

# Interaction, External Representation and Sense Making

David Kirsh (kirsh@ucsd.edu)  
Dept of Cognitive Science, UCSD  
La Jolla, CA 92093-0515

## 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 eight ways external representations enhance cognitive power: they provide a structure that can serve as a shareable object of thought; they create persistent referents; they change the cost structure of the inferential landscape; 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.

## Keywords

External representations, interactivity, sense making, cost structure.

## Introduction

Here is a basic puzzle about sense making. In a closed world, consisting of a person and a representation – a diagram, illustration, spoken instruction or written problem statement – why do people so often perform actions to help them understand? If we assume there is no one to ask, no tool to generate new results, no clock to provide chronometric input, no process to run and then observe the outcome, then nothing changes in the environment other than what that person changes. If all the information needed for full understanding is logically present in mind and initial representation, then in principle, the environment contains no additional information after a person's actions than before. Yet people make marks, they gesture, point, mutter, manipulate the inert representation, they write notes, annotate, rearrange things and so on. Why not just 'think'? Why interact?

Figure 1a illustrates a simple case where interaction is almost inevitable. A subject is given the sentence "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 make sure they understand? After re-reading the sentence a few times, if they have an excellent imagination and some knowledge of geometry, they just think about the sentence and come to believe they know what it means. They know how to make sense of it without interacting with anything external. Most of us, though, reach for pencil and paper, and sketch a simple diagram, such as in figure 1a or 1b, to better understand the truth of

the property. 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.




Figure 1. By drawing an example of 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 1b the equalities are explicitly marked and the claim is even easier to read and helps hint at problem solving approaches.

This essay is an inquiry into why we interact with the world when we try to make sense of things. There are, I believe, two major types of interaction concerned with external representations. The first, and most familiar type – the only one I will examine in this article – concerns our reliance on tools, representations and techniques for solving problems and externalizing thought. In the right-angled triangle case, for example, we make an illustration to facilitate understanding. We then perform a variety of operations, mental or physical, on that external representation. My discussion of this resource-oriented sort of interaction focuses on the power of physical sorts of operations – ways sense makers interact to change the terrain of cognition.

The second, and less well-documented type of interaction concerns those things we do to *prepare ourselves* to use external representations, things we do to help us project cognitive structure. They are activities that help us tie external representations to their referents. For example, before we use a map to wayfind, we typically orient or 'register' the map with our surroundings to put it into a usable correspondence with the world. Many of us also gesture, point, talk aloud, and so on. These sorts of 'extra' actions are pervasive when people try to understand and follow instructions. They are not incidental and quite often vitally important to sense making, though rarely studied.

The theme unifying both the first and second types of interaction is their connection with our ability to project structure onto things and then modify the world to materialize or reify our projection. This core interactive process – *project then materialize* – underlies much of our

epistemic and pragmatic engagement of the world. I briefly discuss this dynamic at the end of this article.

## Coupling with External Representations

The most natural way to externally represent two and three-dimensional spatial structures in a persistent re-usable manner is by drawing them using lines or by making physical models. The drawing in figure 1a is a *canonical* way of representing 2-D things; it is flat and uses thin clear lines and shapes. A linguistic statement of the same geometric property, by contrast, is more general but also less 'natural'. It is harder to make sense of.

The view many of us embrace is that people interact and create external structure when thinking because:

*Through interaction it is easier to process more efficiently (i.e. with better speed-accuracy), more deeply, more precisely, and often more broadly than by working inside the head alone. [Kirsh & Maglio 94, Kirsh 95, Clark 08.]*

Operating with external *material*, with pen, paper, ruler, and then working to meet one's goal or subgoals – make the angle 90°, cut the hypotenuse in half – is a process that provides constraint and hints that usually help cognition. Just by grappling with external material – rulers, making lines intersect – our understanding of the properties and possibilities of forms are excited.


The intuitive reason, I suggest, that making the drawing in figure 1a is a cognitively more efficient mechanism for sense making than just thinking about its linguistic statement is that during the constructive process, most people find it easier to translate and conceptualize the sentence in terms of physical lines than in terms of the mental counterparts of lines. In both internal and external cases, sense making is a process. But it is easier to perform the process of sense making externally, by constructing a physical drawing and looking at it than it is to construct a geometric form in one's mind's eye. Drawing is likewise easier than mentally computing a conceptual whole from the semantics of linguistic parts. To be sure, performance varies among individuals. People must have learned to draw. And clearly some people can do things in their heads that others cannot. But there is always a point where cognitive powers are overwhelmed and physical realization is advantageous. [Kirsh 09]. Thus, although from a purely logical point of view a closed system of world and person contains no additional information after drawing or writing down sentences than before, there are important changes wrought by interaction that 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 memory, what's primed;
- What's **persistent outside**, and in the visual or tangible field; and

- How information is encoded, both inside and outside: **internal and external representational form.**

The upshot is that, often, humans are able to improve their comprehension by creating and using external representations and structures. This may be obvious, yet it is sufficiently foundational to deserve analytic and empirical exploration.

Let me press this idea further with another spatial example. In figure 2, a geometric form has been added to the body position of a dancer. Points on the body were first identified, then, line segments such as the line between elbow and hand were superimposed, and then finally the segments were joined in a three dimensional trapezoid and the effect of movement visualized in terms of deformations of this trapezoid. The commentator, in this case the choreographer himself, Bill Forsythe, presupposed listeners could index into the visible annotation, the trapezoidal structure. In his discourse the choreographer referred to these properties to explain such concepts as torsion, sheer and body axes. [Forsythe 08].



**Figure 2.** Bill Forsythe, a noted contemporary choreographer, has begun documenting certain concepts and principles of choreography in film. Here he explains torsion.

What sort of aid to understanding do these types of external visualizations provide? They again concern geometric forms but they are here used in the service of sense making in dance and choreography.

**Shareable and identifiable objects of Thought:** One virtue of this particular annotation is that by having defined the structure to be discussed and then *visibly* locating it on a body, the choreographer and anyone looking at the video, *knows* they can refer to various parts of the trapezoid and anyone present will understand their reference. Everyone can rely on shared knowledge of the visible properties of the shape and ask, in a rather specific way, how they figure in what the speaker is saying or how they figure in some abstract idea. For instance, once there are external lines and planes anyone can ask the speaker, or themselves, which body positions keep the volume of the shape constant, or 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.

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. Even though annotation is not necessary for everyone – some people can grasp the structure of a superimposed trapezoid by imagining or projecting an invisible structure onto the body just by listening to the speaker and watching his gestures – nonetheless materializing a projection through annotation supplies unambiguous affordances that are not literally present if merely projected.

For instance, when the lines of a form 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. Granted, it is still possible, though not easy, to measure the length of mentally projected lines, provided they are appropriately anchored to visible points: a choreographer can refer to the length of someone's forearm through language or gesture without annotation. But can he or she refer to the length of lines connecting the top and bottom planes without having those planes visibly present? Those lines have to be anchored on the body. So a complex structure like a truncated pyramid must be constructed in an orderly manner, much as Forsythe did in his annotated video. This doesn't definitively prove 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, and especially if invisible anchors are required. Once the form is made manifest in visible lines, however, all such elements can be explicitly referred to; they can be measured or intentionally distorted, and the nature of their deformation over time can be considered. They become shared objects of thought.

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 for agreeing on attributes of the referent. For instance, Quine, following Strawson, argued that objects must have identity conditions, as in his motto "No entity without identity". Entities have to be identifiable, re-identifiable and individuable from close cousins. Would the structures and annotations in figure 2 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 often hold clear ideas of projected structure, and use them to think with. As long as there is enough stability in the 'material anchors' [Hutchins 05] 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 a purely empirical claim. But most of us find that reifying structure by adding visible or tangible elements to the environment makes those ideas more vivid, more robust, clearer and easier to work with. Most of us need to see the lines and shapes to see subtle geometric relations between

them. By materializing our initial projections, by creating traces of them through action, most of us find we have created something that can serve as a stepping-stone for our next thoughts. This is why the interactive strategy of *project then materialize* is so powerful. It applies to all humans, but exactly how complex things must get before it is necessary surely varies with imaging capacities and expertise.

All too often the extraordinary value of externalization and interaction is reduced to a 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. Something quasi-symbolic that comes in visually enters some sort of memory system (visuospatial store) but it need not be the same memory system, or encoded in the same way as the information stored in whatever type of working memory is used during problem solving. [Logie 95]. So it cannot be assumed that the costs are always lower in perceptually retrieving information than 'internally' retrieving information. 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. Objects of thought that are perceived must still be gestalted, grasped and conceptualized. External referents are often more costly to grasp than internal ones.


More importantly, the biggest bang from externalizing thought and structure usually flows from differences in the form and properties of internal and external representations. External computation is a far more interesting source of cognitive power than the simple fact that useful information can be stored externally. For example, outer forms can be manually duplicated and rearranged. They are extended in space, not just in time, and can be operated on in different ways. This is an extraordinary gift. By reordering physical tokens of statements, for instance, it is possible to discover aspects of meaning and significance that were hard to detect from an original statement viewed in isolation. This is a big deal. It is one of the next sources of interactive power that I turn to now.

**Rearrangement.** The power of physical rearrangement is that it lets us visually compare statements written later with those written earlier and it let's us manipulate spatial relations to improve perception of 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 by writing them near the givens, or by juxtaposing them with statements written previously in our proof process, we can make them local 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 rationale behind proofs, levels of abstraction, the lisp programming language, and indeed symbolic computation more generally.

The power of *rearrangement* is shown in figure 3 where the problem is to determine whether the six pieces on the left are sufficient to build the image on the right. Since the problem is well posed and self contained, the question again, is why do people not just work in their mind? The answer: because it is easier to do it in the world!

As with proofs, reorganizing pieces in physical space makes it possible to examine relations that before were distal. 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 looking. Thus, interaction has converted the world from a place where internal computation was required to solve the problem to one where the relevant property can be perceived. 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.



**Figure 3.** Can the jigsaw images on the left be perfectly assembled into the picture on the right? Answer: no. Can you see why without moving them? Perhaps. There is at least one simple test that does not require *constructing* the answer. What if the answer were ‘yes’. Could you know without constructing the complete image? Rearrangement of statements, like rearrangement of puzzle pieces, serves to make it easier to notice key attributes.

Needless to say the power of rearrangement can be increased dramatically by digital means. By automating the operations of ordering, sorting and filtering, the cost structure of external operations can be significantly altered. When a workplace has been augmented with tools such as wizards, 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. I resist discussion of these tools here, however, because, in my opinion, all these qualitative changes can be shown equally well through simple everyday examples.

**Persistence and Independence.** Rearrangement would be impossible if the pieces to be arranged were not simultaneously present. In figure 2 the ‘3D’ trapezoid is shown in stop action. Measurements can be made because the structure can be frozen for as long as it takes to perform


the measurements. Architects exploit the power of persistence when they build models. Scale models are tangible representations of an intended design that let an architect and client explore the structure at will. A model is independent enough from its author that it can serve as a shared object of thought. It can be manipulated, probed, and observed independently of its originator’s conception. The same idea also applies to 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. Without the stability of reproducibility and persistence it would be virtually impossible to reason about certain temporal dynamics of a structure and explore some of the complexities of its 3D form.



**Figure 4.** A 3D model permits architects to view a form from arbitrary angles. It allows them to measure, compare, and look for violation of constraints.

Persistence is presupposed in most external operations, though not all. Among experts and in certain everyday

contexts, gesture and linguistic reference can be sufficient for listeners or viewers to project a structure that is mentally persistent and specific enough for speaker and hearer to share reference and a range of operations on that structure. An architect can talk of the curve swept out by a door. See figure 5. Ever after, when in the company of her fellow architects, she can assume her audience will understand reference to that arc, without gesture, and without explicit markings on the floor or plans. There are, however, always limits to this capacity to augment the external world with mentally projected structure that one can assumed is shared.



**Figure 5.** Spatially literate people can project the arc swept out by opening a door without having to mark the external environment.

**Reformulation.** A third source of the power of interaction relies on our ability to restate ideas. Representations encode information. Some forms encode their information more *explicitly* than others [Kirsh 90]. 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  $x = \sqrt[3]{28,561} + \sqrt{2209}$  is trivial to solve once the

appropriate values for  $\sqrt[3]{28,561}$  and  $\sqrt{2209}$  have been substituted, as in  $x = 13 + 47$ .<sup>1</sup>

Much cognition can be understood as a type of external epistemic activity. If this seems to grant the theory of extended mind [Clark 08] 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 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. Again, there are probably some people who can do them. But again, 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

<sup>1</sup> 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.

rules governing reformulation are based on rewrite rules. The revisions and interactions soon become too complex to expect anyone to detect or remember them.

**Natural Encoding.** Persistence, reordering and reformulation largely explain why externalizing information and representation may increase the efficiency, precision, complexity and depth of cognition. To see why external processes may also increase the *breadth* of cognition 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 figure 6 – 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<sup>2</sup>.



**Figure 6.** Imagine hearing 10 seconds 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 heard sound than is represented in the notation. Sound is the natural representation of music.

A further reason it is sometimes advantageous to externalize content and manipulate it outside, then, is that the natural representation of that 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 08] – 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.

**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

<sup>2</sup> 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; the music the conceptualizing structure that maps the sound.

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.

**Construction.** The penultimate 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 visual annotation of the dancer shown in figure 2, a geometric construction to figure 1, the layovers in figure 4, or building a prototype of a design – there is magic in actually making something in the world. Unlike reformulation, a construction is not logically contained in the deductive closure of the representation it extends. When a geometer adds a new line to a triangle the new line is permitted by the logic of shapes but it is not logically implied. It is consistent but not entailed. This opens up a new realm of possibility. A problem solver can add things to the situation in the hope that this extra structure will facilitate discovering the target property or theorem. But like Wittgenstein’s ladder, once the new structure has been used it can be thrown away. This is a remarkable idea. By assuming something we know is not false but also is not something derived from our givens, we are able to discover a truth constructively.

Construction is the closest thing in the deductive world to experimentation. It is also the most tough minded. In the foundations of mathematics there is a dispute whether any mathematical truth can be genuinely proved unless it is reached step by step through a constructive proof. To be constructive a proof must actually display the thing to be proved. It is not enough to prove there exists a solution or an entity, e.g., that the set has a largest element, you must actually show what the element is.

Construction is mechanical and concrete and as such relies on rules of allowable construction. If the construction is defined over a representational system the rules are formal, but if the construction is defined over a more tangible structure the rules are determined by physical principle. In a children’s game of connect the dots the implicit pattern is invisible until the dots have been connected. Augmenting the stimulus with straight lines

reveals significant elements that were otherwise meaningless or invisible.

**Simplifying control.** I close with a virtue of external representations lying at the heart of distributed cognition – offloading control, process management. Starting with the obvious: a list of to do’s relieves a person from remembering what to do next. But so does a table of rows and columns if it is 90% complete and requires values for the remaining cells. A user of the table knows what has to be done to make the table usable. Representations often contain an implicit control structure that behaves like a set of implicit instructions telling us what has to be done. This is part of the support system that external representations provide.

	20-29 yrs	40-49
Male	183 lbs	196
Female		168

## Conclusion

In order to extract meaning, draw conclusions, and deepen our understanding of representations 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. Can’t we just think in our heads? Because nothing comes without a cost, a useful approach to understanding epistemic interaction is to see it as a means of reducing the cost of projecting and imagining by creating external structure that can support more complex projection, more efficient computation, deeper and broader sense making, and more shareable 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.

## References

- Clark, Andy. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. OUP. 2008
- Hutchins, E. Material anchors for conceptual blends. by: Edwin Hutchins. *Journal of Pragmatics*, Vol. 37, No. 10. (October 2005), pp. 1555-1577
- Forsythe W. [http://www.youtube.com/watch?v=0P\\_4D8c2oGs&feature=related](http://www.youtube.com/watch?v=0P_4D8c2oGs&feature=related)
- Kirsh, D. When is Information Explicitly Represented? *The Vancouver Studies in Cognitive Science*. (1990) pp. 340-365. Re-issued Oxford University Press. 1992.
- Kirsh, D. The Intelligent Use of Space. *Artificial Intelligence*, Vol. 73, Number 1-2, pp. 31-68, (1995).
- Kirsh, D. and P. Maglio. On Distinguishing Epistemic from Pragmatic Actions. *Cognitive Science*. Vol. 18, No. 4: pages 513-549. (1994).
- Kirsh, D. *How Interaction Improves Sense Making*. Draft. 2009.
- Logie, R.H.(1995). *Visuo-spatial working memory*, Hove, UK: Lawrence Erlbaum Associates.
- Sacks, O, *Musicophilia : tales of music and the brain*. New York : Alfred A. Knopf, 2008.