition means it must be a combination of old parts and/or from outside the domain. Of course, this configuration must be organized. Rather than focusing on the *organizing process* (heuristics like Lenat’s and Davis’s for controlling attention), O&R’s manufacturing model uses templates. Indeed, one must be struck by the unity of O&R’s view: According to them, what happens in the brain is very much like what we see in the external behavior of scientific discovery—it’s all topological puzzle solving, templates and pieces being fit together.

An opposing view of memory, knowledge, and reasoning (detailed in Clancey, 1997, 1999) is that meanings change as concepts are brought into new relations. Indeed, the very process of subsuming an idea under a broader concept changes the more general concept.

Consider how the meaning of “planet” is changing as we are discovering remote spherical objects between the size of ten times Jupiter and brown dwarf stars. The very idea of “planet” is changing: Must a planet revolve around a star? If a planet is a source of heat (like Jupiter), then how is it different from a star? Could a planet be much larger than any planet in our solar system? Could stars develop into planets? These questions involve causal stories and blended classifications, changing meaning relative to each other. It appears odd to talk about “parts, interrelated in particular ways” (p. 7).

Insofar as concepts are discrete entities (and it is of course sometimes useful to use that metaphor), are they unchanging parts? How does relating concepts change them? What constitutes a conceptual space? One explanation for cognitive change is that concepts in different modalities (verbal, visual, auditory, physical-manipulative-postural-gestural) are re-coordinated (Clancey, 1999). Thus, we step outside of a verbal space of terms, definitions, and associations by coordinating them with, for example, visual relations, providing the basis for a new articulation (Arnheim, 1969; Miller, 1984). Ironically, O&R cite the Gestalt view of cognition (p. 7, 34), which emphasizes the non-descriptive nature of perception and nonverbal aspect of understanding, but they don’t build on this aspect of Gestalt theory in their consideration of scientific discovery.

The template view also misconstrues the nature of human problem-solving in a given field. I agree that “The process of constructing a novel conception in field X cannot be entirely reduced to the application of what is already known about X” (p. 33). But the same thing can be said about day-to-day problem solving in field X. Reasoning and creative thinking are on a continuum; all experience changes our understanding. There is no such thing as merely “applying” knowledge, because remembering is conceptualizing, and hence meanings and conceptual coordinations are always adapting to circumstances. Again, this point has been well articulated in the situated cognition literature (e.g., see Clancey, 1997, Chapter 3).

Returning to O&R’s assumption: “To operate beyond the boundary of current knowledge, one has to be guided by something in addition to the facts and skills already acquired” (p. 32). One is reminded of Plato’s unchanging forms in the mind, the templates by which ideas develop—something we know, beyond what we have learned from experience. O&R resolve the quandary by saying that besides facts and skills, there are abstractions in the mind that are *physically separated* from domain-specific knowledge. I resolve the quandary by saying that facts and skills are associational, sequential, and compositional relations in different modalities, and that conceptualization is a process of re-coordinating these relations, which changes what the concepts mean and the work they do.

The role of experience, identity, ambition

Part of the scientific puzzle described by O&R is how “two researchers who were at the beginning of their careers and who... had not previously identified the structure of some large molecule” could have solved the DNA structure problem (p. 15). Indeed, the pattern is well known: Two guys who had never invented a computer before or run a company formed Apple. A lone programmer who had never written a large program and didn’t have a college degree invented Napster—a conceptual breakthrough that turned the music industry upside down.

Certain forms of mathematical-mechanical-logical construction do not rely on decades of experience in a given field. But does this mean that research labs should hire promising high school students instead of expensive PhDs? In some respect, breakthroughs are indeed possible by ignoring (or not knowing) the approaches other people take for granted (and demonstrably haven’t worked). Yet, Watson and Crick didn’t work in a vacuum; they undoubtedly had much help by knowing what approaches failed and by using clues others had generated. A social-historical analysis would consider more broadly how knowledge developed in the community.

Motivation must also be considered. The desire and energy to succeed, to make a name for yourself, count for a lot. This is not to discount in any way the essential role of facts, methods, models, etc. relevant to the domain of inquiry, but it does say that more than knowledge is important for explaining human performance. This is another conclusion from situated cognition (e.g., see Wenger, 1998).

Conclusion: The thesis is not new, and it’s too limited

O&R’s experiment goes to great lengths to show something already consistent with knowledge-based programs. AM showed that mathematical abstractions were essential for discovery in a formal domain. DENDRAL showed that topological constraints were important for determining molecular structure. Many other computational models show the varieties of abstraction and reasoning relevant to scientific discovery.

What did O&R accomplish?

1. They started by assuming that formal (mathematical) insights are essential in learning and that hierarchical, procedural, and causal relations referring to domain concepts are not abstractions.
2. They reduced a natural science situation to a formal analog (where mathematical relations are salient).
3. They showed that formal problem solving is assisted by formal cues.
4. They concluded that hierarchical, procedural, and causal relations referring to domain concepts are often not “the main determinant” in scientific discovery.

Given the simplifications of steps 1 and 2, the conclusion in step 3 is not surprising, and the conclusion in step 4 is unwarranted. O&R’s distinction in the specificity hypothesis between “specific knowledge” and abstraction is a false dichotomy. A more magnanimous view of expert systems informs our understanding of scientific discovery as involving many varieties of abstract reasoning. Better yet, situated cognition has begun to reveal the mechanism

of conceptualization, the physical-interactive aspect of problem solving, and the social aspect of creativity. Thus a broader view of cognition shows more completely the nature of abstraction and its role in scientific discovery.

Abstraction as cycles of doing and reflecting: van Oers

Van Oers' nemesis, and rightly so, is the "transmission model that provides abstract models to pupils" (p. 25). By van Oers' definition, abstract thinking is a process of contextualizing an experience, providing a framework by which particulars become "a situation" and hence action may be organized—"ascending to the concrete." Such thinking is natural and need not be taught by transmitting abstractions; the understanding develops interactively, in a discursive process by which meaning is negotiated. Accordingly, students should be given "a perspective on where they are going" by having a role in the contextualisation process itself (p. 25).

Van Oers' examples show us what abstract thinking looks like; how it develops and its sequential structure. Abstraction is an activity, something people do, a behavior. Abstraction often involves the manipulation of physical materials and cycles of re-perceiving to discover new features and conceptual reframing. Abstraction takes time. In response to van Oers' rhetorical question, whether "we are human beings because we can abstract," (p. 1), I would say, yes, because we have a form of consciousness that is different from other primates. Human higher-order consciousness involves perceptual-motor and conceptualization cycles of doing and reflecting, by which conceptualizations are held active over time, so they may be compared, sequenced, and composed (Clancey, 1999). van Oers' analysis can be understood as emphasizing how focus details in our experience (perceptions) are related, such that higher-order conceptualizations are brought into play, giving a contextual order to perceptual experience and guiding subsequent attention and action.

In contrast, the classical view of "abstract thinking," as outlined by van Oers, emphasizes creating types, a process of generalizing by *removing* circumstantial aspects of time and place. The Piagetian notion of abstraction is better (p. 3), characterizing abstraction as a scheme of coordinated actions, of operations, and not just a scheme of objects. Vygotsky's notion of stability of (word) meaning across situations is a manifestation of abstraction, revealing that the process of relating doesn't necessarily involve a fixed, stored template (word meaning) that is matched against present experience. Rather comprehending is inherently a learning process.

Van Oers wishes to emphasize that abstraction is not exclusively a reduction (from concrete to abstract) but operates in both directions:

"So the development of thinking cannot be adequately described as a process going from the concrete to the abstract, or going from the contextualised to the decontextualised. The ... relationship between the concrete and the abstract can be seen in a more dialectic relationship, that is elaborated in the course of a conversational process of focusing and object construction" (p. 4).

Van Oers view of abstraction and his theoretical stance places him in the avant-garde of cognitive science. Consequently, the first four parts of his paper are likely to be difficult to understand for people who haven't learned about this perspective elsewhere. In contrast, the examples and analysis in section 5 are wonderful. It would have been better to place the examples first.

My other quibble is with the use of the term “discursive,” which usually means verbal or argumentation (as van Oers indicates). The understanding process of activating and adapting abstractions may more broadly be characterized as *dialectic*, such that the abstraction and the particulars are mutually influencing each other. The process of looking, manipulating materials, conceiving meaning, looking again, etc. is interactive, or more precisely transactional⁵. Describing and arguing play an important part, but it is misleading to characterize the overall process as discursive, for perception and conception always have tacit, nonverbal aspects.

In the sections that follow, I present my own examples that I hope will clarify van Oers' theoretical presentation, as well as develop further the idea of abstraction as an activity.

Example of abstraction in programming languages

I will present two examples involving Excel programming. The first example concerns the process by which a “dialog sheet” in Excel's programming language became a “user module” in Visual Basic for Applications. The second example concerns how the programming language is used.

The Excel language is useful for understanding the nature of abstraction because:

- Language constructs (e.g., worksheets) have a visible, physically manipulable manifestation—showing how abstractions play a role in creating programs
- The language is formal and has mathematical aspects (e.g., the hierarchy of workbooks and worksheets)—making the example accessible to those who think of abstraction in mathematical terms
- The language is internally organized as objects that have behaviors (i.e., “methods” such as “move Row”)—so meanings of terms (“Row”) are often topological, but relations between terms are not reducible to propositional definitions.

Excel is also valuable because it brings out the contrast between instructional-theoretical and developmental-historical perspectives—the contrast between how the language would be described and taught today versus how it came about.

Originally, Excel code was written as macros (subroutines) on Excel worksheets, which came to be known as “macro sheets.” In Excel 5.0 these were moved behind the scenes to be “modules,” so a workbook consisted of multiple worksheets and modules. A special kind of worksheet was again introduced, called a “dialog sheet” by which a form could be created with

⁵ See my detailed discussion of Dewey's theory in *Situated Cognition*, Chapters 4 and 9.

buttons (e.g., “OK” and “Cancel”) and other means for user input (e.g., check boxes and text entry). In the next version of Excel, dialog sheets were also moved behind the scenes, so that there were now two kinds of modules: Code and User Forms. Finally (?), with the introduction of object-oriented programming into Visual Basic, a third kind of module, a “class,” was introduced (in fact, a code module whose variables correspond to class properties and whose subroutines correspond to class methods).

The reformulation of a dialog sheet into a “module” in later versions of Excel is an example of a conceptual change. Something specific in Excel, a kind of worksheet, became part of the programming environment available to other Microsoft Office products (Word, Powerpoint). The conceptual change removed features incidental to Excel to create a programming construct that actually has more features. The conceptual change, the abstraction of “dialog sheet” into “user form” is not a process of “focusing on the essence,” something inherent or contained within “dialog sheet” but creating a more general, encompassing formal construct of structures and processes. In this conceptual reformulation, the meaning of “module” changed, too.

These abstractions – user form, code module, class module – as Casirer/van Oers says, are “not the recognition of new, previously unknown general characteristics of the objects involved” (p. 7) but attributes and objects added to our thinking, “a process of enrichment.”

The development of the Visual Basic framework involves sorting and *re-interpreting* formal constructs, such that *previously known entities are transformed*. Amazingly enough, two general entities—dialog sheet and module—are changed such that one becomes a *type*; a user form is now a kind of module. The original notion of module, for code, is now just a subtype. (Charts were then added as yet another kind of worksheet and have not yet been subsumed as being modules, perhaps because they are created and manipulated directly by Excel’s end-users and not only programmers.) Note that all of the terms mentioned here are abstractions and formal in character. They are inherently part of a programming domain, specific to Visual Basic, and yet independent of any domain of application (e.g. finance, weather forecasting). They are “domain general” abstractions.

Concepts vs. representational artifacts

Aside from perceptual-motor “direct” grasping of the world⁶, all thoughts are conceptual and hence “abstract.”

“The thought elements in perception and the perceptual elements in thought are complementary. They make human cognition a unitary process, which leads without break from the elementary acquisition of sensory information to the most generic theoretical ideas. The essential trait of this unitary cognitive process is that at every level it involves abstraction. Therefore the nature and meaning of abstraction

⁶ I associate the notion of “grasping” with Gibson and ecological psychology in general (Clancey, 1997, Chapter 11).

must be examined with care.” (Arnheim, 1969, “What Abstraction is Not,” p. 153)

Concepts, which I take as embracing all modalities of thought—visual, gestural, auditory, and verbal—are more or less abstract relative to each other, depending on how they are composed. In the realm of *external representations* (descriptions, diagrams, pictures), abstraction is also relative. Here the notion of abstraction as “taking away” developed, as we say that a sketch is more abstract than a photograph because it leaves out details.

Understanding abstraction is complicated because these two realms—concepts and representational artifacts—are of different character, yet they are co-determined (mutually influenced). A serious error occurs in equating representational artifacts (e.g., rules in an expert system, words in a dictionary) with conceptualizations. There is a relation of course, but it is not correspondence, let alone functional equivalence (as a strict or “strong” interpretation of knowledge-based systems would attest). Perhaps the most telling point is that verbal conceptualization, which obviously has a strong relation to verbal behavior (what we say), is not the only kind of conceptualization in human consciousness. Thus, equating verbal representations (words, statements, notations) with conceptualizations is a fundamental error in cognitive modeling.

In particular, to have a conception is not necessarily to have a “criterion from which one can decide” (p. 5) something, in the manner that a written law allows deciding on the validity of late overseas ballots during an election. We must be careful in using words like “decide,” “rule,” and “criterion” that we distinguish conceptualizing from verbal argumentation. Here are some statements in van Oers’ paper where confusion is possible:

- “Attaching a meaning” (p. 5) must be understood as creating a meaning. For example, understanding why a dialog sheet should be implemented as a module—a conceptualization that occurred in the revision of the Excel language—involved *seeing* a dialog sheet as *being* a module, thus changing both the meaning of “dialog sheet” and “module.” “Seeing as” is the conceptual aspect of the subsequent articulated definitions. The interpretive meaning remains open and tacit, not necessarily exhausted by the definition, because it has nonverbal aspects.⁷
- “The general is not the end result of an abstraction process; rather, some general principle is always the beginning. An abstract concept is not so much a reproduction of reality, but actually establishes a **point of view** that guides our thinking” (p. 6, emphasis in original). This non-correspondence theory of the relation of knowledge and reality is perhaps better stated as a “both-and” relation: The general is both the beginning and the “end.” The previously understood (general) notion is changed by the abstraction process.
- Similarly, saying that “redness was not an inherent quality of the objects themselves” goes too far. Both physical processes in the world and neu-

⁷ The contrast between tacit “open texture” of meaning and articulated definitions is emphasized by Wittgenstein (1953; Kolenda, 1964).

ral processes are involved (see Clancey, 1997, Chapter 4). On the other hand, “purely mental constructions that are created by abstraction” (p. 7) may lack any physical referent, but still have a perceptual referent *in the imagination* (Dewey, 1938a).

Contextualization

In saying “context...is itself a theoretical system” we mean that it is a *conceptual* system. When van Oers says that understanding “represents this original situation in an unspecified way” (p. 9), (I trust) he means that the conception of the activity has nonverbal aspects.

A simple example is the activity of answering a telephone. When a phone rings, we (usually) conceive it as requiring us to answer. We begin a sequence of acknowledging our presence by saying “hello,” and are then engaged in the sociocultural activity of answering the phone. But this activity is undeveloped, we are not goal-directed in the sense in which cognitive scientists characterize human problem solving (Newell & Simon, 1972). Yet, we are abstracting, by which we are making sense of the concrete situation and constructing a framework by which our ongoing behavior is organized (hang up immediately? Say, “sorry, wrong number”? Say, “Can I call you back in five minutes”? Carry out a long conversation?). Thus, in conceptualizing what the call is about and organizing a response we are contextualizing the concrete situation—such that what the situation is and how we respond to it are arising together in our understanding, just as our understanding and actions are influencing each other (again, see discussion of Dewey in Clancey, 1997).

“Contextualization” is “making sense of a concrete situation” by grasping/viewing it (van Oers says “translating it”) with respect to or as “a particular (but still abstract) sociocultural activity from which new actions can emerge” (p. 9). The essential notion is *conceptualizing* the activity, *proceeding from concrete aspects* of a situation (“multifaceted, but organized whole to which the human action is intrinsically related”).

A physical-interactive aspect is inherent (“intrinsic”) and nonverbally associated with the phone answering activity conceptualization. Without having to verbalize what we are doing, we simply get up and answer the phone. By the phone answering conceptualization, perceptual-motor actions, objects, attributes, operations, roles, names, causal stories, notations, etc. come to be coordinated. The phone answering activity is at first “‘abstract’ because it is still undeveloped, that is, not yet particularized in specific goal-directed, tool-mediated human actions” (p. 10).

Schön (1979) would say “framing” where van Oers says “contextualize”, e.g., making a gate can be framed as “going to a shop” or “carpentry” (p. 11) – which are activity conceptualizations involving different objects and operations.

What constitutes a *task* in studies of human problem solving arises within a conceptualization of an activity. Contextualizing a situation precedes the definition of a goal and then means (action possibilities). As Lave (1988) emphasized, the framing of what constitutes the situation and how it might be approached is left out of studies of human problem solving that give

subjects problems to be solved. Crucially, the role of abstract thinking is not fully appreciated in traditional cognitive modeling. Abstract thinking is not only a process of solving formal problems in isolation (e.g., cryptograms, chess, physics)—reducing value-laden situations to puzzles—but especially the role of abstract thinking is to transform a concrete situation into *possible goals* and actions. By this analysis, mental models become instruments for organizing a situation, of particularizing it in a certain way, so it can be operated upon in the manner Newell and Simon showed.

“Abstract thinking is always itself contextualised in a system that regulates this mental activity of abstract thinking” (p. 10). Abstract thinking is “not a detached way of acting” but a cultural activity that constructs an understanding of reality (van Oers says “appropriation of reality” p. 10). Similarly, how we respond to the phone call depends not only on the particulars of the call but the ways in which we are then conceiving of our activity. Are we eating? In a meeting? Driving a car?

To pick up on the Visual Basic example again, developing a user interface may be framed as creating a “dialog box” or as creating a “wizard” (a series of dialog windows with similar style and conventional user options for continuing or ending the interaction). The wizard concept frames the design problem in a particular way. It is an object by which the elements in a concrete situation can be put together (p. 14); the wizard is an instrument for transforming a situation into a *definable problem*, an arrangement of unknowns that can be resolved.

Contextualization is basically a *dialectic* process, often using discursive argumentation as a means of negotiating new meanings of categories and particulars (contrast with van Oers, p. 12, “basically a discursive (argumentative) process”). *Descriptions are instruments* for developing an understanding (Dewey, 1938a, b). The verbal interpretation is a means, not an end, for accomplishing meaningful interaction, for conceiving of an activity. Nevertheless, it is appropriate to say that problem solving activities in school are basically discursive because they inherently involve descriptions of objects, events, and goals (“conversations between teacher and the pupils” p. 19). My point is that problem solving is just one kind of activity in which people engage, which is to say it is just one target framework for contextualizing a situation. A religious ceremony is another framework people use for ascending to the concrete.

Types of activity foci

“Focusing on particular aspects of a situation” (p. 12) may involve *constructing* not merely selecting foci. For example, when Schön’s (1979) painters were trying to understand a malfunctioning synthetic paint brush, they worked with the gloppy paint on the wall and focused on the flow, such that they came to see the space between bristles as *channels*. Thus, perceptual features may be created dialectically in physical interaction, changing what perceptually constitutes the concrete situation. Just as abstracting is not removing details, contextualizing may create new perceptual details (contrast with “focusing” and “attending”). The perceptual process creates features, rather than merely selecting or pointing to them (Clancey, 1997, Chapter 9).

The relation of perception, action, and conception is structurally *dialectic* (mutually influencing) at the neural level, and “discursive” in how we have culturally developed linguistic behavior for negotiating meanings (perhaps most obvious in court rooms). In analyzing a classroom van Oers is right that the resultant process is “basically discursive” (talk dominates the organization of activity).

Figure 1 illustrates the three aspects of focusing in abstract thinking:

- “Pronounced focus” (p. 14) is a specified (named) activity
- “Intended focus” (p. 15) is a cultural (normative) activity
- “Attended focus” (p. 15) is a task (functional) activity

Crucially, the intended focus is *experiential*, an activity-understanding that the participants have previously used to organize their behavior. The sequence of pronounced-intended-attended focus can be understood as nested. In van Oers’ example, the intended shoe shop activity becomes constrained by the articulation of measuring, by which the shoe shop play activity is constrained to comparing and counting shoes. Viewing shoes as types and quantities involves another kind of normative, experiential activity, with operations of classifying, counting, and drawing (abstract thinking proper).

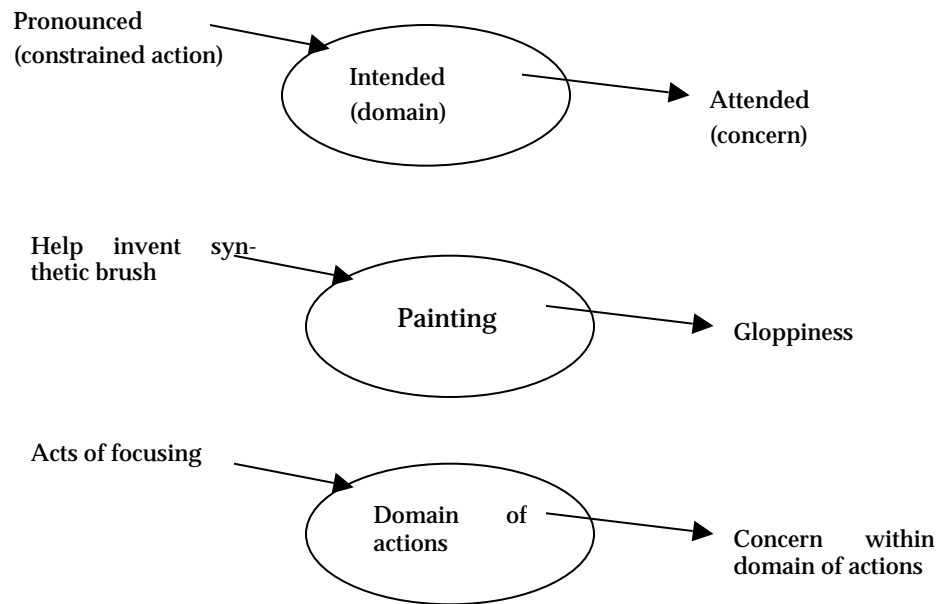


Figure 1. Types of foci in abstraction activities, with example from Schön's paint-brush inventors

Not all aspects of the intended focus are verbal, just as not all aspects of conceptual, experiential knowledge have been articulated or could in principle

be inventoried. Specifically, the pronounced foci in the example are: corner, shoe, pile, and histogram. The attended foci (domains of activity) are: measuring, classifying, piling, counting, drawing, and counting (on paper). The example reveals abstraction as a mediator, an instrument, a model for transforming attention, ideas, and materials.

Similarly, my second Excel programming example can be analyzed by this framework. The pronounced focus pointed to a general difficulty beginners had using a Visual Basic program. The intended focus could be stated as “How should the interface be fixed?” “To deal with the concrete” (p. 23) I employed the abstraction of a wizard (an attended focus), from which I created multiple concrete objects, such as a “financial statement panel.”

Now I was engaged in a problem-solving activity—how to manage the operations of a given wizard panel. What were the user operations and their effects? This was the new pronounced focus as I created buttons and menus for adding, deleting, and renaming statements. The activity of constructing a wizard transformed into a subsumed activity of creating an internal model of changes requested by the user—a model that future programmers would be able to understand and modify if necessary. The domain of “design for ease of use” shifted to the domain of “internal code organization and ease of modification.” The intended focus shifted from future users to future programmers.

The organizing construct that came into play (the new attended focus) was “a class module”—a kind of data structure. Just as the students in van Oers’ example moved from an externally visible and physical interactive world of “playing shoe store” in a corner to an essentially mathematical and computational world, my activity shifted from a visible user-interface world (organized by the wizard metaphor) to a behind-the-scenes programming implementation world (organized by the change data structure model).

The shift from designing an interface to designing a program is “a process of continuously progressive recontextualizing” (p. 23). This “nested series of foci” (p. 22) constitutes “an implicate order of abstractions,” such that “a variety of elements of the concrete activity setting can be seen as related.” Through the progressive nesting, ultimately, the “user interface difficulty” becomes related to a wizard panel, a drop down menu, an internal data structure, and a special class module. Different view points, levels of design, are coordinated and managed over time, such that the complexity of relating hundreds of external and internal objects is cognitively tractable and a coherent, successful artifact results.

By “implicate” (p. 23), van Oers means that the conceptualization of activity tacitly organizes our perception and action, which is sequential, compositional, and generative (Schön, 1979). Crucially, the abstraction is not merely (or at all) “level-raising” but grounded in particulars, thus the phrase “ascend to the concrete.”

A key point (cf. O&R’s analysis) is that both the wizard and internal change (data structure) models are abstractions. Both are means of organizing a situation, indeed creating what the person can view as being a situation, by locating the momentary interaction within an experiential, normative context (respectively, user-interface design and code design). This attended focus (a name for “what I am doing now,” cf. Clancey, 1999) becomes a source for defined goals, potential actions, and evaluative methods. The domain of action (the intended focus) comprises these, plus perceptual fea-

tures, language, scenarios, resolution schemas, previous cases, roles, etc.—a complex stance, a way of being, indeed much more than a “view point” (p. 19).

Being in school—primary abstractions

We must be careful not to go too far in idealizing the students’ role. In particular, we must contrast the experience of the student to the inventor or scientist whose work is being (directly or indirectly) presented to the student. Some abstractions that have been named over time and formalized into methods are indeed “held out by the teacher.” For example, I did not invent the idea of a wizard; I read about it in a programming manual. The ideas of “graphical user interface” (GUI), wizard, and “change model” (data structure) were taught to me. Davydov (1996) calls these “primary abstractions.” In van Oers’ example, graphing has a similar special, epistemological status, an object that the society wishes the student to appropriate. Students do not have to invent graphing, as I did not have to invent the idea of a wizard.

Van Oers emphasizes that the real question is how students come to such abstractions by virtue of starting with the concrete, and how this occurs discursively in a process of progressive focusing (p. 24). I would strike a perhaps somewhat different balance by posing the question of how students come to *value* such abstractions, while distinguishing between the *articulated theory* (e.g., the documented nature of a wizard) and the *conceptualization* the students will individually form. The students have a complex of abstractions, names, and artifacts to coordinate. They do not simply “come to this ‘primary abstraction’ in the first place starting from the concrete.” They are starting from a *named target* (e.g., “histogram”). In learning carpentry, the teacher holds out a variety of tools, many of which you do not know how use at first. You do not have to invent the tools or even invent how they are to be used. Some you will be shown in a detailed way: To do this, use this tool thusly. Primary abstractions such as a wizard or a histogram are tools. We are surely *given* them as much as we “come to them” in our understanding.

Further, for van Oers’ example, I would stress the larger conceptualization of activity of “being in school” by which students come to attend to teachers, follow their direction, and expect to be taught something that goes beyond the particular games and artifacts they employ (playing with shoes in the corner is not just about shoes). In pointing to this broader intended focus, I am entirely within van Oers’ theory, but emphasizing that the students are playing a special role, and the students’ knowing this is presumably important for meeting the schools’ criteria for success (Lave, 1988).

Conclusions about van Oers' paper

To summarize van Oers' point, "ascending to an understanding of the concrete" (p. 9) is not just instantiating a generalization in an example or real situation. Rather a conception of the activity transforms a concrete situation by providing context for acting ("meanings that help to transform this concrete situation in an activity that will function as a context for his actions," p. 9). The abstraction constrains our behavior (making, saying, doing) by providing a *tool* (instrument, guide, standard, method) to accomplish tasks that become defined within the conception of activity. For example, using the wizard conception defined a layout of multiple dialog boxes, a configuration of buttons, a right for the user to cancel or finish the sequence at any time, etc.

There are two kinds of abstraction discussed here, perhaps easily confused by the repeated term "focus." The first kind is *experiential* and largely tacit. For example, the first focus in the programming example (an intended focus) involved usability and interface design; these concepts are not fully defined or reducible in principle to a single set of theories. The second kind of abstraction is *formal* and tool-like. It involves invariant standards or methods by which some activity may be organized. These are verbally defined structures/processes, often articulated as rules (mathematical or heuristic). The second focus in the programming example (an attended focus) involved a wizard, whose layout and operation is given in manuals. The meaning and intent of wizards is still open to interpretation and use, but the construct is not primarily experiential; it is a deliberately designed entity. The histogram is an analogous abstraction in van Oers' example.

A reader wishing to better understand the value and indeed revolutionary import of van Oers' presentation should read the work of Dewey, Schön, and Bamberger. This work has striking similarity to van Oers' concepts, but he does not cite them. For discussion of the "conversational process of focusing" (p. 4) and "thinking as a conversational process" (van Oers cites Billigs p. 12) see Bamberger and Schön (1977, 1983). For an example of allowing "pupils to present their own renderings of the situation (the activity) and encourage them to reveal their initial symbolic representations" (p. 24), see Bamberger's (1991) use of bells and notations in teaching musical theory. For elaboration of the notion that "contexts are not given as such" (p. 22) see Lave (1988).

Abstraction as composed conceptualisations and interactions: Dreyfus, Hershkowitz, and Schwarz

I will use the paper by Dreyfus, Hershkowitz, and Schwarz (DH&S) as an opportunity to bring out distinctions about mental processes, verbalized mathematical theory, and artifacts. In doing this, I am not so much criticizing their work, as using their presentation as a springboard for a more theoretical discussion.

To begin, DH&S "define abstraction as an activity of vertically reorganizing previously constructed mathematical knowledge into a new structure"

(p. 1). The question I asked in reading O&R's paper arises here: Is abstraction only relevant to mathematical knowledge? Shouldn't the definition (and paper title) say "mathematical abstraction" to make clear that abstraction in other domains (science, medicine, finance, sports, religion, etc.) is not being considered? DH&S later point out (p. 5) that they believe their definition and model to be general. This generality may have several dimensions. I suggest that we distinguish mathematics as a *discursive domain* (e.g., algebra problems in school), mathematics as a *tool* (e.g., using topological relations, such as layering, to understand geological processes on Mars), and mathematics as an *inherent relation in human conceptualisation* (e.g., sequencing, composition, substitution; Clancey, 1999).

The focus of DH&S's paper is clearly on peer interaction, but is intended to elucidate the nature of abstraction as an activity. DH&S "investigate the distribution of the process of abstraction in the context of peer interaction" (p. 1). Thus, the paper's hypothesis is that abstraction is a visible phenomenon and occurs by means of interactive social (not only mental) processes. Furthermore, peer collaboration provides a context that presumably influences individual experience and contributions.

The particular problem given to the students involves inferring the pattern (algebraic formula) in a spreadsheet computation. The students are required to learn two things: the computation rule (the extended distributive law) and the idea of using algebra to "justify the correctness of a numerical property" (p. 2) (e.g., the spreadsheet pattern). DH&S apply a previously validated model to analyze the students' interaction, showing when abstraction is occurring and how it is nested. Differences between student pairs allows DH&S "to identify interaction patterns that are compatible with abstraction processes" (p. 2).

DH&S embrace activity theory because it relates goals to more general "motives," as emphasized in van Oers' analysis. But as brought out in van Oers' discussion, Davydov's dialectic theory of abstraction (p. 4) is not the same as Ohlsson's notion of "assembly of existing ideas" (DH&S p. 3). The "vertical" relation of abstractions is much more than "accessibility" or "putting together" mathematical "elements" (cf. my discussion of abstraction in developing the "module" construct in Visual Basic). The ideas of "attuning practices" and "restructuring knowledge" are more telling.

Saying that "abstraction is not an objective, universal process" (p. 3) is right in referring to abstraction as an activity, an interactive process, involving tools, language, and procedures. However, I believe it is just as important to remember that abstraction is also a neural process, and almost certainly does have universal characteristics in people as patterns in grammars attest (Clancey, 1999).

Sorting out knowledge, theories, and artifacts

Central to DH&S' analysis is the notion of "epistemic actions": *Constructing* ("assembling knowledge artefacts to produce a new structure" p. 3), *Recognising* (relating an object to a known mathematical structure via analogy or specialization), and *Building-With* ("combining existing artefacts...to satisfy a goal" p. 4).

In reading DH&S' definitions, many questions arise. Is a "knowledge artefact" an object in the world (what is what artifacts normally are) or a conceptualization? Put another way, are mathematical structures, which are properties of visible artifacts (e.g., spreadsheets), the same as mathematical concepts? Is *Constructing* the same as developing a new concept (e.g., the idea of "module")? How does assembling (mathematical) *structures* relate to the "overall motives" that activity theory emphasizes? How are structures and motives related during learning? Are motives *structures*? Does mathematical knowledge only consist of structures? Are the "attuned practices" stored as structures in memory? Or is RBC analysis only focusing on mathematical terms and laws ("structures"), which are figurative parts of a mathematical artifact (analogous to design layout problems, as in O&R's experiment)? Or does "structure" here refer more generally to any organizational scheme, a concept for sequencing and composing activity?

I become very uncomfortable when knowledge (abstractions), descriptions (mathematical terms and laws), and objects in the world (artifacts) are loosely equated. So I began reading the paper quite confused about the terms "structure" and "entity." When we say "a need for a new structure" (p. 4) do we mean a conceptualization? A formal description within a written mathematical theory (such as a new definition)? Or a construct within a mathematically based artifact (e.g., a formula in a spreadsheet cell)? If all of these, then how are they related? (Of course, this is a central question in cognitive science.)

On the other hand, the purpose of the paper is quite clear: RBC analysis aims to characterize behaviors in peer interactions, by labeling actions as having different epistemic bases and effects, and then using the analysis as a way of explaining collaborative results.

Early on, DH&S distinguish between mechanical or rote learning and learning abstractions (p. 5). This distinction appears relative—a given concept is not coordinated with certain actions, so the concept is not involved. DH&S present the good example of solving a puzzle, involving much *Recognising and Building-With*, but no *Constructing* (of a new method/strategy). But the distinction is also potentially misleading because no human performance is strictly mechanical. A dialectic theory of abstraction assumes that conceptual change is always occurring, even when it is not visible in behavior. At the simplest level, in learning to subtract four digit numbers (p. 5), behavior patterns are being reinforced (habit sequences); this changes the procedural process, but not its meaning. But at a higher, *conceptual* level, the conception of the activity incorporates repetition and cycling as an operator (e.g., "when the teacher gives us problems that look like this, this is what we should do"). Also, what puzzle solving means to the person is perhaps changed. For example, when recently playing a "15" puzzle on a Personal Digital Assistant, I was amazed at how mechanically and quickly I moved; but I also realized (a reflection on the activity) that this was not how I wished to spend my time. In contrast, when playing solitaire on the same PDA a few moments later (the first time I played the game in perhaps 30 years), I realized how pleasurable it was, and played many times in the ensuing weeks.

To recap: From the perspective of mathematical theory, one may precisely label descriptions as being abstractions (or not) relative to each other. But we cannot so quickly apply such talk (“not an abstraction” p. 5) to mental processes. At least, some acknowledgement of the shift in domains is required, with an explicit claim about how conceptual processes, mathematical theory, and artifacts are related. DH&S’ important point, which is consistent with my caution, is that we must be careful in analyzing performances with respect to (mathematical) concepts, even when formal problem solving is involved (as in solving some puzzles), Procedural proficiency, conceptual understanding of a mathematical theory, and the mathematical theory itself (as an artifact) should not be equated.

Comparison of analysis across the three papers

The experiments described in DH&S’ paper show that “student pairs constructed more abstract representations of the [spreadsheet] problems than individual students” (p. 5). DH&S emphasize that students at a certain level of development tend to rely on “inductively checking a number of cases” (p. 6) rather than looking for an algebraic formulation and proving that it explains an input-output pattern. The focus on the students’ “psychological need” (p. 6) and use of spreadsheets to construct data and formula relations are particularly interesting. Analyzing the collaborative interactions is difficult, but DH&S adapt Resnick et al.’s (1993) formulation to good effect. For example, one can see in the diagram for Segment 4 that the students are not relating to or continuing ideas (p. 17). And strikingly, “individual elaborations occurred successively in segment 3 but simultaneously in segment 5” (p. 17). Building on this notation, the subsequent RBC analysis shows that “the constructing action may be quite long and contain shorter segments, which themselves are constructing actions” (p. 19).

At this points, one realizes that such an elaborate analysis would have been useful, at least in part, to better understand peer interaction during both O&R’s and van Oers’ experiments. Furthermore, if we try to apply van Oers’ analysis of types of foci we find strikingly different activities and hence contexts for learning. First, in O&R’s archaeology puzzle, we find that the text creates a background for the problem, but the background is not an activity (the students do not pretend to be either archaeologists or Egyptian priests, etc.). Nor are any facts or relations in the background story relevant to solving the problem (!). Second, in DH&S’ spreadsheet puzzle, there is no non-mathematical background provided at all. The pronounced activity is to learn algebra in an Excel learning environment (p. 11). In contrast, playing in van Oers’ shoe store corner constitutes an activity that engages the students in an experiential conception. Within their personal understanding of objects, descriptions, actions, and roles, which frames and organizes their behavior, they are guided to compare and then measure shoes, and then to use mathematics to count and graph relations. All of the experiments are contextualized in a literal sense (archaeology, shoe stores, spreadsheets), but only van Oers relates mathematics to *an experienced background* from which the students may draw ideas to organize their understanding.

DH&S' interpretation of activity theory is more narrow than van Oers', viewing abstraction as an activity that with observable aspects during peer interactions. The notation reifies interactions, from which properties may be observed and articulated (e.g., "successive" vs. "simultaneous" collaboration) and related to outcomes. In fact, we could use DH&S' paper itself as an example of abstraction at work.

In van Oers' terms, the pronounced activity in DH&S' project is to examine an experiment with spreadsheets; the intended focus is the domain of peer interactions; the attended focus is the categorization and relation of utterances. Subsequent foci are segmented transcripts (which are contrasted and compared), utterance categories ("conversational moves" p. 17), and then diagrams and tables of utterance categories. The analysis is repeated to apply the RBC abstraction, now revealing a new property, nesting of epistemic actions (p. 18). With RBC, the intended domain shifts from categorizing and relating utterances to relating kinds of epistemic actions. The essential claim is that the conception of the extended distributed law (designated C2) occurs within the broader conception that "justification may be expressed algebraically" (p. 21). Crucially, the relation is dialectic, such that the meaning of the extended distributed law and the idea of making an algebraic argument (the motive designated C1) are developing together (p. 21 of DH&S' paper bears repeated reading to understand this dialectic relation). The construction C2 gives meaning to C1, while C1 is guiding the invention of C2. DH&S characterize this relation as "a deep holistic construction" (p. 21), such that the nested understanding that results goes beyond what was required to explain the Diagonal Product Property (DPP). They conclude that "C1 is an activity of vertically reorganizing previously constructed mathematical knowledge into a mathematical structure, which fits our definition of abstraction" (p. 21).

Conceptualization of mathematical theory

The reader is probably not surprised that I very much like DH&S' example and analysis. We can now revisit my opening questions. We see that DH&S refer to C1 as a construction, an activity, a level, something expressed by (C2) actions, a nesting, a segment (p. 20), a constructing action (p. 19). All of this suggests that in my terms, C1 is a conceptualization of an activity (a mental construct). C1 has characteristics I presented in the discussion of van Oers—it frames perception and action over time, it relates talk, ways of looking, ways of manipulating materials, student roles, etc.. As DH&S show, C1 can be formally associated with patterns of utterance categories, as well as nesting of other conceptualizations (notably C2).

To answer my other question, DH&S equate mathematical structures with constructions, so they are conceptualizations, too (and not just articulated terms or laws). Specifically, C2 is the conceptualization of the extended distributed law, which comprises how it is used, talked about, expressed in writing, applied to the particular DPP problem, associated with spreadsheet puzzle-solving, etc. Crucially, the notion of "abstraction" is applied to C1 and C2, not just the written formulation of the distributed law (contrast with saying " $X(A + B) = XA + XB$ is a mathematical abstraction").

However, we must remember that “Constructing” is being used in a specific way here. When DH&S characterize Ra’s justification of a property “as building-with rather than as constructing” (p. 21), they mean that Ra doesn’t form any “new” mathematical ideas, not that his existing understanding is not changed from this experience.

At another level, we might find further clues about conceptualization from the mistakes the students make. For example, Ha & Ne together make the interesting mistake of rewriting $(X + 6)(X + 2)$ as “6X times X and 6X times 2” (p. 25).. When Ha says in H135, “This is 6X,” is this perceptually driven by association with already learned pattern, $aX(bX + c)$, so $X+6$ is manipulated as 6X? Importantly, this conceptual jump of recognizing $X + 6$ as a single unit, although mistaken, “gave the girls the necessary point of view to realize the applicability of the [distributive] law” (p. 30). Although a kind of visual illusion, seeing $X + 6$ as a unit brought the right construct into play, and through a subsequent reformulation (H139 “No, because this is X plus 6, this is not 6X, it’s different” p. 25) they were able to construct the correct (extended) distributive law. Thus, recognizing may be inappropriate in specific terms, but provide materials that can be manipulated to construct a new understanding (a generalization of the previous understanding and use of the distributed law). I point this out because such a move is not just a mistake, but as a too-hasty simplification was pivotal in developing the correct mathematical formulation. The incremental nature of learning, often involving failures, has been emphasized in models of learning (e.g., Winston, 1970). I am especially interested in how a perceptual association (“seeing as”) brought an existing conceptualization into play that was literally enacted, and then had to be generalized in order to be correctly applied. How general is this particular pattern in human learning? In scientific theory formation?

The structure of peer interactions

DH&S conclude with many interesting observations about the structure of the peer interactions. Most interesting is the question of whether the “common RBC flow” (p. 35) can be characterized as being produced by a single “social being.”⁸ A key point is that if we say that the “students participate in the same activity,” then “in the course of their actions they share the same motive” (p. 37). However, with respect to C2, this is more than being “conscious of the overall goal imposed by the task” (p. 37). We must find evidence that they are co-constructing C2, which is not given or imposed, but discovered. Without engaging in the nested subactivity, C2, the broader C1 activity would lose coherence. The nature of C1’s coherence appears stronger than DH&S claim.

The subsequent discussion of symmetry and alignment (Section 5.2, p. 38) is very useful. The conclusion is that these concepts are fluid; the character of the interaction changes as the activity develops. For example, Yo&Ra’s interaction, although coherent, is symmetric, but not aligned be-

⁸ Besides the cited papers concerning interaction in conjunction with learning, DH&S might have mentioned the work of Greeno and Hall (1992; Hall, 1995), Goldman and McDermott (1994; MMAP), or Greeno and Roschelle (1987).

cause “they do not collaborate during the construction” (p. 38). A larger implication is that we should be careful not to characterize every successful interaction as being a “collaboration,” or at least recognize the varying degree of symmetry and alignment during an episode sequence (“segment”).

As an example of social-cognitive causal relations in peer interactions, consider the important observation that Ne’s pushing Ha to explain forces her to “communicate her ideas in a new perspective” (p. 41). Thus social interaction is causing an individual to rearticulate and hence reconfigure a conceptual relation—adding aspects, emphasizing what’s important, relating to other examples, etc. Thus the social interaction framework is a necessary causal aspect in the story of Ha’s conceptual change. Similarly, “paying attention to mutual results they reach” (p. 42) indicates that the cognitive process is not merely individual, but a social, interactive process.

Rather than saying that the outcome is to “vertically reorganize knowledge into new structures” (p. 42), I would say “compositionally reorganize conceptualizations” (referring to the students’ activities of spreadsheet problem solving, using algebra to justify patterns, and using the extended distributive law). The compositional nature (more than vertical) provides the overarching activity organization, which sustains coherence in even asynchronous and parallel interactions. Thus, the understanding of the activity, the conceptual motive, makes the interaction coherent. This “activity-based, dynamic meaning of abstraction” (p. 36) is substantially different from the traditional view of human behavior as driven by and fully understandable in terms of goals, inferences, and mental models. The study of peer interaction reveals that much more is required to sustain and organize learning. Individual understanding of the group’s shared context as being engaged *in a certain kind of problem-solving activity* is essential.

Summary conclusions

My objection to O&R’s paper is not that the claim is wrong; obviously formal insights influence topological problem solving. Rather the paper is too limited as an appraisal of scientific practice and unfairly interprets knowledge-based research. The broader focus I advocate is nicely illustrated by van Oers’ paper, where discovery is placed in an activity context, which progressively focuses on different perceptual properties and formal means of characterizing the objects being manipulated. I illustrated the same point in the Excel examples (which I used to verify that I understood and agreed with van Oers’ analysis—illustrating how van Oers’ abstractions enable me to reformulate my own experience and vice versa). Finally, DH&S use a very similar activity-based framework to develop a notation for describing peer interaction, such that the development of understanding can be understood with respect to specific behaviors and utterances. I suggested that both van Oers and DH&S’ analysis could be improved by including “being in school” as a higher-order conceptualization of the context that organizes the students’ activity.

Throughout, I have emphasized that general conceptual processes are at work, and that a dichotomy between mathematical and other abstractions is

not warranted or necessary. The repeated shortcoming of educational theory has been to give primacy to the formal aspects of our own analyses and the domains we operate within, leading us to assume that fixed patterns, definitions, and rules generate behavior (Clancey, 1997). Even in DH&S' analysis, we might question whether the learning method "proof via algorithm" is best characterized as a mathematical abstraction being directly applied or is the manifestation of something more general, such as using a model and making inferences (*using algebra*)?

It is time in cognitive science to developmentally relate—in causal stories of physical/social activity—conceptual processes, descriptions/notations, and artifact manipulation (such as Watson and Crick's metal model). The common theme in these papers is that something other than specific domain facts and rules organizes changes in human behavior. This organizing influence can be characterized as an abstraction, the intended focus, the conceptualization of activity, the understanding of context. The key theoretical notion is not in viewing behavior as derived from a repertoire of stored knowledge (whether mathematical or not), but in analyzing behavior as *progressing from previous organizations of behavior*: Learning is a transformative process and occurs as and through action itself. The question is not, why do the students do X, in isolation. But rather, why and how did the students shift from doing/saying X to doing/saying Y? The notion of abstraction as transforming behavior starts to bring together the mental, descriptive, and interactive aspects of cognition. To go further, I suggest that we view knowledge and learning in terms of conceptual coordination—a neurally based process in which perceptual-motor systems and conceptual systems in different modalities are blended, sequenced, and composed.

References

- Arnheim, R. 1969. *Visual Thinking*. Berkeley: University of California Press.
- Bamberger, J. 1991. *The mind behind the musical ear*. Cambridge, MA: Harvard University Press.
- Bamberger, J., and Schön, D. A. 1977. The figural formal transaction: A parable of generative metaphor. Unpublished technical report. Division for Study and Research in Education.
- Bamberger, J., and Schön, D. A. 1983. Learning as reflective conversation with materials: Notes from work in progress. *Art Education*, (March): 68-73.
- Bickhard, M. H., and Richie, D. M. 1983. *On the Nature of Representation: A Case Study of James Gibson's Theory of Perception*. New York: Praeger Publishers.
- Bickhard, M. H., and Terveen, L. 1995. *Foundational Issues in Artificial Intelligence and Cognitive Science: Impasse and Solution*. Amsterdam: Elsevier Science Publishers.
- Buchanan, B.G. 1992. DENDRAL. In Stuart C. Shapiro (ed), *Encyclopedia of Artificial Intelligence Second Edition*, p. 331. New York: John Wiley.
- Buchanan, B.G. and Wilkins, D.C. 1993. *Readings in Knowledge Acquisition and Learning*. San Mateo: Morgan Kaufmann.
- Clancey, W. J. 1989. Viewing knowledge bases as qualitative models. *IEEE Expert*, (Summer 1989), 9-23.
- Clancey, W. J. 1997. *Situated Cognition: On Human Knowledge and Computer Representations*. New York: Cambridge University Press.

- Clancey, W. J. 1999. *Conceptual Coordination: How the Mind Orders Experience in Time*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Davis, R. 1982. Teiresias: Applications of Meta-Level Knowledge. In Davis, R. and Lenat, D.B. 1982. *Knowledge-Based Systems in Artificial Intelligence*. New York: McGraw-Hill International Book Company, pp. 229-490.
- Dewey, J. [1938a] 1981. The criteria of experience. In *Experience and Education*, pp. 23-52. New York: Macmillan Company.
- Dewey, J. 1938b. *Logic: The Theory of Inquiry*. New York: Henry Holt & Company.
- Feigenbaum, E.A., Buchanan, B.G., and Lederberg, J. 1971. On generality and problem solving: A case study using the DENDRAL program. In B. Meltzer and D. Michie (eds), *Machine Intelligence 6*, pp. 165-190. New York: American Elsevier.
- Feltovich, P.J., Spiro, R.J., and Coulson, R.L. 1989. The nature of conceptual understanding in biomedicine: The deep structure of complex ideas and the development of misconceptions. In D.A. Evans and V.L. Patel (eds), *Cognitive Science in Medicine*. pp. 113-172. Cambridge, MA: Bradford.
- Gentner, D. and Stevens, A.L. 1983. *Mental Models*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Hall, R. 1995. Realism(s) for learning algebra. In C. B. LaCampagne, W. Blair, and J. Kaput (Ed.), *The Algebra Initiative Colloquium*, (Vol. 2, pp. 33-51). Washington D.C.: U.S. Department of Education, Office of Educational Research and Improvement.
- Goldman, S. and McDermott, R. 1994. Is it math yet? Kids' communication and math in framework-inspired classrooms. Paper presented at the American Educational Research Association annual meeting, New Orleans, LA.
- Greeno, J. 1994. The situativity of learning: Prospects for syntheses in theory and practice. Paper presented at the American Psychological Association, Los Angeles, CA.
- Greeno, J. and Hall, R. 1992. Participatory Design of Assessment Systems. Paper presented at the Inaugural Conference of the Center for the Study of Education and Child Development in a Diverse Society, University of California at Los Angeles.
- Kolenda, K. 1964. *The Freedom of Reason*. San Antonio, Tx: Principia Press of Trinity University.
- Lave, J. 1988. *Cognition in Practice*. Cambridge: Cambridge University Press.
- Lenat, D. 1982. AM: Discovery in Mathematics as Heuristic Search.. In Davis, R. and Lenat, D.B. 1982. *Knowledge-Based Systems in Artificial Intelligence*. New York: McGraw-Hill International Book Company, pp. 1-225.
- Lenat, D. and Feigenbaum, E.A. 1987. On the thresholds of knowledge. *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*. 2:1173-1182.
- Miller, A.I. [1984]. 1986. *Imagery in Scientific Thinking: Creating 20th Century Physics*. Cambridge: MIT Press.
- MMAP, url. Middle-school Mathematics through Applications Project. <http://mmap.wested.org/>
- Newell, A., and Simon, H. A. 1972. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall, Inc.
- Roschelle, J. and Greeno, J.G. 1987. *Mental models in expert physics reasoning*. ONR Report GK-2. University of California, Berkeley.
- Schank, R. C. 1980. Language and Memory. *Cognitive Science*, 4(3): 243-284.
- Schön, D. A. 1979. Generative metaphor: A perspective on problem-setting in social polity. In A. Ortony (ed.) *Metaphor and Thought*, pp. 254-283. Cambridge: Cambridge University Press.
- Vaughn, J. 2000. "Mars Global Surveyer: Performing daily, 687 days a year," *The Planetary Report*, XX(4), July/August.

Wenger, E. 1998. *Communities of Practice: Learning, meaning, and identity*. Cambridge, UK: Cambridge University Press.

Winston, P.H. [1970] 1975. Learning structural descriptions from examples (PhD Dissertation). In P.H. Winston (ed), *The Psychology of Computer Vision*, New York: McGraw-Hill, Chapter 5.

Wittgenstein, L. [1953] 1958. *Philosophical Investigations*. New York: The Macmillan Company.