3 Person-plus: a distributed view of thinking and learning

D. N. Perkins

Underlying psychology's multitude of investigations are a few broad and abiding questions. One of these – or perhaps two-in-one – is, How do thinking and learning happen? Efforts to reply include those of the behaviorist tradition, with its theory of classical and conditioned reflexes; now-classic cognitive theories, with their visions of problem spaces and schemata; and the more recent perspective of parallel distributed processing, with its holographic concept of how the mind captures and deploys information.

Whatever theory you pick, there is a notable and in some ways peculiar asymmetry between the posture taken toward the person and toward the physical arena in which the thinking and learning occur. Consider, for example, a student in a course on medieval history who has developed careful, well-organized notes about 1066 and all that. Most theories of learning would say that what the student has learned lies in his or her head. Whatever is in the notebook that is not also in the student's head is not part of what the student has learned.

Not, of course, that the notebook is deemed irrelevant. The student's effort to keep the notebook in a thoughtful, organized way will have resulted in better mental encoding of a good many of the ideas also represented in the notebook, including superior understanding and retention, because of the "elaborative processing" (e.g., Baddeley, 1982; Craik & Lockhart, 1972; Pressley, Wood, & Woloshyn, 1990). Nonetheless, the notebook itself would not usually count as a container of what the student had learned, even though it

represents considerable cognitive investment in a well-organized memory bank, and a bank that will pay dividends when, for instance, the student writes a term paper, drawing on this well-organized resource for ideas.

Another view of the matter is certainly possible. We could take as our unit of analysis not the student without resources in his or her surround – the *person-solo* – but the person plus surround, or *person-plus* for short, in this case the student plus the notebook. We could say that this person-plus system has learned something, and part of what the system has learned resides in the notebook rather than in the mind of the student. Moreover, in learning about 1066 and all that, this system thought hard, the notebook serving as a thinking scratch pad as well as a repository of conclusions.

What should we make of so odd a way of putting the matter? A reasonable attitude might be this: Certainly one can speak of the person plus surround as a compound system that thinks and learns – but is it particularly illuminating to do so? Do insights result that might otherwise pass us by?

Distributed cognition

The general view adopted here takes as one point of departure a conception of "distributed intelligence" articulated by Roy Pea (Chapter 2, this volume). Pea urges that we do well to reconsider human cognition as distributed beyond the compass of the organism proper in several ways: by involving other persons, relying on symbolic media, and exploiting the environment and artifacts.

The present view also reflects the distinction drawn by Salomon, Perkins, and Globerson (1991) between effects *with* and *of* information-processing technologies, effects *with* being amplifications of the user's cognitive powers during the use of a technology and effects *of* being cognitive spinoff effects that occur without the technology. The present discussion focuses on effects *with* – not only with high technologies but with what in general will be called the *physical* distribution of cognition – onto things such as computers, to be sure, but also pencil and paper or the simple tactic of leaving a folder in front of the door to remind yourself to take it to work. There will be some attention also to the *social* distribution of cognition.

The posture taken here can be summarized as follows:

1. The surround – the immediate physical and social resources outside the person – participates in cognition, not just as a source of input and a receiver of output, but as a vehicle of thought.
2. The residue left by thinking – what is learned – lingers not just in the mind of the learner, but in the arrangement of the surround as well, and it is just as genuinely learning for all that.

Indeed, in the person-plus spirit we might venture a rather brash claim called the *equivalent access hypothesis*. This hypothesis asserts that thinking and learning for the person-plus depend only on what might be called the "access characteristics" of relevant knowledge – what kind of knowledge is represented, how it is represented, how readily it is retrieved, and related matters – and not whether the knowledge is located in the person or the surround. If, for example, the student can access ideas in that notebook about 1066 fairly easily, having organized it so well, what does it matter whether the ideas lie inside or outside the student's cranium?

Of course, the case cannot be pressed too far. The claim is certainly not that a set of notes in the best indexed notebook or even a rapid retrieval electronic database is exactly functionally equivalent to a well-memorized battery of facts in long-term memory. Indeed, there are a number of trade-offs between the two. The real claim is more a point of principle: the litmus to be applied is function – a matter of the access characteristics of the information – not locus – a matter of which side of one's skull hosts the information.

Cognition as information flow

Let us sharpen this notion of person-plus on the stone of a very abstract notion: a knowledge-processing system. This system might be a person filling out income tax forms, a computer in an insurance company calculating risks, or DNA replication. In such a system, a typical information-handling episode picks up knowledge from various places in the system and operates on it, often incrementing the knowledge of the system. For example, a person-plus consisting of a person, pencil, tax forms, and instructions would, at a certain point, increment the system's knowledge with the total decrement due to the federal government.

For a broad-stroke analysis of such an episode, one might look to four categories: *Knowledge* concerns what kinds of knowledge are available, including declarative and procedural knowledge, facts, strategies, and skilled routines, in other words knowledge in the broadest sense. *Representation* concerns how the knowledge is represented – in particular, whether in ways that make it easily picked up, transported in the system, and recoded. *Retrieval* concerns whether the system can find the knowledge representations in question, and how efficiently. *Construction* concerns the system's capacity to assemble the pieces of knowledge retrieved into new knowledge structures.

The four together comprise the *access characteristics* of the system – what knowledge it includes access to, via representations that afford what access to information, by way of what retrieval paths for accessing the information, and with what access to further constructions based on that knowledge. Because of the emphasis on access, the entire perspective is called the *access framework* (Perkins & Simmons, 1988; Perkins, Crismond, Simmons, & Unger, in press).

These four categories have been chosen partly because they represent a fairly intuitive partition of the facets of an information-handling episode that can be applied to any system, involving a human or not, examining whatever kinds of knowledge, representation, retrieval mechanisms, and construction mechanisms serve the system. The access framework amounts to a general framework for what might be called an "information flow" analysis. Although all this may sound quite computeresque, no such restriction or even emphasis is intended. DNA replication or the generation of antibodies by the immune system in response to invasion are both processes that could be analyzed in such terms. Both involve certain kinds of information, encoded in certain ways and retrieved by certain paths, and the construction of new (or replicated) information structures.

However, in part the four categories were selected because they prove useful for sorting a number of findings from contemporary psychology about the conditions for good human learning. The following samples show how this sorting works.

Knowledge. Understanding a discipline typically involves not only "content-level" knowledge – facts and procedures – but what might be called "higher-order" knowledge about problem-solving strategies, styles of justification, explanation, and inquiry characteristic of

the domain (Perkins & Simmons, 1988; Posner, Strike, Hewson, & Gertzog, 1982; Strike & Posner, 1985). In many learning situations, neither the learner nor the surround contains much of this higher-order knowledge, a situation that often obscures the meaning and motive of particular facts and procedures.

Representation. A considerable body of work suggests that visual mental models aid us in understanding complex and novel concepts. Able learners may construct such models or something of similar function for themselves, but less able learners benefit more when models are provided (e.g., Gentner & Stevens, 1983; Mayer, 1989; Perkins & Unger, 1989; Salomon, 1979).

Retrieval. Research shows that typical patterns of learning lead to "inert knowledge," that, although forthcoming on the fill-in-the-blank quiz, is not retrieved under authentic conditions of use. That is, such knowledge is represented in the system but with inappropriate retrieval characteristics. Problem-based learning, among other tactics, can help to prime such knowledge for contextually appropriate retrieval (Bransford, Franks, Vye, & Sherwood, 1986; Perfetto, Bransford, & Franks, 1983; Sherwood, Kinzer, Bransford, & Franks, 1987).

Construction. A number of developmental studies suggest that limitations in short-term memory create a processing bottleneck that makes certain concepts inaccessible to the learner. However, a well-designed surround can provide a surrogate short-term memory and support learners in attaining some of these concepts (Case, 1985; Fischer, 1980; Halford, 1982).

Notice that the access framework and general considerations such as those just mentioned offer an analysis of a thinking-learning system somewhat "above" the level of particular psychological theories of mechanism. The access framework does not provide a detailed theory of cognition, but rather an encompassing outline of a cognitive system at a rather high level of description. We do not have to know how the mind does what it does to profile the access characteristics of a person-plus. We only need to recognize the "black box" operating characteristics of the system and to ask whether the hoped-for pattern of information flow can occur.

For example, one does not require a detailed theory of knowledge representation to make the point that, in many learning situations,

higher-order knowledge about the subject matter is *nowhere represented* in the system. One does not need a detailed theory of working memory to make the point that complex tasks and concepts are likely to overload the person-solo.

It is not the intent of the access framework to offer an account of underlying mechanism. Rather, the aim is to achieve an "information flow" analysis, and the claim is that interesting strengths and shortfalls of thinking and learning settings emerge at that level of analysis.

The distribution of thinking and learning generally

The distributed thinking and learning of the person-plus perhaps comes most to the fore in situations of authentic and extended inquiry – a student or a professor developing an essay, an advertising executive contriving a campaign, a director mounting a play, an engineer designing a bridge. Such creative processes have been studied (e.g., Gruber, 1974; Perkins, 1981; Tweeney, 1985, 1992), but they are hardly in the mainstream of psychological inquiry. Let us use the engineer as a focus. The categories of the access framework aid in surveying the ways in which the engineer distributes the thinking and learning that must be done.

Knowledge. From the standpoint of knowledge, the engineer-solo, of course, brings to the task a rich technical repertoire in long-term memory. But also very important are books with tables of materials strength, formulas about stress on beams, regulations governing construction in the state where the bridge will be built, descriptions and images of the bridge site, and so on.

Representations. Besides mental representations, the engineer employs text, mathematical tables and formulas, and drawings to explore both structural and aesthetic alternatives. It is likely today that the engineer would utilize computer-aided design, with its powerful capacity to render and rotate three-dimensional visualizations of the project.

Retrieval. The engineer employs tables of contents and indexes in books, conventional table lookup processes for reading numerical tables, the retrieval resources of the computer-aided design system, and perhaps key words to probe bibliographic databases for the latest information on some point of construction.

Construction. The engineer works amid a surround providing massive short-term and long-term memory support through drawings and notes on paper and through the computer-aided design system. Memory aside, the setting affords computational support for a number of valuable cognitive operations. Hand calculators enable the engineer to make simple computations. The power of a computer-aided design system permits the engineer to view the evolving design from different angles with full precision, an achievement possible with pencil and paper only by tedious redrawing or building a three-dimensional model, the classic and quite serviceable approach.

Moreover, these points about the knowledge, representation, retrieval, and construction address the physical, not the social, surround. Almost certainly, the engineer is part of a team, and its collaborative processes contribute to the picture. The team, too, is part of this engineer's person-plus. Indeed, perhaps a better phrase is "people-plus" – the functioning cognitive unit is the team, plus its physical support system of scratch pads, technical tables, computer-aided design systems, and so on.

The tacit views of psychology and education

Such stories are easy to tell whenever complex inquiry occurs. Moreover, other activities that are less thinking-intensive in the usual sense also typically involve massive environmental support. For example, the bustle of a cook in a kitchen, where not only the cookbook but the presence of implements in stored positions or out on the table ready to be used, or placed in the sink for later washing, constitutes a kind of cognitive scaffold that would make it difficult indeed for the cook to lose his or her place in the process.

In contrast, typical psychological and educational practices treat the person in a way that is much closer to person-solo. The usual laboratory subject is rarely equipped with more than pencil and paper to support cognition. This would serve nicely if studying cognition meant no more than studying the Platonic mind abstracted from the physical world. However, the claim here is that person-plus situations have emergent characteristics that substantially change the information-processing capacities of the system and that warrant investigation and understanding.

Schools mount a persistent campaign to make the person-plus a person-solo. "Person plus pencil, paper, text, almanac, encyclopedia" and so on is fine for *studying,* but the target performance is typically "person plus paper and pencil." And the pencil and paper are conceived not so much as thinking aids as a hopper into which the person-solo can pour concrete evidence of achievement.

Certainly there is justification for some concern with the person-solo. But so much seems quite misguided, for at least two reasons: (1) If part of the mission of schools is to prepare students for out-of-school performance, this perseveration on the person-nearly-solo is not "lifelike"; (2) most students have much to learn about the art of distributing cognition, and schools should help.

The cause of overemphasis on the person-solo may be this. There is a widespread belief in what I have previously called the "fingertip effect": Simply make a support system available and people will more or less automatically take advantage of the opportunities that it affords (Perkins, 1985). Were the fingertip effect a reality, there would be little need for education to worry about students learning to make the best use of support environments – ones as simple as pencil and paper or as complex as word processor, outliner, or hypertext environment.

However, considerable evidence argues that the fingertip effect is a sham. For example, investigations of the impact of word processors on students' writing have shown that the powers of structural transformation of the text afforded by word processors are hardly used at all. Instead, the students utilize this powerful mechanism primarily to make minor stylistic, grammatical, and spelling corrections and to get nice printouts (Cochran-Smith, 1991; Daiute, 1985, 1986). Experienced writers *do* use the resources for structural revision, and were doing so more painfully by hand before they began to use word processors.

But we do not have to turn to high technology to make the case that people miss some of the best uses of the physical support structures at their disposal. Research on reading strategies shows that readers can benefit enormously by taking advantage of abstracts, tables of contents, section headings, and captions in previewing an article they are going to read and by being aware of the kinds of text structures they are reading (e.g., Higbee, 1977). Yet without instruction

in reading strategies, students do the straightforward thing: They read linearly from beginning to end.

Or, for example, conventional linear note taking in a class is arguably less efficacious than recourse to notational techniques that show the structure of knowledge in a better way, such as concept mapping (Novak & Gowin, 1984) or the use of a variety of graphic organizers for capturing particular patterns of ideas, such as narrative, compare–contrast relationships, or argument–counterargument (Jones, Pierce, & Hunter, 1988–9; McTighe & Lyman, 1988).

In summary, two principal points invite recognition. First, in rich contexts of inquiry and indeed in most everyday activities we find immense physical support systems for cognition; these support systems speak to all four facets of the access framework, providing (1) needed knowledge, (2) accessible representations, (3) efficient retrieval paths, and (4) constructive arenas (scratch pads, work benches, etc.) that support the making of things and the structuring of ideas. Second, the best use of these physical support systems is an art. It is not so commonly found. And conventional instruction does little to acquaint students with this art, mistakenly expecting the fingertip effect to do the job.

The distribution of the executive function

We do not have to untangle the paradoxes of free will to recognize that cognitive organisms – even machines – have an executive function. That is, there are routines that do the often nonroutine job of making choices, operating at decision points to explore the consequences of options and select a path of action. While the preceding section examined the distribution of thinking and learning in general, it is worthwhile to focus for a while on this special case – the executive function and its distribution in various versions of the person-plus.

To relate the executive function to the access framework, making choices in complex circumstances is plainly a highly constructive act; consequently, the executive function inevitably draws on knowledge and representational, retrieval, and constructive resources. Sometimes, however, the executive function is fulfilled in a more straight-

forward way through memory for previous choices at similar choice points, mostly a matter of retrieval from internal or external representations of the knowledge stored about the choice point.

The executive function of a person-plus during thinking and learning can be distributed in a number of ways. We most often envision a person deciding for him- or herself. But many other scenarios occur. For example, often during instruction the teacher decides what would be best to do. The learner, to be sure, decides whether or not to go along. A text or workbook or computer-aided instruction (CAI) program has an implicit set of executive decisions built in: Read this chapter, then do this exercise. Solve this problem; depending on how well you do, the computer will provide another problem.

In cases like these, it would be easy to sloganeer about a learner's loss of autonomy, certainly an important issue. But it is definitely *not* the presumption here that this taking over of the executive function by the learning support system (teacher, book, computer, etc.) is generally a bad thing. All depends on the wisdom of the support and on whether the learner eventually has a chance to develop whatever executive functions are needed to gain from the learning experience.

Indeed, ceding the executive function to the surround is often one of the most powerful moves we can make. If the directions for assembling the components of a new stereo system are clearly written, we do best to follow them. When concerned with the capriciousness of human judgment in cases of conflicting interest, we make written contracts and laws that freeze certain patterns of decision making. Of course, all this is usually done with some latitude or power of override left to a human-solo or a social group (e.g., juries, judges), but that should not obscure the basic tactic of ceding considerable executive function to the physical surround.

There are also interesting mixed cases. The menu systems commonly used in computer interfaces leave choices to the user, but organize the options on pull-down menus that anticipate the user's likely priorities. Thus, the surround undertakes part of the normal executive function – constructing a representation of the option space. To turn for a moment to the social distribution of intelligence, clinicians commonly avoid taking over executive function for their clients, because they want to build the clients' autonomy. But they

scaffold and nudge the clients in the construction of the option space. To turn to education, Mark Lepper's studies of expert tutors disclose a complex pattern of interaction in which the tutor leaves the student feeling empowered but subtly exercises enormous control on the student's path through questions and challenges of various sorts (Lepper, Aspinwall, Mumme, & Chabay, 1990).

Granted that there are many sound distributions of executive function between the person and surround, in some circumstances there is an executive function gap: Learners do not automatically know how to handle distributed executives. For example, following directions with precision (a ceding of the executive function to the source of the directions) is a very useful skill; but many learners do not seem to muster related skills of self-monitoring, checking, and attentional control and so do not track well directions that ask for high precision. For a socially oriented example, some people of all ages seem to have difficulty making decisions in group contexts; indeed, sorting out priorities in a group involves a multitude of complications and a maze of cross-talk not encountered in solo decision making.

So education might in principle give students more help in the art of distributing the executive function. In practice, however, instruction generally has its own executive shortfalls. Many instructional designs may leave students – especially weaker students – with inadequate executive function: The learner does not know and cannot readily figure out quite what to do, and the surround does not provide enough help. This is commonly the case in open-ended learning situations, such as the use of Logo when teachers are not skilled in the art of scaffolding students' activities (Papert, 1980; Pea & Kurland, 1984a, b; Salomon & Perkins, 1987). The implication is *not* that such environments should involve a strong executive function in the surround, telling learners what to do, but rather that such environments should involve enough support specifically for the executive function that students can find their way into worthwhile activities.

For example, Harel (1991) reported an experiment in which youngsters used Logo to develop simple instructional software about fractions. As is not always the case in Logo settings, care was taken to create a support structure around the students sufficient to sustain

fairly systematic progress through a long-term project. While the students had considerable autonomy, notebooks, discussions, and other mechanisms scaffolded good task management. The students progressed well on the projects and gained dramatically in both programming skills and fractions understanding.

Transitions of the executive function during learning

This brings us to the point that the distribution of the executive function during learning can change in various ways. In the most familiar pattern, the learner cedes executive function to the surround and gradually gets it back as he or she gains mastery over the knowledge and skills in question. The catch, in much educational practice, is that the student never gets back much autonomy at all. The educational surround typically maintains extensive executive control through the formal learning process. Then the learner leaves the educational surround to function alone, suddenly responsible for an executive function but entirely unprepared for it.

A classic example is problem selection. Conventional education does virtually all problem selecting for students, deciding which problems are worth doing and, often, in what order. Then the assignments stop. And we are puzzled when students do not see opportunities in everyday life to apply what they have learned. Such a mishap is commonly called "lack of transfer." But this is something of a misdiagnosis, because it fails to recognize that the students have never had a chance to learn the process we are hoping they will transfer – problem selection. The surprising thing is not that learners commonly miss "real-life" applications, but that from time to time students find some. This is, if anything, evidence of the remarkable reach of transfer under uncongenial conditions (Perkins & Salomon, 1988; Salomon & Perkins, 1989).

So investing the learner with needed executive function is an important, yet neglected educational agenda. At the same time, there is no intent here to beat an ideological drum to the tune of total learner autonomy. Depending on the nature of the learning objective, the learner may not ever need executive control. Consider, for example, some of the CAI environments designed to routinize skills, such as typing, word recognition, or spelling. They exercise executive

function to lead the learner through the learning process, but enhanced executive functioning may not be important to the automatized skill itself.

For example, research suggests that, for some slow arithmetic learners, automatization of basic arithmetic skills is a critical bottleneck that can be eliminated by drill and practice under time pressure (Hasselbring, Goin, & Bransford, 1988). Coercive as this may seem, it is not the executive function of the student that needs developing in this case. A learning experience that pays no attention to the student's executive function but simply develops the student's automaticity and stops serves perfectly well.

Finally, it is important to recognize that in some learning situations the learner moves toward ceding more executive function rather than less. A manager learns to cede executive function to capable subordinates. A museum goer, after some experience with self-directed tours versus the use of audiopacks provided by the museum, may learn to cede executive function to the audiopack, which he or she finds can provide a better tour than a self-constructed one, at least until the person gains more experience. A married couple, thrashing out some problems, may fashion written rules for themselves, such as "We talk about finances for no more than an hour on Saturday morning." In general, in the course of learning, executive function may appropriately flow toward or away from the learner, depending on the circumstances.

To summarize, there is a complex tale to be told about the social and physical distribution of the executive function. We cede executive function to the physical, never mind social, surround much more often and for much better reasons than might at first be expected. The person-plus is often substantially empowered by ceding the executive function.

At the same time, the nuances of the game are all-important. Not infrequently a person-plus fails in an activity because neither the person nor the surround nor the two in combination provides for a good executive function for the activity. Often, instruction seizes executive control when it might be better to scaffold the executive function of the student, helping to decide but not deciding. And often, when the executive function must be transferred to the learner, the instructional surround does nothing to mediate this transition.

The distribution of higher-order knowledge

As mentioned earlier, the "knowledge" category of the access framework distinguishes between content-level knowledge – the facts and procedures of a subject matter – and "higher-order" knowledge, including discipline-appropriate problem-solving strategies and patterns of justification, explanation, and inquiry characteristic of the discipline (cf. Perkins et al., in press; Perkins & Simmons, 1988). Higher-order knowledge in a domain includes, for example, heuristics of problem solving (e.g., Polya, 1954, 1957; Schoenfeld, 1982, 1985) and patterns of explanation, justification, and inquiry (e.g., Schwab, 1978; Toulmin, 1958). Such higher-order knowledge occurs not only in academic domains but in daily life; an example is knowledge about everyday decision making or self-management.

Elsewhere, we have argued that an appreciation of higher-order domain knowledge is very important for learning in a domain (Perkins & Simmons, 1988). Many misconceptions in mathematics and science can be traced in part to the lack of higher-order knowledge that gives the appropriate conceptions a supportive context while disclosing the weaknesses of the inappropriate conceptions.

This higher-order knowledge not only informs the construction of understandings of content-level knowledge but also provides grist for the executive function discussed in the preceding section. Problem-solving strategies and patterns of justification, explanation, and inquiry give the executive major paths of domain-relevant behavior to choose among. Lacking this higher-order structure, the executive is limited in its choices to the retrieval of content knowledge and the execution of routine procedures, such as the algorithms of arithmetic. It is the higher-order aspects of a domain that infuse domain-related activities with significance.

The presence of higher-order knowledge

With such points in mind, it becomes important to ask how higher-order knowledge is distributed in thinking-learning situations. Perhaps the first point to make echoes one underscored for the executive function: In many person-plus situations, there is *no* appreciable representation of higher-order knowledge either in the person

or in the surround. For example, many textbooks in science simply do not touch, in any but the most superficial ways, upon the processes and commitments of science (Evans, Honda, & Carey, 1988). History books commonly say nothing at all about the epistemological basis of history: how historians generate hypotheses about the past and test them against historical evidence. Often, textbooks make little use of "mental-state terms" such as "think," "know," "infer," "assume," "conclude," and "hypothesize" (Olson, & Astington, 1990; Olson & Babu, in press). Students themselves can hardly be expected to conjure up such ideas out of nothing.

It is commonplace to note such shortfalls in conventional instructional materials. However, the point goes well beyond textbook bashing. Many innovative learning environments that dramatically improve some access characteristics of a learning situation nonetheless do not touch on the problem of higher-order knowledge.

To make this proposition concrete, consider the example of the Geometric Supposer (Schwartz & Yerushalmy, 1987), an ingenious computer program designed to restore exploration and discovery to the teaching of Euclidean geometry. The Supposer does this by way of three basic tactics. First, it makes geometric constructions extremely easy: A user can request that a triangle be drawn, an altitude be dropped, a parallel be constructed, and so on. Second, it makes measuring such constructions in order to check conjectures very easy. For example, a student can request a measurement of two sides of a triangle to see whether they are equal. Third, the Supposer makes retesting a conjecture on different versions of the same process extremely easy. For example, having begun by constructing a random triangle, dropping an altitude, and so on, the student can request that the system repeat the entire construction beginning with a new randomly chosen triangle or one the student specifies. Thus, the student can discern whether the same construction on a different triangle yields the conjectured geometric relationship again.

The Geometric Supposer is one of the best-known technological innovations in mathematics education, justifiably renowned for restoring a measure of creativity to the traditional geometry curriculum. From the standpoint of the access framework, it clearly affords improved access characteristics in a number of ways. Basically, the Supposer provides a constructive arena – in the most literal sense. The

operations in this arena are executed with much more fluency and precision than a student could readily muster using a straightedge and compass. Regarding retrieval, the system affords immediate retrieval and re-execution of the previous construction. Regarding representation, the system of course displays constructions, but also, through artful screen layout, makes plain the repertoire of constructive operations that one might apply.

Thus, the student has all the resources to proceed with higher-order inquiry in the domain of geometry. However, remember the fingertip effect, the mischievous assumption that people readily take advantage of opportunities that are there. With the Supposer, many students do not so readily see the opportunities, and many teachers do not know quite what to do to lead students to those opportunities. For the Supposer does not include any knowledge about the higher-order aspects of the domain (Chazen, 1989).

The argument here is not that the Supposer should be improved by somehow building in more higher-order knowledge to the software itself. On the contrary, the Supposer is one of my favorite examples of software design and is fine as it stands. The point, rather, is that the surrounding instruction must – and sometimes does – include articulate attention to the higher-order aspects of geometry.

In general, cognitive opportunities are not in themselves cognitive scaffolds. Thoughtful, innovative technological resources that afford great opportunity for higher-order kinds of thinking and learning in a domain do not in themselves necessarily provide cognitive scaffolding.

The locus of higher-order knowledge

Recognizing the need for higher-order knowledge in the person-plus, we can ask where its locus should be. In general, this essay has pressed the point that locus in itself is not important – whether in the surround or in the person. What counts are the access characteristics – for example, how transparently the needed knowledge is represented and how readily it can be retrieved. This was the equivalent access hypothesis advanced at the outset.

But this hypothesis does not necessarily imply that higher-order knowledge can just as well be located in the surround. All depends on

whether approximate functional equivalence might be achieved – which is not so easily done with higher-order knowledge. By and large, the higher-order knowledge should be in the person (or distributed among the minds of participating persons) rather than physically downloaded.

Why is this? First of all, because higher-order knowledge is referenced more or less continuously by the executive function in complex inquiry activities. It is not like a formula that, checked once a month, might as well be buried in a book. Second, higher-order knowledge is fairly stable, not ephemeral like scratchwork, and so it might as well sit in long-term memory. Third, higher-order knowledge is relatively compact compared with the mass of facts and procedures in a domain. So there is no particular problem with the sheer bulk of it; indeed, the problem is more one of getting it to operate actively in guiding the executive function. Fourth, a person seriously involved in a discipline or caught up in the activities of everyday life functions in many surrounds – at his or her desk, the walls crowded with reference books; at meetings, with notebook in hand; washing the dishes or mowing the lawn; or hashing over a personal problem with the help of a close friend and a couple of beers. The higher-order knowledge, needed moment-to-moment in diverse settings, cannot readily be relegated to a particular physical storage system.

Accordingly, contrary to the general spirit of distributed cognition, the internalization of higher-order knowledge in a domain is particularly to be urged.

What is a person as a cognitive agent?

We began with an asymmetry. Most views of thinking and learning lean toward the person-solo, neglecting the ways in which people employ the surround (including other people) to support, share, and undertake outright aspects of cognitive processing. In contrast, one can take a person-plus perspective on thinking and learning, treating the person plus surround as one system, counting as part of the thinking what gets done or partly done in the surround, counting as learning traces left in the surround (assuming it stays accessible) as well as the person, and in general picking the lock of a person-solo view of thinking and learning.

So when we pick the lock, do we find anything interesting in the larger space we enter? The case was made that genuine contexts of inquiry typically involve massive distribution of thinking and learning between the person and the surround. Active thinkers assemble around themselves a rich surround and interact with it in subtle ways to achieve results that would be difficult for the person-solo. Unfortunately, schools show a strong bias toward the person-solo. They rely on the "fingertip effect" assumption, presuming that people will automatically take effective advantage of the surround just because it is there. They thereby miss the opportunity to cultivate all sorts of skills concerning the artful distribution of thinking and learning.

Thinking and learning often involve ceding the executive function to the surround in worthwhile ways. No dogged vision of person-solo autonomy seems warranted. At the same time, a prevalent problem of thinking and learning occurs when neither person nor surround nor the two together supports an effective executive function. Many open-ended instructional settings suffer from this problem. Another mishap occurs when the executive function is ceded to the surround for a while during the early stages of learning, but the learner never gets it back.

Higher-order knowledge informs the executive function in important ways. Arguably, most higher-order knowledge ought to find its locus in the person; as explained earlier, it should be continuously on tap, not buried in a sourcebook or other surrogate memory. In many thinking and learning situations encountered in schools, necessary higher-order knowledge is to be found neither in the persons nor in their surrounds. Even innovations strikingly insightful in boosting other access characteristics for the person-plus commonly do not address the presence of higher-order knowledge in the system.

All of this demonstrates that a story can be told about the happenings and mishappenings of thinking and learning, employing the notions of distributed thinking and learning and the access framework to see the person-plus rather than the person-solo as the key player. This story makes salient some neglected features of cognition and throws into relief certain shortfalls of typical and even innovative education.

As Pea (Chapter 2, this volume) points out, a further potential payoff of perspectives emphasizing distributed cognition is an enlarged

concept of human development. The Piagetian perspective, for example, has highlighted assimilation of and accommodation to the environment by the organism, as though the environment were given and the person there to learn to deal with it. Of course, this is true to a considerable extent. But also, people select and build their physical and social environments, and do so in part to support cognition. In this sense, there is mutual assimilation and accommodation between the person and the surround – a complex equilibration process, if you like, in the person-plus.

A Vygotskian perspective would highlight the learner's assimilation from the social surround of patterns of cognition (Vygotsky, 1962, 1978). The notion of distributed cognition would also mark the person's modifying influence on the social surround. Moreover, it would emphasize the importance of the *physical* surround alongside the social as a major factor in the cognition of the person-plus system.

Finally, many contemporary developmental perspectives make much of limitations in working memory as a developmental bottleneck, and various experimental results suggest that physical support in the surround can enable the person-plus to deal with some complex concepts that would be unmanageable for the person-solo. It would be interesting to investigate to what extent the available physical supports in person-plus settings generally absorb some of the cognitive burden of the thinking youngsters and whether they could easily absorb more of it with some adjustments.

In short, a person-plus view suggests that some basic parameters and trajectories of human development might change according to what might ordinarily be considered nuances in the surround and the person's relation to it. This is surely something we need to understand better. And it is surely possible to envision an educational process oriented more toward the person-plus, empowering learners to capitalize with greater awareness and art upon the cognitive resources afforded by the physical and human resources around them – indeed, empowering learners to construct around themselves their personal "plus," their own surround for an agenda that will evolve with that surround.

Such an educational tactic surely would be in keeping with the human trend from one-pebble-per-sheep accounting systems to hiero-

glyphics and on. It is notable how vigorously we human beings, given half a chance, function as agents recruiting into the cognitive enterprise not only other people but the insentient physical things around us, arranging them and refashioning them so that they become "partners in cognition" (Salomon et al., 1991).

Reciprocally, it seems worth reflecting that at the center of every person-plus is, of course, at least one person. Indeed, any person alone is the intersect of the set of person-pluses in which that person participates. A person alone, then, becomes the queen bee in a hive of innumerable participations.

So what is the person proper – the person-solo? The tendency of our language and much of educational practice and psychological research is to say yes, the person proper is the person solo. But this paradigm needs to be rethought. Perhaps the person proper is better conceived not as the common core but the set of interactions and dependencies; not as the intersection but the union of involvements; not as the pure and enduring nucleus but the sum and the swarm of participations.

References

Baddeley, Alan. (1982). *Your memory: A user's guide.* New York: Macmillan.

Bransford, J. D., Franks, J. J., Vye, N. J., & Sherwood, R. D. (1986, June). *New approaches to instruction: Because wisdom can't be told.* Paper presented at the Conference on Similarity and Analogy, University of Illinois, Urbana.

Case, R. (1985). *Intellectual development: Birth to adulthood.* New York: Academic Press.

Chazen, D. (1989). *Ways of knowing: High school students' conceptions of mathematical proof.* Unpublished doctoral dissertation, Harvard Graduate School of Education, Cambridge, MA.

Cochran-Smith, M. (1991). Word processing and writing in elementary classrooms: A critical review of related literature. *Review of Educational Research, 61,* 107–55.

Craik, F. I., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior, 11,* 671–84.

Daiute, C. (1985). *Writing and computers.* Reading, MA: Addison-Wesley.
    (1986, May). Physical and cognitive factors in revision: Insights from studying with computers. *Research in the Teaching of English, 20,* 141–59.

Evans, R., Honda, M., & Carey, S. (1988). *Do theories grow on trees?* Unpublished manuscript, Harvard Graduate School of Education, Educational Technology Center, Cambridge, MA.

Fischer, K. W. (1980). A theory of cognitive development: The control and construction of hierarchies of skills. *Psychological Review, 87*(6), 477–531.

Gentner, D., & Stevens, A. L. (Eds.). (1983). *Mental models.* Hillsdale, NJ: Erlbaum.

Gruber, H. (1974). *Darwin on man: A psychological study of scientific creativity.* New York: Dutton.

Halford, G. (1982). *The development of thought.* Hillsdale, NJ: Erlbaum.

Harel, I. (1991). *Children designers.* Norwood, NJ: Ablex.

Hasselbring, T., Goin, L., & Bransford, J. (1988). Developing math automaticity in learning handicapped children: The role of computerized drill and practice. *Focus on Exceptional Children, 20*(6), 1–7.

Higbee, K. L. (1977). *Your memory: How it works and how to improve it.* Englewood Cliffs, NJ: Prentice-Hall.

Jones, B. F., Pierce, J., & Hunter, B. (1988–9). Teaching students to construct graphic representations. *Educational Leadership, 46*(4), 20–5.

Lepper, M., Aspinwall, L., Mumme, D., & Chabay, R. (1990). Self-perception and social perception in tutoring: Subtle social control strategies of expert tutors. In J. M. Olson & M. P. Zanna (Eds.), *Self-inference processes: The Ontario Symposium* (Vol. 6, pp. 217–37). Hillsdale, NJ: Erlbaum.

Mayer, R. E. (1989). Models for understanding. *Review of Educational Research, 59,* 43–64.

McTighe, J., & Lyman, F. T. (1988). Cueing thinking in the classroom: The promise of theory embedded tools. *Educational Leadership, 45*(7), 18–24.

Novak, J. D., & Gowin, D. B. (1984). *Learning how to learn.* Cambridge University Press.

Olson, D. R., & Astington, J. W. (1990). Talking about text: How literacy contributes to thought. *Journal of Pragmatics, 14*(15), 557–73.

Olson, D. R., & Babu, N. (in press). Critical thinking as critical discourse. In S. P. Norris & L. M. Phillips (Eds.), *Foundations of literacy policy in Canada.* Calgary: Detselig.

Papert, S. (1980). *Mindstorms: Children, computers, and powerful ideas.* New York: Basic Books.

Pea, R. D., & Kurland, D. M. (1984a). On the cognitive effects of learning computer programming. *New Ideas in Psychology, 2*(2), 137–68.

(1984b). *Logo programming and the development of planning skills* (Rep. No. 16). New York: Bank Street College.

Perfetto, G. A., Bransford, J. D., & Franks, J. J. (1983). Constraints on access in a problem solving context. *Memory & Cognition, 11*(1), 24–31.

Perkins, D. N. (1981). *The mind's best work.* Cambridge, MA: Harvard University Press.

(1985). The fingertip effect: How information-processing technology changes thinking. *Educational Researcher, 14*(7), 11–17.

Perkins, D. N., Crismond, D., Simmons, R., & Unger, C. (in press). Inside understanding. In D. N. Perkins, J. Schwartz, M. M. West, & M. S. Wiske (Eds.), *Teaching for understanding in the age of technology.* New York: Teachers College Press.

Perkins, D. N., & Salomon, G. (1988). Teaching for transfer. *Educational Leadership, 46*(1), 22–32.

Perkins, D. N., & Simmons, R. (1988). Patterns of misunderstanding: An integrative model of misconceptions in science, mathematics, and programming. *Review of Educational Research, 58*(3), 303–26.

Perkins, D. N., & Unger, C. (1989, June). *The new look in representations for mathematics and science learning.* Paper presented at the Social Science Research Council conference "Computers and Learning," Tortola, British Virgin Islands.

Polya, G. (1954). *Mathematics and plausible reasoning* (2 vols.). Princeton, NJ: Princeton University Press.

(1957). *How to solve it: A new aspect of mathematical method* (2d ed.). Garden City, NY: Doubleday.

Posner, G. J., Strike, K. A., Hewson, P. W., & Gertzog, W. A. (1982). Accommodation of a scientific conception: Toward a theory of conceptual change. *Science Education, 66*(2), 211–27.

Pressley, M., Wood, E., & Woloshyn, V. (1990). Elaborative interrogation and facilitation of fact learning: Why having a knowledge base is one thing and using it is quite another. In W. Schneider & F. W. Weinert (Eds.), *Interactions among aptitudes, strategy, & knowledge in cognitive performance* (pp. 200–21). New York: Springer.

Salomon, G. (1979). *Interaction of media, cognition, and learning.* San Francisco: Jossey-Bass.

Salomon, G., & Perkins, D. N. (1987). Transfer of cognitive skills from programming: When and how? *Journal of Educational Computing Research, 3,* 149–69.

(1989). Rocky roads to transfer: Rethinking mechanisms of a neglected phenomenon. *Educational Psychologist, 24*(2), 113–42.

Salomon, G., Perkins, D. N., & Globerson, T. (1991). Partners in cognition: Extending human intelligence with intelligent technologies. *Educational Researchers, 20,* 2–9.

Schoenfeld, A. H. (1982). Measures of problem-solving performance and of problem-solving instruction. *Journal for Research in Mathematics Education, 13*(1), 31–49.

(1985). *Mathematical problem solving.* New York: Academic Press.

Schwab, J. (1978). *Science, curriculum and liberal education: Selected essays* (I. Westbury & N. J. Wilkof, Eds.). Chicago: University of Chicago Press.

Schwartz, J. L., & Yerushalmy, M. (1987). The Geometric Supposer: Using microcomputers to restore invention to the learning of mathematics. In D. N. Perkins, J. Lochhead, & J. Bishop (Eds.), *Thinking: Proceedings of the second international conference* (pp. 525–36). Hillsdale, NJ: Erlbaum.

Sherwood, R. D., Kinzer, C. K., Bransford, J. D., & Franks, J. J. (1987). Some benefits of creating macro-contexts for science instruction: Initial findings. *Journal of Research in Science Teaching, 24,* 417–35.

Strike, K., & Posner, G. (1985). A conceptual change view of learning and understanding. In L. H. T. West & A. L. Pines (Eds.), *Cognitive structure and conceptual change* (pp. 211–32). New York: Academic Press.

Toulmin, S. E. (1958). *The uses of argument.* Cambridge University Press.