

Thinking with external representations

David Kirsh

Received: 4 November 2009 / Accepted: 20 December 2009 / Published online: 10 February 2010
© The Author(s) 2010. This article is published with open access at Springerlink.com

Abstract Why do people create extra representations to help them make sense of situations, diagrams, illustrations, instructions and problems? The obvious explanation—external representations save internal memory and computation—is only part of the story. I discuss seven ways external representations enhance cognitive power: they change the cost structure of the inferential landscape; they provide a structure that can serve as a shareable object of thought; they create persistent referents; they facilitate re-representation; they are often a more natural representation of structure than mental representations; they facilitate the computation of more explicit encoding of information; they enable the construction of arbitrarily complex structure; and they lower the cost of controlling thought—they help coordinate thought. Jointly, these functions allow people to think more powerfully with external representations than without. They allow us to think the previously unthinkable.

Keywords External representations · Thinking · Interactivity · Sense making · Cost structure

1 Introduction

This essay is an inquiry into why thinking and sense making, so often, is interactive. By ‘interactive’ I mean a back and forth process: a person alters the outside world, the changed world alters the person, and the dynamic continues. Reading a text silently is not an interactive

process, for my purposes here, though it is extremely active. Reading and underlining the text, or reading and summarizing it, even reading and moving one’s lips, are.

The puzzle that interaction raises about sense making and thinking can be posed like this. In a closed world, consisting of a person and an external representation—a diagram, illustration, spoken instruction, or written problem statement—why do people do more than just think in their heads? If we assume there is no one to ask, no tool to generate novel results, no clock to provide chronometric input, no process to run and observe, then there is nothing external, no oracle or tool, that a person can consult or manipulate, that yields new information. The environment contains nothing that could not be inferred through reflection, at least in principle. So why bother to mark, gesture, point, mutter, manipulate inert representation, write notes, annotate, rearrange things, and so on? Why not just sit still and ‘think’?

Figure 1a illustrates a simple case where interaction is likely. A subject is given the sentence, S_1 :

A basic property of right-angled triangles is that the length of a median extending from the right angle to the hypotenuse is itself one half the length of the hypotenuse.

What do people do to understand S_1 ? After re-reading it a few times, if they have a good imagination and some knowledge of geometry, they just think. They make sense of it without physically interacting with anything external. Most of us, though, reach for pencil and paper to sketch a simple diagram such as Fig. 1a or b. Why? If the sentence were “The soup is boiling over” or “A square measuring 4 inches by 4 inches is larger than one measuring 3 inches by 3 inches,” virtually no one would bother. Comprehension would be automatic.

D. Kirsh (✉)
Department of Cognitive Science, University of California,
San Diego, La Jolla, CA 92093-0515, USA
e-mail: kirsh@ucsd.edu




Fig. 1 By drawing a right angle triangle and median, it is easier to understand the claim ‘in a right-angled triangle, the median of the hypotenuse is equal in length to half the hypotenuse’. The illustration does not carry the generality of the linguistic claim, but it is easier to convince ourselves of its truth. In **b**, the equalities are explicitly marked and the claim is even easier to read; it helps hint at problem solving approaches

Anyone who believes in situated (Robbins and Ayedede 2009), distributed (Hollan et al. 2000), or extended cognition will have a ready explanation (Clark 2008). Cognitive processes flow to wherever it is cheaper to perform them. The human ‘cognitive operating system’ extends to states, structures, and processes outside the mind and body (Giere 2004). If it is easier to understand a sentence by creating a diagram to help interpret it, then one does that instead of thinking internally alone. The analogy is with a computer system that has memory systems and scratch pads in different media and locations. The decision whether to work out a computation in one or more scratch pads is determined by the type of operators available in each, the cost of operating in each pad, and the availability of space to work in. Processes should migrate to wherever they are best, or most easily, performed.

Figure 2 is suggestive of this view of extended or distributed cognition. Because people are embedded in their environments, they are densely coupled to the outside. Cognitive processes drift to wherever they are more cost effective (Russell et al. 1993; Pirolli 2007). It’s all about the cost structure of computation in each of the interconnected sub-systems. Evidently, when pen and paper is handy, and when the sentence is complex enough, it pays to make a good illustration; it reduces the overall cognitive cost of sense making.

Although I believe this is, essentially, a correct account, it is only one of the reasons people interact with external representations. The others have to do with ways changing the terrain of cognition can do more than change cost




Fig. 2 This image of a coupled system represents the state space trajectory over time of certain cognitive processes. Processes readily move from one side to the other, wherever the cost of an operation is lower

structure. Chiefly, these involve a) access to new operators—you can do something outside that you cannot inside; b) you can encode structures of greater complexity than you can inside, external mechanisms allow us to bootstrap to new ideas and new ways of manipulating ideas; or, c) you can run a process with greater precision, faster, and longer outside than inside—you can harness the world to simulate processes that you cannot simulate internally or cannot simulate as well. In short, these other ways are ways of concern changing the domain and range of cognition. This is a striking claim. It suggests that as our environments and technology changes, we will be able to think about things that today are unthinkable.

There is a further reason why people interact with external representations: to *prepare themselves* to coordinate internal and external states, structures, and processes. This feature of interaction is fundamental to our understanding of external representations but rarely studied. See Kirsh (2009a, c). For example, before subjects use a map to wayfind, they typically orient or ‘register’ the map with their surroundings; they put it into a usable correspondence with the world (Koriat and Norman 1984). Many people also gesture, point, talk aloud, and so on. In principle, none of these actions are necessary to establish a correspondence between elements in the map and the things those elements refer to. Eye movements, mental projection, and other non-interactive techniques may suffice for map-based navigation. But external interactions are commonplace, and a major aspect of understanding representations.

I have found these ‘extra’ actions also pervasive when people try to understand and follow instructions. In pilot studies, we found that subjects engage in ‘interpreting’ actions when they follow origami instructions. They re-orient or register the origami paper with the instruction sheet; they point to elements on the instruction sheet and then focus attention on the counterpart property of the paper; they mutter, they gesture, they move the paper about. This activity is part of processing the meaning of the instructions.

To a lesser degree, the same thing often happens when non-expert cooks follow recipes. They keep place with their finger; they arrange the ingredients to encode their order of use (Kirsh 1995); they read the recipe aloud, ask themselves questions about ingredients, or mutter reminders. We observe similar behavior when people assemble furniture. Far from just thinking and then executing instructions, people perform all sorts of apparently ‘superfluous’ actions that facilitate comprehension. They point, mumble, move the instruction manual around, encode order of assembly in the arrangement of pieces. These actions are not incidental, they are often vitally important to sense making and effective action.

One function of these extra actions is to help people anchor their mental processes on external features or processes. Another is to help them tease out consequences, to




Fig. 3 This illustration suggests that there are three cost structures: the cost of inner operations on states, structures, processes, the cost of outer operations on states, structures, processes, and the cost of coordinating inner and outer processes, which includes the cost of anchoring projections, and the cost of controlling what to do, when, and where to do it

deepen their semantic and pragmatic processing of the instructions. In both cases, people need to establish a coordination between what goes on inside their heads and what goes on outside. They construct a correspondence, a coordination relation, synchronization. Because these coordination processes are not cost-free, Fig. 2 overly simplifies the relation between internal and external processes. A further process needs to be included: the coupling process itself, the special actions performed to establish a cognitive link. Figure 3 illustrates this added cost-laden process: anchoring (see Kirsh 2009b, for an initial discussion of this third cost space). See also Hutchins (2005) for an account of perceptual anchoring and Fauconnier and Turner (2002) for anchoring in mental spaces.

As important as this anchoring or grounding process is I restrict my focus, in the remainder of this work, to ways we interact with representations to alter the cognitive terrain rather than the interactions we perform to prepare ourselves to engage the external part of that terrain through anchoring.

2 Materiality and its consequences

The argument others and I have long advanced is that people interact and create external structure when thinking because:

Through interaction it is easier to process more *efficiently* and more *effectively* than by working inside the head alone (Clark 2008; Kirsh 1995, 1996, 2009; Kirsh and Maglio 1994).

Efficiency usually translates as speed accuracy. Interactive cognition enhances efficiency because it regularly leads to fewer errors or to greater speed.

Effectiveness usually translates as coping with harder problems. Interactive cognition enhances effectiveness because it regularly helps subjects to compute more deeply, more precisely, and often more broadly.

The idea is that by operating with external *material*, pen, paper, ruler, and then working to meet one's goals and sub-

goals using that external material—*draw* a triangle, *mark* the half point of the hypotenuse—subjects benefit from physical constraint and visual hints that help cognition (Scaife and Rogers 1996). This plays out in a few ways. For instance, the constructive process helps drive interpretation. Because action is primarily serial, it is incremental; a structure emerges step-by-step and a subject must resolve specific problems. What size should the base and height be? Does it matter? Does the median bisect the right angle? Working with tools and external structure has the effect of grounding interpretation in an ever more constrained case study. After choosing the size of the right angle triangle, the requirement to split the hypotenuse in half is fully concrete. It is now 'split this hypotenuse'. This incremental, interactive process, filled with prompts, hints, visible possibilities and impossibilities, provides more constraint than mentally computing a conceptual whole solely from the semantics of linguistic parts. The linguistic formulation is more general, but it is also less constrained (see Figs. 1 and 4).

A second way materiality figures in cognition is by explicitly involving visual and motor cortex. When a structure is viewable and drawable, its properties *prime* a constellation of associations. Just by grappling with external material—using rulers, making lines intersect—and then looking at the results, a set of properties and possibilities of forms are encountered and primed. For instance, if two lines intersect then they define a set of angles. It is natural for visual attention to focus on estimating angles. Are they equivalent? If the triangle has a right angle, then automatically a network of spatial concepts related to right triangles are activated, particularly associations derived from previous work with diagrams of right triangles. These visual and physical associations may be different and more extensive than associations derived from verbal accounts. This is apparent whenever a tool is in hand. Rulers prime measuring actions and thoughts; protractors encourage thoughts of angles and degrees.

The benefits of interacting with an external representation are especially clear for complex structures. As the complexity of a linguistic specification of a visual structure increases, it becomes more rewarding to make sense of the sentence by constructing a physical drawing and looking at it than by constructing that geometric form in one's mind's eye and making sense of the sentence internally. Most people find it easier to think in terms of physical lines than




Fig. 4 Choices must be made when drawing a triangle. Should the triangle be long and short? Isosceles? Will any of these choices affect the truth of the sentence? By having to resolve these questions, subjects are helped in the problem solving process

in terms of the mental counterparts of lines, particularly the more lines there are, or the more complex the structure. Even though some people can do things in their heads that others cannot, there is always a point where internalist cognitive powers are overwhelmed and physical realization is advantageous (see Kirsh 2009b). Thus, although from a purely logical point of view, a closed system of world and person contains no additional information after that person has drawn an interpretation than before, there nonetheless are important changes wrought by interaction that can positively alter the cognitive terrain. Specifically, these interactive changes concern:

- What's *active inside* the person's head—what's being attended to, what's stored in visual or motor memory, and what's primed—an external structure encourages a visual scanpath that activates expectations, drawing the structure displays angles, lengths, and will cause distant cognitive associations in motor and visual cortex;
- What's *persistent outside*, and in the visual or tangible field—an external structure holds a structure constant until it is added to; the structure does not decay the way mental structures and processes do, and it supports repeated perceptual inquisition;
- How information is encoded, both inside and outside in virtue of interaction. Because there is an external structure present, subjects can try out different *internal and external representational forms*, the two forms can play off each other in an interactive manner, leading to new insights.

The upshot is that, often, humans are able to improve their thinking and comprehension by creating and using external representations and structures. By working outside, they change what is inside and interactively they can reach new thoughts. This may be stunningly obvious, yet it is sufficiently foundational and far-reaching to deserve analytic and empirical exploration.

Let me press this idea further by turning now to seven distinct benefits that externalization of structure confers.

3 Shareable and identifiable objects of thought

When someone externalizes a structure, they are communicating with themselves, as well as making it possible for others to share with them a common focus. An externalized structure can be shared as an object of thought. This reification of internal object—this externalization—has benefits for both parties.

Here is an example. In Fig. 5, an explicit geometric form has been added to the body position of a dancer. Using a video to demonstrate torsion, Bill Forsythe, a noted choreographer, had his colleagues visually annotate key body features on the video as he spoke. He first identified points on his body, orally, then, as he turned his discussion to line segments, such as the line between elbow and hand, these were superimposed on the video, and finally he talked of joining the segments into a three-dimensional trapezoid, and his viewers saw a representation of the three-dimensional form appears on screen. It was then easy for viewers to see the effect of movement on deformations of the trapezoid. Forsythe relied on his listeners seeing the visible annotation, the trapezoidal structure, as he explained the ideas of torsion, sheer, and body axes. (Forsythe 2008).

One virtue of this particular annotation is that by having verbally defined the structure to be manipulated, and then visibly locating it on his body, the choreographer and anyone looking at the video, knows that if they refer to any visible part of the trapezoid their reference will be understood. They can ask pointed questions about how the shape figures in what the speaker is saying, or even how some specific feature—the apex or base—figures in an abstract idea. For instance, once there are external lines and planes anyone can ask the speaker, or themselves, which body




Fig. 5 Bill Forsythe, a noted contemporary choreographer, has begun documenting certain concepts and principles of choreography in film. Here he explains torsion. The annotation makes it easy for the audience to refer to otherwise invisible structures

positions keep the volume of the shape constant, or which movements ensure the top plane remains parallel to the bottom plane. Choreographers find such questions helpful when thinking about body dynamics and when they want to communicate ideas of shearing and torsion to their dancers. But they are hard to understand if the group does not share a visual or projected image of a transforming shape.

Physically reifying a shape through annotation adds something more than just providing a shared reference; it provides a persistent element that can be measured and reliably identified and re-identified. Measurement is something one does after a line or structure has been identified. This need not always require an external presence. Some people are able to grasp the structure of a superimposed trapezoid purely by mentally projecting an invisible structure onto the body. They listen to the speaker; watch his gestures, and project. But even for these strong visualizers, annotating still helps because once something is externalized it has affordances that are not literally present when projected alone.

For instance, when the lines of a shape are externalized, we can ask about the length of the segments and their angles of intersection. We know how to measure these elements using ruler and protractor. Lines afford measuring. Granted, it is still possible, though not easy, to measure the length of mentally projected lines if the subject is able to appropriately anchor the projected lines to visible points. A choreographer, for instance, can refer to the length of someone's forearm through language or gesturally mark a structure without having to annotate a video. But can he or she refer reliably to the length of lines *connecting* the top and bottom planes of a complex structure without having those planes visibly present? Those lines have to be anchored on the body. If a structure is as complex as a truncated pyramid, which has eight anchor points, it must be constructed in an orderly manner, much as Forsythe did in his annotated video, else there is too much to keep track of mentally. This does not decisively show that such structures cannot be identified and marked out by gesture and posture without visible annotation. But the complexity of mental imagery and mental projection goes way up as the number of anchors increases; or when the target body moves the anchor points; or, worst of all, if invisible anchor points are required, as would be the case if the conceptualized pyramid were to extend right up to its apex. The peak itself would be floating in air, unconnected to anything material. Imagine trying to use that invisible anchor as an anchor for something else. By contrast, once the form is made manifest in visible lines, all such elements can be explicitly referred to, even visibly labeled; they can be located, measured, intentionally distorted if so desired, and the nature of their deformation over time can be considered. They become shared objects of thought.

This is worth elaborating. To say that something is, or could be, an object of thought implies the thinker can mentally refer to it—in some sense the thinker can grasp the referent. A shared object of thought means that different thinkers share mechanisms of reference and for agreeing on attributes of the referent. For instance, Quine (1960), following Strawson (1959), argued that objects must have identity conditions, as in his motto “No entity without identity”. Entities have to be identifiable, re-identifiable, and individuatable from close cousins. Would the structures and annotations in Fig. 5 meet those criteria if imagined or projected mentally? It depends on how well they are anchored to physical attributes. Certainly there are some people—choreographers, dancers, and people with wonderful imaging abilities—who can hold clear ideas of projected structure, and use them to think with. As long as there is enough stability in the ‘material anchors’ (Hutchins 2005) and enough expertise among the subjects to ensure a robust projection, the lines and shapes these experts project onto the visible environment meet most criteria of ‘entification’, though, of course this is purely an empirical claim. But most of us find that it is easier to think about a structure that has been reified by adding visible or tangible elements to the environment. The structure is more vivid, more robust and, clearer—a better object of thought. Almost everyone needs to see the lines and shapes to see subtle geometric relations between them. So, we create external structure. It is by this act of materializing our initial projections, by forming traces of those projections through action, or material change, that we create something that can serve as a stepping-stone for our next thoughts.

This interactive process of *projecting structure then materializing it*, is, in my opinion, one of the most fundamental processes of thought. When we interact with our environment for epistemic reasons, we often interact to create scaffolds for thought, thought supports we can lean on. But we also create external elements that can actually serve as vehicles for thoughts. We use them as things to think with.

All too often, the extraordinary value of externalization and interaction is reduced to a boring claim about external memory. “Isn't all this just about offloading memory?” This hugely downplays what is going on. Everyone knows it is useful to get things out of the head and put where they can be accessed easily any time. It is well known that by writing down inferences, or interim thoughts, we are relieved of the need to keep everything we generate active in memory. As long as the same information can be observed and retrieved outside, then externalizing thought and structure does indeed save us from tying up working memory and active referential memory.

But memory and perception are not the same thing. Treating information to be the same whether it is outside or

inside ignores the medium-specific nature of encoding. The current view in psychology is that when we visually perceive an external structure, the information that enters is stored first in visuo-spatial store (Baddeley 2000; Logie 1995), before being processed for use in later mental processes. Since the form a structure is encoded in profoundly affects how easily it can be used in a process, it is an open question how much internal processing is necessary to convert an external structure into an internal structure that is usable. Accordingly, it cannot be assumed, without argument that the costs are always lower in perceptually retrieving information than ‘internally’ retrieving information, even if that information is complex and voluminous and something we would normally assume is more efficiently stored externally. The strength of this concern is obvious if the information element to be perceived is buried in visual clutter. Much will depend on visual complexity, the form information is encoded in, how easy it is to perceive the structure when it is wanted, and so on. Even when an object of thought is present in a clear and distinct way—as Forsythe’s graphical annotations are—it still must be perceived, then gestalted, and conceptualized. Do we really know the relative cost of grasping an externally represented content versus an internally represented one?

The implication is that using the world as external storage may be less important as a pure source of cognitive power than using the world for external computation. Things in the world behave differently than things in the mind. For example, external representations are extended in space, not just in time. They can be operated on in different ways; they can be manually duplicated, and rearranged. They can be shared with other people. Tools can be applied to them. These differences between internal and external representations are incredibly significant. They are what makes interactivity so interesting.

I turn now to another of these differences: the possibility of manually reordering physical tokens of statements. Because of rearrangement, it is possible to discover aspects of meaning and significance—implications—that are hard to detect from an original statement when viewed in isolation. By reordering and rearranging what is close to what, we change a token’s neighborhood; we change the space of what is cognitively near.

4 Rearrangement

The power of physical rearrangement, at least for vehicles of propositions, such as sentences, logical formulae, pictorial narratives, is that it lets us visually compare statements written later with those written earlier; it lets us manipulate what is beside what, making it easier to perceive semantically relevant relations. For instance, we can take lemmas that are non-local in inference space—inferences that are logically downstream from the givens and usually discovered later, hence written further down the page—and rewrite them so they are now close to earlier statements. Statements that are distant in logical space can be brought beside each other in physical space. If we then introduce abbreviations or definitions to stand in for clusters of statements, we can increase still further the range of statements we can visually relate. This process of inferring, duplicating, substituting, reformulating, rearranging and redefining, is the mechanism behind proofs, levels of abstraction, the lisp programming language, and indeed symbolic computation more generally.

The power of *rearrangement* is shown in Fig. 6. The problem is to determine whether the six pieces on the left are sufficient to build the form on the right. What do you need to do to convince yourself? Since the problem is well posed and self-contained, the question again, is ‘why not just work things out in your mind?’ In Fig. 6, you have no choice: because the pieces are not movable, no doubt, you will confine you’re thinking to looking and imagining the consequences of moving and rotating them. But, if the problem were posed more tangibly, as a jigsaw puzzle with movable tiles, wouldn’t it be easier to try to *construct* an answer in the world than to think through an answer internally?

Reorganizing pieces in physical space makes it possible to examine relations that before were distant or visually complex (e.g., rotations and joins). By re-assembling the pieces, the decision is simply a matter of determining whether the pieces fit perfectly together. That is a question resolvable by physically fitting and visually checking. Interaction has thus converted the world from a place where internal computation was required to solve the problem to one where the relevant property can be



Fig. 6 Can the jigsaw images on the left be perfectly assembled into the picture on the right? If you could rearrange the pieces, the answer would be trivial. The answer is no. Can you see why? Why, in general, is it easier to solve jigsaw puzzles tangibly?

perceived or physically discovered. Action and vision have been substituted for imagery, projection, and memory. Physical movement has replaced mental computation. Instead of imagining transformations, we execute them externally.

It is tempting to interpret the benefits of rearrangement entirely in cost structure terms: processes migrate to the world because they are cheaper or more reliable there. Evidently, physical manipulation, at times, is cognitively more efficient and effective than mental manipulation. So, on those occasions, it is rational to compute externally.

And sometimes that is all there is to it. For example, in Tetris, subjects can choose between rotating a tetrazoid in their heads and rotating it in the world (Kirsh and Maglio 1995). Since physical rotation is, in fact, a bit faster than mental rotation, the cost incurred by occasionally over-rotating a piece in the world is more than made up for by the benefits that come from the faster and less error prone decision making based on vision.

Yet, it is not always so. In solving jigsaw puzzles, more is at stake than cost alone. As the descriptive complexity of the board state increases, there comes a point where it is hard, if not impossible, for someone to hold the complete structure in mind. The very act of trying out a move mentally causes the internally maintained structure to degrade. Imagine trying to assemble twenty separate pieces in your mind, and then checking to see if the twenty-first will fit anywhere in the mentally sustained assembly. The twenty-first piece may be the last straw, total overload, causing the whole mental structure to lose integrity.

The analogy is with swap space in a computer. Once a threshold of complexity is reached, a computer begins to degrade in its performance. In fact, if its flailing is serious enough, it reaches a standstill, where it takes so much of its working memory to hold the board state, that the simple act of changing that state exhausts memory. The system lacks the resources to keep track of what it has tried already and what remains to be tried. It has to place in long-term memory the part of the board state it is not currently checking, so that it can process the steps in its program telling it what to do next. Then, to do the next thing, it has to bring back part of the board state in long-term memory, and swap out the control state. The result is that the system may cycle endlessly in the same subset of states, never canvassing the part of the state space where the solution is to be found. Zero progress.

It is not quite like that in the world. Because of physical persistence, the board remains the same before and after a subject thinks about moves. Unlike the mental realm, the stability of a physical state is not significantly affected by its complexity. A twenty-piece assemblage is just as stable as a ten-piece assemblage.

There are limits in the physical world too. Once a board arrangement has been changed physically, the previous state is lost, unless a further trace, an annotation was created, or a digital image taken. So searching for a solution in the world, as opposed to in the head, is not always better. But with enough externalization of state—enough external record keeping—there are jigsaw puzzles that can be solved physically that would be impossible to solve in the head, through mental simulation alone. We can push the complexity envelope arbitrarily far. This cannot be done in the head alone.

I will return to this topic of in principle differences between mental and physical simulation at the end of the essay.

5 Physical persistence and independence

Both rearrangement and having stable objects to think with both rely on physical things being persistent. Accordingly, the next key difference between internal and external representations that should be on our list is their difference in stability and persistence over time. Rearrangement of jigsaw pieces is possible because the different pieces to be arranged are simultaneously present. If six pieces were present before rearrangement, there are six after. Pieces can be moved nearer to each other without destroying their integrity. Even though things are not quite as simple with thinking with physical tokens of sentences, we still can be confident that we have the same thought before and after moving a sentence. Because it is easy to detect differences in a sentence token simply by comparing the original with its copy—that is, before and after copying a sentence inscription—we depend on physical persistence to ensure that we do not change the object of thought just by copying or moving tokens. Other things equal, the sentence ‘this sentence has five words’ means the same whether printed on the right or the left of the page, and whether printed yesterday or today.

The case is rather different for mental representations. How can a subject be sure that the mental image in mind at time t_1 is the same as the one at t_0 ? And how can a subject know whether the addition of another mental image, or a simple rotation of a mental image, has not changed the original image? The only reliable test is whether the image is caused by the same external structure on both occasions. If that structure is not present, there is no objective touchstone to decide sameness. There is just subjective belief. For Wittgenstein (1953), this was a source of skepticism concerning the possibility of knowing one’s mental state without outside referents to ground it. No inner state or inner process without outer criterion. Hence, without external support, there might be no way of knowing whether one has the same thought on two occasions.

The brute fact of physical persistence, then, changes the reliability, the shareability, and the temporal dynamics of

thinking. It is easier to have the same thought tomorrow, if the vehicles encoding the thought, or the cues stimulating the thought, are the same as today's. That's why writing helps. When the vehicle is external, we can also count on other people ratifying that it remains the same over time. So, we can be confident that if we think we are reading the same sentence on two occasions, there is a fact of the matter. Similarly, we can be confident that if we interact with an external representation and we think we have left it unchanged; our judgments are more reliable than those concerning our beliefs about internal representations. In the outside world, there is widespread empirical agreement on the effect of interaction—we know there is a broad class of transformations that leave structures invariant, for example: rotation, translation, lighting change, and so forth. There is no comparable principle for internal representations. We have no way of knowing the constancy of our inner life. This means that we have a better idea of the effect of interacting with external representations than with internal ones.

Physical persistence also differs from mental persistence, and transient mental presence, in increasing the *range of actions* a subject can perform on the underlying thing encoding the representation—the vehicle. In Fig. 5, for example, the truncated '3D' trapezoid is displayed as a line drawing on the choreographer's body. It is shown in stop action. Measurements can be made because the visible structure—the trapezoid—can be frozen for as long as it takes to perform the measurements. Tools can be deployed. The materiality of external representations provides affordances internal representations lack.

Architects, designers, and engineers exploit the benefits of persistence and material affordance when they build models. Models have a special role in thinking, and can for our purposes be seen as two, three, or even four-dimensional external representations: paper sketches—2D; cardboard models, cartoons, and fly-throughs—3D models in space or time; and dynamically changing three-dimensional spatial structures—4D models. To see the extra power, these sort of external representations offer let us look at the scale models that architects build (see Fig. 8a).

Scale models are tangible representations of an intended design. They serve several functions:

1. They can serve as a shared object of thought because they are logically and physically *independent* from their author. They can be manipulated, probed, and observed independently of their author's prior notion about how to interact with the model. This is vital for talking with clients, displaying behavior, functionality and detecting unanticipated side effects. It makes them public and intersubjective.
2. Models enforce consistency. The assumption behind model theory in mathematics is that if a physical

structure can be found or constructed, the axioms that it instantiates must be consistent (Nagel and Newman 1958). Unlike a description of the world, or a mental representation, any actual physical model must be self-consistent. It cannot refer to properties that are not simultaneously realizable, because if it is a valid model it counts as an existence proof of consistency. In a many part system, part A cannot be inconsistent with part B if they both can simultaneously be present in the same superstructure. Similarly, the movement of part A cannot be inconsistent with the movement of part B if the two can be run simultaneously. Build it, run it, and thereby prove it is possible. Inconsistency is physically unrealizable. There are few more powerful ideas in the history of thought than this.

3. Models reveal unanticipated consequences. To say that an external model is independent of its creator is to emphasize that other people can approach the model in ways unconstrained by its creator's intention. Once a structure is in the public domain it has a life of its own. This is well appreciated in the case of ambiguous objects. Look at Fig. 7a. Its author may have intended it to be a convex cube with two concave sides extending to the bottom right. But a viewer may initially see those concave sides as a convex cube with a corner pointing outward. Look at the image longer and other interpretations should appear. Studies on mental imagery have shown that subjects who have not yet detected an ambiguity by the time they create a mental image are not likely to realize the ambiguity inherent in their image (Chambers and Reisberg 1985). It is as if they are sustaining their image under an interpretation, a prior conception. And so, they are closed to new interpretations. When externalized and in the visual field, however, the very processes of vision—the way the eye moves and checks for consistency—typically drives them to see the




Fig. 7 Here are some ambiguous objects. In **a**, a variant of the Necker cube is shown where some corners that start looking convex (outward pointing) change to concave. **a** Ambiguous in several ways. Can you see at least four interpretations? In **b**, the middle element will seem to be a B or 13 depending on whether you read vertically or horizontally. How an object visually appears depends on how an agent looks at it, and this can be affected by how the structure is framed, how it is contextualized, how the agent feels, or what an agent is primed to see




Fig. 8 A 3D model permits architects to view a form from arbitrary angles. It allows them to measure, compare, and look for violation of constraints. By approaching the model from odd angles they can see occlusions and relations that would be extremely hard to see otherwise. In **b**, we see a near perfect coronet formed by a drop of milk in a famous photograph by Harold Edgerton. And in **c**, we see a famous stop frame image of a golfer swinging and then hitting a golf

ball (Reformulation is not limited to formal problem solving. The statement “Police police police police police” is easier to understand when restated at “Police who are policed by police, also police other police”. Most people would not break out their pens to make sense of that statement, but few of us can make sense of it without saying the sentence out loud several times.) (Densmore Shute Bends the Shaft, 1938, © Dr Harold Edgerton, Silver Gelatin Print)

ambiguity.¹ When a structure is probed deeply enough, relations or interactions between parts, that were never anticipated may be easy to discover. Thus, an author may be able to discover interpretations he or she never considered. Whether the thing externalized is a representation of a thought, image, or mental animation, its persistence and independence means that it may be reconsidered in a new light, and interacted with in a new manner.

The power of modeling is a topic of its own. Another special property is one that is made explicit in mathematical simulations that can be run back and forth under a user’s control. Such simulations provide persistence and author independence because they can be run forward, slowed down, stopped, or compared snapshot by snapshot. All normal sized models support our physical interaction. We can move them in ways that exposes otherwise hard to see perspectives and relations. When our interaction is controlled precisely, or interpreted as movement along a timeline, we can juxtapose snapshots in time for comparisons that would simply be impossible otherwise. Without the stability of reproducibility and persistence, some of the ideas we form about the temporal dynamics of a structure would be virtually unthinkable (see Fig. 8a, b).

6 Reformulation and explicitness

A fourth source of the power of interaction relies on our ability to externally restate ideas. Sometimes it is easier to perform restatement externally than in our heads.

¹ Although this argument concerns visual processing, it applies equally well to the physical interactions we can perform on the physical object.

Representations encode information. Some forms encode their information more *explicitly* than others (Kirsh 1992). For example, the numerals ‘ $\sqrt{2209}$ ’ and ‘47’ both refer to the number 47, but the numeral ‘47’ is a more explicit encoding of 47. Much external activity can be interpreted as converting expressions into more explicit formulations, which in turn makes it easier to ‘grasp’ the content they encode. This is a major method for solving problems. For instance, the problem $x = \sqrt[4]{28,561} + \sqrt{2209}$ is trivial to solve once the appropriate values for $\sqrt[4]{28,561}$ and $\sqrt{2209}$ have been substituted, as in $x = 13 + 47$.²

Much cognition can be understood as a type of external epistemic activity. If this seems to grant the theory of extended mind (Clark 2008) too much support add the word ‘managing’ as in ‘much cognition involves managing external epistemic activity’. We reformulate and substitute representations in an effort to make content more explicit. We work on problems until their answer becomes apparent.

The activity of reformulating external representations until they encode content more transparently, more explicitly, is one of the more useful things we do outside our heads. But why bother? Why not do all the reformulation internally? A reason to compute outside the head is that outside there are public algorithms and special artifacts available for encoding and computing. The cost structure of computation is very different outside than inside. Try calculating $\sqrt{2209}$ in your head without relying on a calculator or an algorithm. Even savants who do this ‘by just thinking’ find there is a limit on size. Eventually, whoever you are, problems are too big or too hard to do in the head. External algorithms provide a mechanism for manipulating external symbols that makes

² To see why music can be both referent and representation (terrain and also map) ask whether there is a difference between hearing sound and hearing sound as music. The sound is the terrain; music is the conceptualizing structure that interprets the sound; it maps it.

the process manageable. Indeed, were we to display the computational cost profiles (measured in terms of speed accuracy) for performing a calculation, such as adding numbers in the head vs. using algorithms or tools in the world, it would be clear why most young people can no longer do much arithmetic in their heads. Tools reshape the cost structure of task performance, and people adapt by becoming dependent on those tools.

A second reason we compute outside rather than inside has to do with a different sort of complexity. One of the techniques of reformulation involves substitution and rewriting. For instance, if asked to find the values of x given that $x^2 + 6x = 7$, it is easiest if we substitute $(x + 3)^2 - 9$ for $x^2 + 6x$. This is a clever trick requiring insight. Someone had to notice that $(x + 3)^2 = x^2 + 6x + 9$, which is awfully close to $x^2 + 6x = 7$. By substituting, we get $(x + 3)^2 = 16$, which yields $x = 1$ or -7 . Could such substitutions be done in memory? Not likely. Even if there are people who can, as before there always comes a point, where the requisite substitutions are too complex to anticipate the outcome ‘just by thinking’ in one’s head. The new expressions have to be plugged in externally, much like when we swap a new part for an old one in a car engine and then run the engine to see if everything works. Without actually testing things in the physical world, it’s too hard and error prone to predict downstream effects. Interactions and side effects are always possible. The same holds when the rules governing reformulation are based on rewrite rules. The revisions and interactions soon become too complex to expect anyone to detect or remember them.

7 Natural encoding

Persistence, reordering, and reformulation largely explain why externalizing information and representation may increase the efficiency, precision, complexity, and depth of

cognition. And if these aspects of interaction with external representations do not explain the extra power to be had then simulation does. Still, there is another aspect to consider: how external processes may increase the *breadth* of cognition. To explore this aspect consider again, why we prefer one modality to another for certain types of thinking.

Every representational system or modality has its strengths and weaknesses. An inference or attribute that is obvious in one system may be non-obvious in another. Consider Fig. 9—a musical notation. The referent of the notation is a piece of music. Music is sound with a specific pitch or harmony, volume, timber, and temporal dynamics. The ‘home’ domain of music, therefore, is sound. Visual notation for music is parasitic on the structure of sound. Prima facie, the best representation to make sense of musical structure is music itself; we go to the source to understand its structure (see footnote 2).

If there are times when the source medium is required to represent the content of a thought, then a further reason to externalize content and manipulate it outside is that for some problems, the natural representation of the content only exists outside. Arguably, no one—or at best only a few people—can hear music in their head the way it sounds outside. Mental images of sounds have different properties than actual sounds. Even if it is possible for the experience of the mental image of music to be as vivid and detailed as perception of the real thing, few people—other than the musically gifted, the professional musician, or composer (Sacks 2008)—can accurately *control* musical images in their heads. It is far easier to manifest music externally than it is to do so internally. So, for most people, to make sense of music the first thing to do is to play it or listen to it.

This raises a further requirement on the elements of thought. If a representational system is to function as a medium of thought, the elements in the system must be sufficiently manipulable to be worked with quickly. Spoken and written words are malleable and fast. Body movements




Fig. 9 Imagine hearing 12 s of music. Now look at the musical notation shown here. Notation has the value of showing in space a structure that one hears. But there is much more in the sound as heard than is represented in the notation alone. Sound is the natural representation of music. The same is true for dance. Compare Laban

notation for dance with the full body structure of dancers. Even if the joint structure is captured in Laban, how well represented are the dynamics of movement, the feel of the dance, and its esthetic impression?

for dance, gesture, and perhaps the pliability of clay are too. Musical instruments, likewise, permit rapid production of sound. These outer media or tools for creating media support fast work. They enable us to work with plastic media. In this respect, they enable us to work outside much like the way we work inside, using visual or auditory images for words or ideas, which most of us work with at the speed of thought. If external manipulability matches the internal requirements on speed, then an external medium has the plasticity to be a candidate for thinking in.

8 Using multiple representations

Despite the value of listening to music, there are times when notation does reveal more than the music one has listened to—instances where a non-natural representation can be more revealing and intuitive than the original representation. Because a notational representation uses persistent, space consuming representations, early and later structures can be compared, superimposed and transformed using notation specific operators. As with logic and jigsaw puzzles, it is useful to have tangible representatives that can be manipulated. In these cases, a subject who moves from one representation to the other may extend cognition. By moving between listening to music, and writing it down in a notation, or listening and then reading the notation, or sometimes vice versa, a composer or listener may be able to explore certain elements of musical structure that are otherwise inaccessible. The more complicated the structure of the music, the more this seems to be true. Without interacting with *multiple representations* certain discoveries would simply be out of reach. Visual designers who move between pen and paper, 3D mockups and rapid prototypes are familiar with the same type of process.

9 Construction and tools

The final virtue of external interaction I will discuss is, in some ways, the summation of persistence, rearrangement, and reformulation. It may be called the power of construction. In making a construction—whether it be the graphical layovers of the dancer shown in Fig. 5, the geometric construction of Fig. 1, or building a prototype of a design as in Fig. 8a—there is magic in actually making something in the world. As mentioned in the discussion of scale models, by constructing a structure, we prove that its parts are mutually consistent. If we can build it, then it must be logically and physically viable. If we can run it, then the actions of those parts are consistent, at least some of the time; and if we can run it under all orderings then it is consistent all of the time. The physical world does not lie.

The constructive process has a special place in human thinking because it is *self-certifying*. In mathematics, constructive reasoning means proving a mathematical object exists by showing it. For example, if it were claimed that a given set has a largest element, then a constructionist proof would provide a method for finding the largest element, and then apply the method to actually display the element.

Not every form of human reasoning is constructive. Humans reason by analogy, by induction, they offer explanations, and they think while they perform other activities, such as following instructions, interpreting a foreign language, and so on. None of these are constructive methods in the mathematical sense. However, because of the incremental nature of construction, the effort to construct a solution may also be a way of exploring a problem. When students look for a constructive proof to a geometric problem, they use the evolving external structure to prompt ideas, bump into constraints, and realize possibilities. When they write down partial translations of a paragraph, they rely on explicit fragments to help guide current translation.

The question that begs to be asked is whether thinking with external elements is ever necessary. Can we, in principle, do everything in our heads, or do we need to interact with something outside ourselves in order to probe and conceptualize, and get things right? In mathematics, externalization is necessary, not just for communication, but to display the mathematical object in question. It is like measurement: you cannot provide the value of a physical magnitude without measuring it. You cannot show the reality of a mathematical object (for constructivists) without revealing a proof that parades it. Yet, during the discovery process might not all the thinking be internal, the result of an interaction between elements inside the head? Where is the proof that, at first, all that probing and conceptualizing might not be the outcome of a purely internal activity? Might it not be that all the ‘real’ thinking lives internally, and that the internal activity is simulating what it would be like to write things down outside? Or perhaps that the internal activity amounts to running through how one would present one’s idea to others? Mightn’t the truth be that we needed the outside world to teach us how to think,³ but once we know how we never need to physically

³ Vygotsky among others has suggested that we mastered thinking externally, by conforming our behavior to social norms of rational inquiry, and that what we learned to do first on the outside we came to do on the inside. Thus, the reason we can do math in our head is because we can do math in the world. The same applies to thinking internally in auditory images. We think in words internally, using auditory images of sounds, because when we think in public we speak. Thinking internally is simulating what we do externally, though Vygotsky did believe that inner speech of adults would be much compressed and unintelligible to anyone except the thinker. [Vygotsky 1986].

encounter tangible two- or three-dimensional structures to epistemically probe the ‘world’?

I believe this is wrong: physical interaction with tangible elements is a necessary part of our thinking process because there are occasions when we must harness physical processes to formulate and transition between thoughts. There are cognitive things we can do outside our heads that we *simply cannot do inside*. On those occasions, external processes function as special cognitive artifacts⁴ that we are incapable of simulating internally.

To defend this hypothesis is harder than it might seem. In practice, few people can multiply two four-digit numbers in their heads. And, if they can, then increase the problem to ten digit numbers. This ‘in practice’ limitation does not prove the ‘in principle’ claim, however, that normal human brains lack the capacity to solve certain problems internally that they can solve with external help, with tools, computers or other people. There are chess masters who can play nearly as well, blindfolded as open eyed (Chabris and Hearst 2003).⁵ There is no evidence that a team of chess players is better than an individual.⁶ There are people with savant syndrome who can multiply large numbers in their heads, or determine primes or square roots. Other savants with eidetic memories can read books at a rate of 8–10 s per page, memorizing almost everything (see footnote 5). Tesla said that when he was designing a device, he would run a simulation of it in his head for a few weeks to see which parts were most subject to wear (Hegarty 2004, p. 281, citing Shepherd). Stephen Hawking is said to have developed analytical abilities that allowed him to manipulate equations in mind equivalent to more than a page of handwritten manipulations. For any reasoning problem of complexity n , how do we know there is not some person, somewhere, who can solve it in their head, or could, if trained long enough? To be sure, this says little about the average person. Any given person may reach their computational limit on problems much smaller than n . And our technology and culture has evolved to support the majority of people. So, in practice, all people rely on available tools, practices, and techniques for reasoning. Nonetheless, if a single person can cope with n , then there

⁴ Hutchins (2001).

⁵ Entry from Wikipedia on Kim Peek the inspiration for the character in the movie *Rain Man*: “He reads a book in about an hour, and remembers almost everything he has read (...) His reading technique consists of reading the left page with his left eye and the right page with his right eye and in this way can read two pages at time with a rate of about 8–10 s per page. He can recall the content of some 12,000 books from memory.

⁶ Team chess is virtually unknown, so the evidence for teams being no better than individuals rests, so far, on their being little or no team chess matches. One notable exception is Alexander Alekhine who played 28 teams of four players each, winning 22 wins, three draws and three losses. [http://en.wikipedia.org/wiki/Blindfold_chess].

is an existence proof that the complexity of external simulation does not itself mean that internal simulation is not possible. It suggests that any problem we cannot solve in our heads that we can solve with external help, has more to do with cost structure than with an in principle biological inability.

One way of making the in principle case is to show that there are *operations* that can be performed on external representations that cannot be performed on internal representations, and that, somehow, these are essential. Are there epistemic activities we can perform outside that we cannot duplicate inside, not because of their complexity, but because there are physical properties and technologies available on the outside that we cannot duplicate mentally—operations we cannot mentally simulate with sufficient realism to deliver dependable answers?

Consider Fig. 10. The dots in the two images on the left are related to one another by a rotation of 4° . This is essentially invisible unless the two images are superimposed, as in the image on the right. Superimposition is a physical relation that can be repeated any number of times, as is rotation. Both require control over physical transformations. In the case of superposition, the position of the layers must be controlled precisely, and in the case of rotation, the angle must be controlled precisely. Are there such functions in the brain?

The brain process we are requiring is analog in nature. For over 25 years, a dispute has raged over whether brains support analog processes or whether mental imagery is driven by non-analog means (Pylyshyn 2001). We can sidestep this question, though, by appealing to an in principle distinction between types of processes. In an important paper, Von Neumann (1948) mentioned that some processes in nature might be irreducibly complex. Any description of one of those processes would be as complex as the process itself. Thus, to simulate or model that process one would have to recreate all the factors involved.




Fig. 10 On the left are two collections of random dots. They differ only by the rotation of the plane they are in. On the right, they have been superimposed. Their global relation is now visible. Could this relationship be detected without physically superimposing the patterns? Mental imagery does not support vivid superimposition. And if there are outlier humans who have this odd ability they will necessarily fail as the number of dots or the number of superimpositions increases

This holds whether the simulation or modeling is being performed internally or externally. Von Neumann put it like this:

“It is not at all certain that in this domain a real object might not constitute the simplest description of itself, that is, any attempt to describe it by the usual literary or formal-logical method may lead to something less manageable and more involved.” (p 311)

David Marr invoking the same idea, spoke of Type 2 processes where any abstraction would be unreliable because the process being described evolves as the result of “the simultaneous action of a considerable number of processes, whose interaction is its own simplest description” (Marr 1977). Protein folding and unfolding are examples of such processes, according to Marr.⁷ Other examples might be the *n* body problem, the solution to certain market equilibrium problems, situations where the outcome depends on the voting of *n* participants, and certain quantum computations.

The hallmark of these problems is that there exists physical processes that start and end in an interpretable state, but the way they get there is unpredictable; the factors mediating the start and end state are large in number, and on any individual run are impossible to predict. To determine the outcome, therefore, it is necessary to run the process, and it is best to run the process repeatedly. No tractable equation will work as well.

How are these problems to be solved if we have no access to the process or system itself? The next best thing is to run a physically similar process. For example, to compute the behavior of an *n* body system, such as our solar system, our best hope is to construct a small analog version of that system—an orrery—then run the model, and read off the result (see Fig. 11). Using this analog process, we can compute a function (to a reasonable degree of approximation) that we have no other reliable way of computing.

The implication is that for brains to solve these sort of problems, they would have to encode the initial state of the



Fig. 11 In this mechanical orrery by Gilkerson, housed in the Armagh Observatory, the movement of the planets and their moons are mechanically simulated. It is not possible to access an arbitrary position of the system without moving through intermediate states. This is a feature of simulation systems: they do not have a closed form or analytic solution. To compute the state of the system at t_{12} , one must determine the state at t_{11} and move from there

type II system, and then simulate the physical interaction of its parts. If this interaction is essentially physical—if, for instance, it relies on physical equilibria, or mechanical compliance, or friction—there may be no reliable way of running an internal simulation. We need the cognitive amplification that exploiting physical models provide. We would need to rely on the parallel processing, the physical interaction, and the intrinsic unpredictability of those analog systems. There is nothing in our brains (or minds) like that.

The conclusion I draw is that to formulate certain thoughts and to transition to others, we must either be able to represent arbitrarily complex states—states that cannot be represented in compact form—or we must rely on the external states themselves to encode their values and then use them to transition to later states. These external states we are able to name but never characterize in full structural detail.⁸

10 Conclusion

In order to extract meaning, draw conclusions, and deepen our understanding of representations and the world more generally, we often mark, annotate and create representations; we rearrange them, build on them, recast them; we compare them, and perform sundry other manipulations. Why bother? Minds are powerful devices for projecting structure on the world and imagining structure when it is not present. Our inner mental life is plastic and

⁷ From Marr (1977, p. 38) “One promising candidate for a Type 2 theory is the problem of predicting how a protein will fold. A large number of influences act on a large polypeptide chain as it flaps and flails in a medium. At each moment only a few of the possible interactions will be important, but the importance of those few is decisive. Attempts to construct a simplified theory must ignore some interactions; but if most interactions are crucial at some stage during the folding, a simplified theory will prove inadequate. Interestingly, the most promising studies of protein folding are currently those that take a brute force approach, setting up a rather detailed model of the amino acids, the geometry associated with their sequence, hydrophobic interactions with the circumambient fluid, random thermal perturbations etc., and letting the whole set of processes run until a stable configuration is achieved (Levitt and Warshel 1975).”

⁸ In practice, though not in principle, computers fall into this category. When a workplace has been augmented with tools such as wizards, software agents and the like, it is possible to multiply the potency of basic strategies of interaction to the point where such increases qualitatively change what humans can do, what they can make sense of, and so on. Sometimes our best tools are analog, however, and these are the ones that may provide an augmentation to human thought that in principle cannot be achieved without external help.

controllable, filled with images of speech, visual scene, and imageless propositions. For most of intellectual history, this impressive capacity has been assumed sufficient for thought. Why do we bother to interact so much?

I have argued that much of thinking centers on interacting with external representations, and that sometimes these interactions are irreducible to processes that can be simulated, created, and controlled in the head. Often, the reason we interact with external representations, though, boils down to cost. Nothing comes without a cost. A useful approach to understand epistemic interaction is to see it as a means of reducing the cost of projecting structure onto the world. To solve a geometric problem, we might imagine a structure and reason about it internally; we might work with an illustration and project extensions and possibilities. At some point, though the cost of projection becomes prohibitive. By creating external structure that anchors and visually encodes our projections, we can push further, compute more efficiently, and create forms that allow us to share thought. I have presented a few of the powerful consequences of interaction. It is part of a more general strategy that humans have evolved to project and materialize meaningful structure.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Baddeley AD (2000) The episodic buffer: a new component of working memory? *Trends Cogn Sci* 4:417–423
- Chabris CF, Hearst ES (2003) Visualization, pattern recognition, and forward search: effects of playing speed and sight of the position on grandmaster chess errors. *Cogn Sci* 27:637–648
- Chambers D, Reisberg D (1985) Can mental images be ambiguous? *J Exp Psychol Hum Percept Perform* 11:317–328
- Clark A (2008) *Supersizing the mind: embodiment, action, and cognitive extension*. Oxford University Press, Oxford
- Fauconnier G, Turner M (2002) *The way we think: conceptual blending and the mind's hidden complexities*. Basic Books, New York
- Forsythe W (2008) http://www.youtube.com/watch?v=0P_4D8c2oGs&feature=related
- Giere R (2004) The problem of agency in scientific distributed systems. *J Cogn Cult* 4(3):759–774
- Hegarty M (2004) Mechanical reasoning by mental simulation. *Trends Cogn Sci* 8:280–285
- Hollan J, Hutchins E, Kirsh D (2000) Distributed cognition: toward a new foundation for human–computer interaction research. *ACM Trans Comput Hum Interact* 7(2):174–196
- Hutchins E (2001) Cognitive artifacts. In: Wilson R, Keil F (eds) *The MIT encyclopedia of the cognitive sciences*. MIT Press, Cambridge, MA
- Hutchins E (2005) Material anchors for conceptual blends. *J Pragmat* 37:1555–1577
- Kirsh D (1992) When is information explicitly represented? *The Vancouver studies in cognitive science*. Oxford University Press, Oxford, pp 340–365
- Kirsh D (1995) The intelligent use of space. *Artif Intell* 73:31–68
- Kirsh D (1996) Adapting the environment instead of oneself. *Adapt Behav* 4:415–452
- Kirsh D (2009a) Problem solving and situated cognition. In: Robbins P, Aydede M (eds) *The Cambridge handbook of situated cognition*. Cambridge University Press, Cambridge, pp 264–306
- Kirsh D (2009b) Projection, problem space and anchoring. In: Taatgen N, van Rijn H, Nerbonne J, Schomaker L (eds) *Proceedings of the 31st annual conference of the cognitive science society*. Cognitive Science Society, Austin, TX, pp 2310–2315
- Kirsh D (2009c) Interaction, external representation and sensemaking. In: Taatgen N, van Rijn H, Nerbonne J, Schomaker L (eds) *Proceedings of the 31st annual conference of the cognitive science society*. Cognitive Science Society, Austin, TX, pp 1103–1108
- Kirsh D, Maglio P (1994) On distinguishing epistemic from pragmatic actions. *Cogn Sci* 18:513–549
- Koriat A, Norman J (1984) What is rotated in mental rotation? *J Exp Psychol Learn Mem Cogn* 10:421–434
- Levitt M, Warshel A (1975) Computer simulation of protein folding. *Nature* 253:694–698
- Logie RH (1995) *Visuo-spatial working memory*. Lawrence Erlbaum Associates, Hove, UK
- Marr D (1977) Artificial intelligence: a personal view. *Artif Intell* 9:37–48
- Nagel E, Newman J (1958) *Gödel's proof*. New York University Press, New York
- Pirolli PL (2007) *Information foraging theory: adaptive interaction with information*. Oxford University Press, Cambridge
- Pylyshyn Z (2001) Is the imagery debate over? If so, what was it about? In: Dupoux E (ed) *Language, brain, and cognitive development: essays in honor of Jacques Mehler*. MIT Press, Cambridge, MA, pp 59–83
- Quine WVO (1960) *Word and object*. MIT Press, Cambridge, MA
- Robbins P, Aydede M (eds) (2009) *The Cambridge handbook of situated cognition*. Cambridge University Press, Cambridge, pp 264–306
- Russell DM, Card S, Pirolli P, Stefik M (1993) The cost structure of sensemaking. *Proc CHI* 1993
- Sacks O (2008) *Musophilia: tales of music and the brain*. Alfred A. Knopf, New York
- Scaife M, Rogers Y (1996) External cognition: how do graphical representations work? *Int J Hum Comput Stud* 45:185–213
- Shepard RN (1978) Externalization of mental images and the act of creation. In: Randawa BS, Coffman WE (eds) *Visual learning, thinking and communication*. Academic Press, New York
- Strawson PF (1959) *Individuals: an essay in descriptive metaphysics*. Methuen, London
- Von Neumann J (1948/1961) General and logical theory of automata. In: Taub AH (ed) *John von Neumann: collected works*. Volume V: design of computers, theory of automata and numerical analysis. Oxford, Pergamon Press, pp 288–326
- Vygotsky L (1986) *Thought and language*, revised edn. MIT Press, Cambridge
- Wikipedia on Kim Peek (2009) http://en.wikipedia.org/wiki/Kim_Peek. Accessed Nov 2009
- Wittgenstein L (1953) *Philosophical investigations*. Basil Blackwell, Oxford