David Kirsh

Embodied Cognition and the Magical Future of Interaction Design = Poznanie ucieleśnione i magiczna przyszłość projektowania interakcji

2013

OPEN ACCESS

Embodied Cognition and the Magical Future of Interaction Design[17]

David Kirsh
University of California at San Diego
kirsh[]ucsd.ed

Abstract

The theory of embodied cognition can provide HCI practitioners and theorists with new ideas about interaction and new principles for better designs. I support this claim with four ideas about cognition: (1) interacting with tools changes the way we think and perceive – tools, when manipulated, are soon absorbed into the body schema, and this absorption leads to fundamental changes in the way we perceive and conceive of our environments; (2) we think with our bodies not just with our brains; (3) we know more by doing than by seeing – there are times when physically performing an activity is better than watching someone else perform the activity, even though our motor resonance system fires strongly during other person observa tion; (4) there are times when we literally *think* with things. These four ideas have major implications for interaction design, especially the design of tangible, physical, context aware, and telepresence systems.

Categories and Subject Descriptors: H.1.2 (User/Machine Systems); H.5.2 (User Interfaces): Interaction styles (e.g., commands, menus, forms, direct manipulation).

General Terms: Human Factors, Theory.

Additional Key Words and Phrases: human-computer interaction; embodied cognition; distributed cognition; situated cognition; interaction design; tangible interfaces; physical computation;, mental simulation.

---

[17] The first version of the article was published in ACM Transactions on Computer Human Interaction, Vol. 20, Issue 1, March 2013, Article No. 3. The article is re-published with a kind permission of Holder of the copyright.

1. INTRODUCTION

The theory of embodied cognition offers us new ways to think about bodies, mind, and technology. Designing interactivity will never be the same.

The embodied conception of a tool provides a first clue of things to come. When a person hefts a tool the neural representation of their body schema changes as they recalibrate their body perimeter to absorb the end-point of the tool (Làdavas 2002). As mastery develops, the tool reshapes their perception, altering how they see and act, revising their concepts, and changing how they think about things. This echoes Marshall McLuhan's famous line "we shape our tools and thereafter our tools shape us" (McLuhan 1964). A stick changes a blind person's contact and grasp of the world; a violin changes a musician's sonic reach; roller-skates change physical speed, altering the experience of danger, stride, and distance. These tools change the way we encounter, engage, and interact with the world. They change our minds. As technology digitally enhances tools we will absorb their new powers. Is there a limit to how far our powers can be increased? What are the guidelines on how to effectively alter minds?

Consider a moment longer how coming tools will change us. On the "perception" side, our senses will reveal hidden patterns, microscopic, telescopic, and beyond our electromagnetic range, all visualized imaginatively. On the "action" side, our augmented control will be fine enough to manipulate with micrometer precision scalpels too small for our genetic hands; we will drive with millisecond sensitivity vehicles big enough to span a football field or small enough to enter an artery. Our future is prosthetic: a world of nuanced feedback and control through enhanced interaction. These are the obvious things.

Less obvious, though, is how new tools will redefine our lived-in world: how we will conceptualize how and what we do. New tools make new tasks and activities possible. This makes predicting the future almost out of reach. Designers need to understand the dynamic between invention, conception, and cognition. It is complicated. And changing. Good design needs good science fiction; and good science fiction needs good cognitive science.

Consider next the role the body itself plays in cognition. This is the second clue to our imminent future. The new theory of mind emerging over the last twenty years holds that the physical elements of our body figure in our thought. Unimpaired humans think with their body in ways that are impossible for the paralyzed. If true, this means that thought is not confined to the brain; it stretches out, distributed over body and cortex, suggesting that body parts, because of the tight way we are coupled to them, may behave like cognitive components, partially shaping how we think.

Before the theories of embodied, situated, and distributed cognition "thinking" was assumed to happen exclusively in the head. Voice and gesture were

ways of externalizing thought but not part of creating it. Thought occurred inside; it was only expressed on the outside. This sidelined everything outside the brain. Thus, utterance, gesture, and bodily action were not seen as elements of thinking; they were the expression of thought, proof that thinking was already taking place on the inside. Not really necessary.

On newer accounts, thinking is a process that is distributed and interactive. Body movement can literally be part of thinking. In any process, if you change one of the key components in a functionally significant way you change the possible trajectories of the system. Apply this to thought and it means that a significant change in body or voice might affect how we think. Perhaps if we speak faster we make ourselves think faster. Change our body enough and maybe we can even think what is currently unthinkable. For instance, a new cognitive prosthesis might enable us to conceptualize things that before were completely out of reach. And not just the 10^{20} digit of pi! It would be a new way of thinking of pi; something unlike anything we can understand now, *in principle*. If modern cognitive theories are right, bodies have greater cognitive consequences than we used to believe.

This idea can be generalized beyond bodies to the objects we interact with. If a tool can at times be absorbed into the body then why limit the cognitive to the boundaries of the skin? Why not admit that humans, and perhaps some higher animals too, may actually think with objects that are separate from their bodies, assuming the two, creature and object, are coupled appropriately? If tools can be thought with, why not admit an even stronger version of the hypothesis: that if an object is *cognitively gripped* in the right way then it can be incorporated into our thinking process even if it is not neurally absorbed? Handling an object, for example, may be part of a thinking process, if we move it around in a way that lets us appreciate an idea from a new point of view. Model-based reasoning, literally. Moving the object and attending to what that movement reveals pushes us to a new mental state that might be hard to reach without outside help.

If it is true that we can and do literally think with physical objects, even if only for brief moments, then new possibilities open up for the design of tangible, reality-based, and natural computing. Every object we couple with in a cognitive way becomes an opportunity for thought, control, and imagination. These cognitively gripped objects are not simply thought aids like calculators; things that speed up what, in principle, we can do otherwise. They let us do things we cannot do without them, or at least not without huge effort. The implications of a theory of thinking that allows lifeless material things to be actual constituents of the thinking process are far reaching. They point to a future where one day, because of digital enhancement and good design, it will be mundane to think what is today unconceivable. Without cognitively informed designers we will never get there.

1.1. Overview and Organization

This article has six sections. In the next section, Section 2, I review some of the literature on tool absorption (Maravita and Iriki 2004), and tie this to a discussion of the theory of enactive perception (O'Regan and Noë 2001; Noë 2005), to explain why tool absorption changes the way we perceive the world. The short answer is that in addition to altering our sense of where our body ends each tool reshapes our "enactive landscape"—the world we see and partly create as active agents. With a tool in our hands we selectively see what is tool relevant; we see tool-dependent affordances; we extend our exploratory and probative capacities. This is obvious for a blind man with a cane, who alters his body's length and gains tactile knowledge of an otherwise invisible world three feet away. His new detailed knowledge of the nearby changes his sense of the terrain, and of the shape of things too big to handle but small enough to sweep. He revises his perceptual apprehension of the peripersonal[18] both because he can sweep faster than he can touch and because he has extended his peripersonal field (Iriki et al. 1996; Ladavas 1998). It is less obvious, though no less true, that a cook who is clever with a blade, or knows how to wield a spatula, sees the cooking world differently than a neophyte. Skill with a knife informs how to look at a chicken prior to its dismemberment; it informs how one looks at an unpeeled orange or a cauliflower, attending to this or that feature, seeing possibilities that are invisible to more naïve chefs or diners. The same holds for spatulas. Without acquaintance with a spatula one would be blind to the affordances of food that make them cleanly liftable off of surfaces, or the role and meaning of the way oil coats a surface. With expertise comes expert perception (Goodwin 1994; Aglioti et al. 2008). This is a core commitment of embodiment theory: the concepts and beliefs we have about the world are grounded in our perceptual-action experience with things, and the more we have tool-mediated experiences the more our understanding of the world is situated in the way we interact through tools.

In Section 3, the longest part of the article, I present some remarkable findings that arose in our study of superexpert dancers.

One might think that we already know what our bodies are good for. To some extent, we do. For instance, the by now classic position of embodied cognition is that the more actions you can perform the more affordances you register (e.g., if you can juggle you can see an object as affording juggling) (Gibson 1966). Our bodies also infiltrate cognition because our early sensory experience of things, our particular history of interactions with them, figures in how we

---

[18] Peripersonal space is the three-dimensional volume within arm's reach and leg's reach. Visual stimuli near a hand are coded by neurons with respect to the hand, not the eyes or some other location reflecting egocentric location (Makin et al. 2007).

understand them ever after. Meaning is modal-sensory specific (Barsalou 2008). If we acquired knowledge of a thing visually, or we tend to identify that thing on visual grounds, we stimulate these historic neural connections in the later visual cortex when thinking of it (Barsalou 1999). These visual experiences often activate motor representations too, owing to our history of motor involvement with the things we see. Thus, when thinking or speaking we regain access to the constellation of associations typical of interacting with the thing. Even just listening to language can trigger these activations in the associative cortex. The sentence "the alarm sounded and John jumped out of bed" will activate areas in the auditory and motor cortex related to alarms and jumping out of bed (Kaschak et al. 2006; Winter and Bergen 2012). This is the received embodiment view.

In the findings reported here I discuss additional ways bodies can play a role in cognitive processing, ways we can use the physical machinery of the body and not just our sensory cortex and its associative network. This means that our bodies are good for more things than have traditionally been assumed. More specifically, I discuss howe we use our bodies as simulation devices to physically model things.

For example, we found in our study with dancers that they are able to learn and consolidate mastery of a reasonably complex dance phrase better by physically practicing a distorted model of the phrase than by mentally simulating the phrase undistorted. If all that matters is what happens in the brain we would not observe this difference in learning between simulating in the head and simulating with the body. But somehow, by modeling a movement idea bodily, *even when the model is imperfect*, the dancers we studied were able to learn more about the structure of their dance movement than by simulating it without moving. Perhaps this intuitive. But more surprisingly, the dancers learned the phrase better by working with the distorted model than by practicing the way one intuitively thinks they should: by physically executing the phrase, or parts of it, in a complete and undistorted manner, repeatedly. In other words, our dancers learned best when they explored a dance phrase by making a physical model of the phrase (through dancing it), even though the model they made was imperfect. Standard practice might not be considered to be modeling. No one predicted that finding! The dancers seem to be using their bodies in a special way when they make these imperfect models.

This is not specific to dance. Mechanics trying to understand a machine may sketch on paper an imprecise or distorted model. This can help them explore mechanical subsystems or help them consider physical principles. Architects may sketch in fluid strokes their early ideas to get a feel for the way light pours in, or how people might move through a space. Accuracy is not important, flow is. Violinists when practicing a hard passage may work

on their bowing while largely neglecting their fingers. They are not aiming for perfection in the whole performance; they are fixating on aspects. To fixate on certain aspects it may be easier to work with their body and instrument than to think about those aspects "offline" in their head. These sorts of methods may be common and intuitive; but on reflection, it is odd, to say the least, that practicing (literally) the wrong thing can lead to better performance of the right thing (Kirsh et al. 2012). I think this technique is prevalent, and deeply revealing.

Does anyone understand how or why it works? The knee-jerk reply is that for sketches, at least, the function of the activity is to take something that is transitory and internal—a thought or idea—and convert it into something that is persistent and external—a sketch. This allows the agent to come back to it repeatedly, and to interact with it in different ways than something purely in mind (Buxton 2007). But persistence doesn't explain the utility of making physical actions like gesturing, violin bowing, or dancing, all of which are external but ephemeral. How do we think with these ephemera?

Section 4 explores why such ephemera might be so effective. The answer I offer is that body activity may figure as an external mediating structure in thinking and practicing. The dance practice we observed, called "marking" in the dance world, seems to work well because the dancers model just the *aspects* of a movement they want to think about. This is better than mental simulation alone because making the body move through step one may prime step two more forcefully than just running through step one in the mind's eye. Motor cortex primes motor cortex. Predictably, we also found that practicing the correct movement is a better way to practice than lying down and running through a movement idea mentally. But if getting the body to move for the sake of motor and procedural priming were all that is special about physical practice then why would practicing distorted movements yield better learning than practicing the correct movement?

To explain why working with an imperfect model might be better than working with the real thing, I explore how the body, or physical models more generally, can help people *project* the structure or idea they are most interested in. When a dancer marks a movement with her body she creates a cognitive support for herself that helps her to: (a) manage what she will attend to at each moment, (b) focus her thought on the relevant features or aspects of the movement idea, and (c) compute outcomes and trajectories, ultimately in ways that may be better than through mental simulation, or better even than through correct physical practice (and hence kinesthetic perception too). Sometimes working with a simpler thing, even if it is imperfect, is better than working with a perfect thing.

The comparative advantage of using imperfect models is variable. Sometimes it is best to work directly with real things; to dance the real phrase if you can, to practice the whole musical passage, or to work with real engines. This is probably true for simple dance phrases, simple musical passages, and simple physical objects. Whether it is more effective to work with the real thing or a model depends, of course, on what you are trying to accomplish. Sometimes it is easier to manipulate models than to manipulate the real thing. The real thing may be cumbersome, heavy, or slow and difficult to handle. Sometimes it is better to gesture, sketch, or work with a simplified model. For certain tasks, working with a model has a better cost structure, both physically and cognitively. Similarly, dancing a real phrase may require coping with too many complexities at once. An imperfect model may be more flexible, simple, and adaptable than the real thing.

The same benefits, however, may hold for mental images, which is why sometimes it is so useful to work things through in one's head rather than working directly with real things. Mental images, just as gestures and simplified movements, are fast and flexible. So predictably, sometimes they are the most convenient thing to think with, better than embodied models (that is, gestures and overt movements) and better than working with the real thing. But working with a mental image also has limitations. When an object has a complicated spatial structure, or is highly detailed, it is often easier to simulate outcomes by manipulating either the real physical thing, or an appropriately simplified physical model of the thing than to simulate manipulating that thing through mental imagery (Kirsh and Maglio 1995; Wexler et al. 1998). It all depends on the internal and external cost structure of the manipulation, what is often called the mental and physical costs. The scientific challenge is to determine the right dimensions to measure cost (Kirsh 2010). If we can discover these dimensions, we may be able to predict when working with a basic model is best; that some-times simplified physical models, even biased ones, are better things to think with and practice with than either working with "real" things or working with internal imagery.

The upshot is that, given our case study, it seems that imperfect models can, at times, help us outperform ourselves. Despite our not yet knowing exactly when imperfect models help us outclass, my own belief is that we use our bodies (and hands) far more often for modeling than previously appreciated. This has implications for design. If it is true that imperfect modeling can, at times, facilitate thinking and learning better than imagination or better than working with the real object of knowledge, the question arises as to why *we* must be the modeling agent. Why not watch someone else be the external simulator or watch a computer-created simulation? Maybe it is possible to make our thinking process run faster, or cheaper, or deeper if we piggyback on the actions of others or on the actions of computers.

In Section 5, I present additional results from video analysis of choreographic creation to show that using one's own body to explore a dance movement is a better way to understand a dance movement than watching someone else explore it. This may seem obvious, but the point needs to be made because there has been so much discussion in the neuroscience literature on the power of the motor resonance system (Rizzolati and Sinigaglia 2007; Agnew et al. 2007; Aglioti et al. 2008).

There is extensive neurophysiological evidence of a close link between action observation and action execution. For reviews, see Viviani (2002), Buccino et al. (2004), Rizzolatti and Craighero (2004), Wilson and Knoblich (2005). It has been convincingly argued that we reenact or mimic an actor's movements by covertly behaving as if we are the actor rather than the observer (Sebanz and Shiffrar 2007). These covert actions can be subliminal. The motor system can be activated by "imagining actions, recognizing tools, learning by observation, or even understanding the behavior of other people" (Jeannerod 1994, 2001), as well by the processes of motor preparation that underpin "(intended) actions that will eventually be executed" (Jeannerod 2001). So a covert action is the internal counterpart that may or may not be hooked up to an overt action. As Jeannerod, the originator of the idea, said "Every overtly executed action implies the existence of a covert stage, whereas a covert action does not necessarily turn out into an overt action" (Jeannerod 2001). The surprising thing is that processes in this covert system may be so strong that that even just watching an action may be as powerful a learning experience as performing an action oneself (Cross et al. 2009). This means that we might be able to watch someone else gesture or dance or manipulate gears or sketch a structure and our thinking is driven forward just as effectively as if we were the one overtly gesturing, or dancing, etc. Although the comparison is rarely made, an analogy to listening to someone speak may be apt. When attending to someone talk, if listener and speaker are in tune with them, they seem to synchronize their thinking. To make sense of their speech, their inferences must largely march in step. Might this cognitive resonance also apply by watching others perform action or by watching them manipulate objects?

This ties in with a further thesis of embodied cognition: to fully *make sense* of what we are seeing we need to run our motor system simultaneously with watching to get a sense of what it would be like if we were to perform the action ourselves. Our sympathetic body involvement grounds the meaning of action in a personal way. It activates an ideomotor representation that gives us first-person knowledge of another's action (Shin et al. 2010). Because we see things as if we are the agent we understand the point of the action, the details to be attended to, and the reason it unfolds as it does (Knoblich and Flach 2001). When we experience another's action as if we were that person, we can appreciate why it makes sense to do things that way. We covertly compute the subgoal structure of the action. (Prinz 1997). Evaluat-

ing the scope and limits of this central claim is important in building a balanced view of embodied cognition.

I address this question briefly, again using data from our dance study, by discussing the extra knowledge the choreographer of the piece acquired by executing movement rather than just watching it. I speculate that the key extra he received from overt bodily involvement over and beyond what he would get by simulating an action covertly is knowledge of kinesthetic things that have no visual counterpart: for instance, pain, resistance, gravitational pull.

What is true for choreographic creativity likely applies to other types of creativity. I believe the cognitive importance of overt action generalizes beyond dance and is important for designers to understand. For as we look for new and better ways to extend cognition, we need to know when and how effectively we can piggyback off the efforts of others—how and when we learn by watching—rather than having to be the acting agent who learns by overt doing.

In Section 6, I briefly unify the theory of tools and bodily thinking into an account of how objects, and not just body parts, can be brought into the thinking process. This too is important for designers, since it offers a possible foundation for the power of tangible and physical interfaces. I conclude the article with a brief coda reviewing the main ideas and some of their further implications for interaction design.

2. TOOLS CHANGE OUR BODY, OUR PERCEPTION, OUR CONCEPTION

2.1. The Space Around Us

Studies based on human lesion, monkey neurophysiology, and human imaging, such as fMRI and TMS (Transcranial Magnetic Stimulation), provide evidence that when suitably embodied, human and mammal brains construct multiple representations of space (Colby 1998; Graziano and Gross 1995). Certain brain cells fire specifically when objects approach the space around the body, such as when we see an insect fly toward our face, or when our hands are about to be touched. This near-body region is called peripersonal space. It can be understood informally as the space surrounding an agent that is within easy reach (Ladavas et al. 1998; Brain 1941). In addition to peripersonal representations there are neural representations for personal and extrapersonal space. Personal space refers to the space occupied by the body itself (Vaishnavi et al. 1999; Coslett 1998; Bisiach et al. 1986). Extrapersonal space refers to space beyond the reach of our limbs (Previc 1998; Brain 1941).

2.2. Tools Change Our Body Schema

Tools bear a special relation to peripersonal space since we code the distances of nearby things in manipulation-relative and touch-relative ways (Maravita et al. 2002). That is, we code what is nearby—more precisely, what is "within reach"—in terms of how far we have to move our arms and hands to manipulate or touch things. When we use a tool to reach for a distant object it is as if we are extending our motor capability and we treat our hand as if it is elongated to the tip of the tool. Tool use transiently modifies action space representation by revising what is now within reach. No surprise, then, that humans can quickly adapt their spatial representation to functionally meaningful things such as within fly-swatter distance, tennis reach distance, fencing distance and even pole-vaultable height. As Maravita and Iriki (2004) put it, "neurophysiological, psychological and neuropsychological research suggests that this extended motor capability is followed by changes in specific neural networks that hold an updated map of body shape and posture (the putative 'Body Schema')." Apparently, we change our body schema to include a tool's dimensions (or at least its end-point). We absorb the tool into our functioning body[19]. The original work by Iriki et al. (1996) showed that when Japanese macaques were given a rake and three weeks of training in using the rake to pull in food pellets just beyond their reach, the specific neurons representing the hand and arm, as well as the space around these body parts, changed their firing pattern to include the rake and the space around it. In interview Iriki described it this way:

> *In the parietal association area, there are neurons that compare somatic sensation with visual information and become activated upon recognizing the body. In untrained monkeys, these neurons do not become activated because the rake is nothing more than a foreign implement. After they become able to use the rake as a tool as a result of training, however, the neurons become active as if the rake is recognized as an extension of the hand* (Iriki 2009).

2.3. Extending the Body and Redefining the Peripersonal

The tools we have discussed reshape our peripersonal space by extending it a few feet. Can tools let us extend it to terrains that are geographically remote? This is a useful question for interaction designers. For designers work with a sense of where the body ends and the environment begins. If certain tools can be absorbed, this body boundary becomes an element to be negotiated in design.

---

[19] A further question worth asking is whether our somatic representation of the rigidity and strength of our "extended" limbs is altered when we hold a rigid tool or strap on large skis.

There is ample anecdotal reportage that our sense of where our body boundaries are, and what in space we can affect can be altered through tele-presence and teleimmersion. With digital help we can act on objects arbitrarily distant and then perceptually sense what we are doing. For example, there are tele-presence systems that enable an operator to manage a submersible on the ocean floor, a land vehicle in a war zone, and a scalpel in another town's operating theater, while all the while ensconced in a cozy room some miles away. Given the right sensori-motor hookup the remote human feels as if she is in contact with a robust "enactive landscape" to think, speak, and interact with, as if there. One might think, before studying, that the key success condition is for the tele-agent to have worked in the relevant enactive landscape up close first, using his or her unaided hands and eyes. You need to have worked with a scalpel in your actual hand before mastering it in your digital hand. But this is not really necessary. Pilots of submersibles can be trained on remote enactive landscapes from the start as long as action and feedback are close enough in time. It seems that what falls into your peripersonal space, at one or another moment, can be negotiated early on through practice with tools.

This raises the next question. How different can our remote "body parts" be from our own before we cannot assimilate them? Snap-on arms and legs are one thing. But how about two sets of nine-fingered claws that operate in articulate and continuous ways? Controlling these by means of a piano-like multifingered input device might work for claws with ten or less fingers. But what about twelve-fingered claws, and what about having the fingers work in continuous fashion? Probably not impossible; but clearly an interface challenge. And then there is the question of how different a scene in a virtual world can be before it shatters our situated grasp of things? Can we cope with a world that runs at clock speeds fifty times our own?

A rudimentary start on experimentally determining the constraints on embodied extensions was made by Ikiri et al. (1998) when testing to see if a monkey's sense of hand size changes by replacing the normal image of its hand with an enlarged one. As reported by Blakelee (2004):

> *Dr. Iriki allowed the monkeys to see a virtual hand on a video monitor while the monkey's real hand, hidden from view, operated a joystick. When he made the image of the hand larger, the monkey's brain treated the virtual hand as if it were an enlarged version of its own; the brain's hand area blew up like a cartoon character's hand." Evidently, anatomical mappings can be remapped.*

How far can these remapping transformations go? An enlarged hand seems innocent when compared with some of the mutant alterations we can imagine. Is a person like Edward Scissorhands possible? Are there limits on what can be a prosthetic "body part"? And can these bizarre body parts, especially

the ones that involve distant interaction, be incorporated into our peripersonal space as long as we tightly control them? These are open questions for the embodiment program. They address the core HCI question: what makes a tool, prosthetic, or digital helper work and feel natural? What are the limits to neuroadaptation driven by immersion?

2.4. Tools also Change Our Perception

Perception is altered by our skill in using tools. This is the next implication of extending the embodiment paradigm to include tools. Hills look steeper than normal to subjects wearing a heavy backpack (Proffitt 2006). When a tool is absorbed into our body schema, our perception of height, distance, and related magnitudes changes. The added effort of carrying around weight affects perception. That is just a start. The space in front of a car is affected by the maneuverability, power, and speed of the car. Gibson called this "a sort of tongue protruding forward along the road" (Gibson and Crooks 1938). It is something like the safe operating envelope, the stable handling region, "the field of possible paths which the car may take unimpeded" (Gibson and Crooks 1938). By parity of reasoning we would predict that warehouse staff wearing roller-skates will judge the length of inventory shelves to be shorter, as they speed down aisles looking at the way things can be picked up. Downhill skiers will view the traversability of the terrain differently when wearing skis than when wearing boots, and surfers will view waves differently depending on whether they are on a short or a long board. In all these cases, equipment affects how things are seen because how we act on the world, and the tasks we perform, shape how we perceive.

In the Gibsonion approach to perception (Gibson 1966) the world to be perceived is defined relative to the action repertoire of a perceiver $A\{a_1, a_2, \ldots a_n\}$. Change the repertoire and you change the mode of interaction by which the perceptual world is partly constituted. With a tool, the action repertoire is increased to include tool-enabled actions, so there ought to be new affordances to perceive. Remarkably, Gibson wrote next to nothing on the effect of tools on perception or the relation between tool and affordance.[20] This points to a tension in the classical Gibsonian position. Holding a hammer or carrying a lit cigarette is not a function of untutored human bodies. These behaviors are not in our native action repertoire, our culture free repertoire. But they are natural in an artifactual world, the real world we inhabit. They have consequences Gibson would have appreciated. For instance, as most of us have unfortunately observed, a person who smokes cigarettes will see most physi-

---

[20] See Jones (2003) where the word "tool" does not appear in his discussion of Gibson's evolving conception of affordance over his lifetime.

cal environments as filled with things and areas that afford catching ash, things that can serve as ashtrays. Nonsmokers are blind to them. A stonemason will look at bricks for places to apply cement; when looking at an odd brick he will "see" the particular trowel shape that is needed. A competent tool user may perceive the affordances brought into existence by her use of tools, even when those tools are not in her hands!

Skill is a factor too. A person's skill in using a tool partly determines the conditions in which it can be used successfully. An expert carpenter can use a chisel effectively in more situations than a novice. Accordingly, skill affects what an agent will see in a given situation; skilled tool users detect more tool-relevant features, tool-related affordances, than lesser-skilled users.

2.5. Goals Make Perception Enactive

Goals also figure in perception. This view moves us beyond Gibsonian exegesis to a more enactive paradigm (Varela et al. 1991). The enactive account of perception (Myin and O'Regan 2008; Noë 2005) starts from the Gibsonian insight that perception is active and based on the possibilities of interaction, but it then adds three more things: interests, attention, and phenomenology. These lead to a conception of an environment that is both more and less than Gibson assumed.

When something grabs our attention we often fail to notice things that are visually obvious. This is called attention blindness (Simons 2000). In a famous example, subjects failed to notice a person in a gorilla suit a few feet in front of them because they were concentrating on whether a basketball was being passed legally. They were so focused on the ball they ignored the hairy legs and hands, and the mask.

We also overlook elements in full view when we are distracted by a major change, especially if the "in your face" change occurs simultaneously with the other changes. This is called change blindness (Rensink 2002). Jointly, the effect of this dual blindness is that the world we experience is a tiny fraction of what is there to be perceived. Like a hyperbolic visualization, we exaggerate the parts we are interested in and remain unaware of parts that hold no interest. Because the tools we carry are usually related to our goals and activities, indirectly they shape attention and interest. They narrow and expand our view hyperbolically.

At the same time, though, when we see something, we don't just see what our eyes have taken in; we factor in predictions about what we expect to take in if we continue to look around. Phenomenologically, we experience more of the world than there often is. For instance, when people look at Andy Warhol's Wall of Marilyns they do not saccade to every print of Marilyn (Noë

2005). They look at a few, perhaps examine some quite closely, and peripherally register the rest. Yet their experience is of a complete wall of Marilyns. Somehow their current perceptual experience includes the counterfactual beliefs of what they would see were they to look at each and every print closely.[21]

2.6. Enactive Landscape

Let us introduce the idea of an enactive landscape as the structure that an agent cocreates with the world when he or she acts in a goal-oriented manner. An enactive landscape is meant to capture the goalor activity-dependent nature of the perceptual world. It is the merger of a few ideas: task environment – the states and actions that are related to the achieving the goals and interests of the agent, the broader set of outside things or properties that can be acted on by that agent, and the full range of properties that agent can discriminate. The idea of an enactive landscape is a useful concept for designers to bear in mind when inventing new tools or systems because when a person has a tool in his hands his reshape their enactive landscape: they perceive more things and properties when working with a tool than they would when working unaided. In a sense, designers create enactive landscapes by designing tools.

Take the case of musical instruments. Apart from voice or clapping, music comes into being because musicians use musical instruments. No music makers and no musical instruments then no music.[22] Musical instruments provide the basic physical landscape a musician encounters. But the more skillful the musician, the larger the enactive landscape they inhabit, because skills combine with instruments to constitute a bigger world of possibilities. Music, conceptualized as this bigger world of instrument created possibilities, is an ex-

---

[21] This approach is worth putting in computational terms. To capture the idea that our counterfactual expectations are already factored into our experience we can represent perceptual experience as the current state of a predictive system, a broad-branched Markhov system of some sort, or a predictive state representation. Each branch, each path, represents an action that might be taken: a saccade to the far image, a step to the right and glance forward, and so on. Attached to each action is a probability of what one would likely see or feel. The predictive system should be further constrained by adding biases on the probability of actions that are a function of the goals and interests of the agent. Thus, an art historian, because of his interests, might be more likely to approach etchings closely to examine the printing technique than a casual observer. Or, returning to a cigarette smoker, because her cigarette-related interests are strongly activated when in the act of smoking, she is even more likely to look around for ashtray-like things. This constructed counterfactual space, with a probability distribution over outcomes that factors in the likelihood of a particular person acting in a certain way, is what defines the current state of a perceiver and what determines her experience.

[22] We ignore the Platonist claim that music is part of an ideal realm on par with numbers and other mathematical objects independent of construction.

treme instance of an enactive landscape. To a musician it is the set of possibilities that because of their instrument they can bring into being. An enactive landscape, then, is the set of possibilities that can in principle be brought into being when an agent interacts with an underlying environment while engaged in a task or pursuing a goal.

To complete this picture we need to remember that much of our environment is defined by rules and cultural constraints. Chess, sports and other games depend more on their rules, then on physical things like boards, playing areas, pieces and equipment. Rules and cultural influences mean that the same physical kitchen can constitute many cooking landscapes. The enactive landscape of a cook emerges from the interplay of a cook's interests and the cultural resources – such as recipes, food and taste preferences – with the physical things present – the ingredients, pots and pans, heat and layout of the kitchen. Each chef's vision is primed to notice the details of their physical space as it relates to their current recipe and their cooking style. (Almeida et al., 2008). In fact, looked at more closely, at each moment what a chef sees is partly primed by the tools in their hand. They see the things they might cut when they have a knife in their hands, the places to lay a dirty spatula when they are holding a spatula and so on. The same tunnel vision will apply during medical surgery. We are always primed to see the elements we expect to see as we precede in a task or an activity. (Endsley 1995). This means that the probability distribution that weights the possibilities present in an enactive landscape, will dynamically change as the agent shifts around the goal and subgoal structure of his or her task.

Given, further, that we all multitask during most of our waking life, the actual environment we live in, must be a superposition of dozens of enactive landscapes, each one with its own set of prediction generating elements and attention drawing features, rising and falling with our shifting interests. (Kirsh 2005). In designing a workplace, then, skill resides in blending the many enactive landscapes of its probable inhabitants to minimize error, maximize effectiveness, reduce fatigue and delight aesthetic sensibilities. Understanding the role of tools in shaping these enactive landscapes is a first step. The second step is to understand how co-creation evolves. Embodied cognition offers us new conceptual tools to analyze agent environment interaction.

2.7. Tools Change Our Conceptions

The final way tools change how we engage the world is by reshaping our conception of what is present and what is possible, not just our perception. An agent's immersion in an enactive landscape inevitably leads to concept formation. We are learning engines. Most of the concepts we learn are highly situated and ad hoc (Barsalou 1983). They arise as meaningful elements in the

activity that cocreates an enactive landscape, but may not have obvious natural generalizations. For instance, the way we perceive a beer bottle as we struggle to open it will typically give rise to the concept of trying-to-twista-cap-off. The phenomenon (the trying process) and the concept (the idea of what we are trying to do), are embedded in the cap opening activity. The idea of cap-twisting may eventually be generalized beyond beer bottles to other domains and tasks, losing its ad hoc status. But it started out highly situated in the specifics of beer bottles. When we use tools we multiply our ad hoc concepts because they multiply our enactive landscapes.

Opening beer bottles is typical of everyday tasks. Every task has its ad hoc concepts: washing our hands (ad hoc concept: the idiosyncratic way we each use hand soap), putting on socks (ad hoc concept: how we arrange each sock before slipping our foot in), sitting down (ad hoc concept: the way we stick our bottom out as we bend our knees). In each task there are task-specific things that represent points of learning or indicators of mastery. The hallmark of an ad hoc concept is that there is an attendable something, a potentially meaningful attribute that can be identified, attended to, referred to at the time (at least in thought), that is revealed in the performance of the task. Not everyone will have the same ad hoc concepts, but in any task there are always many things we must attend to and which can become objects of thought. Some are the affordances in the environment, others are the actions we perform, or the special way we perform them.

What does this unending, and potentially idiosyncratic, production of ad hoc concepts mean to designers? It tells us that design is never finished and never truly universal. When agents have an ad hoc concept they are in a position to think explicitly about their situation reflectively. For instance, TV watchers often surf between channels. Channel surfing is an emergent behavior that, once recognized, can drive the desire for change. Without the concept it is unlikely that anyone would identify the standard hassles with channel surfing. For instance, who has not had the irritation of switching from one channel because of a commercial, only to return to it after the program has restarted? This hassle, that is, the difficulty of timing when a commercial has finished – constitutes a design opportunity. In some TV's this need is met by a picture-in-picture feature that permits watchers to monitor the last channel while surfing, then toggling immediately back. The concept of channel surfing is typical of the cycle of how design gives rise to new and emergent behaviors that in turn give rise to new designs. It highlights how learning in our built-up world is continuous and how enactive landscapes are both personal and evolving. This year's cost structure incorrectly measures next year's interactions as learning changes our behavior and cost benefit function (Kirsh 2010).

3. RETHINKING THE ROLE OUR BODY PLAYS IN COGNITION

So far we have discussed how our tools and bodies are used to achieve pragmatic goals. Bodies and tools can be used for nonpragmatic goals as well. Professional dancers, when practicing, use their bodies nonpragmatically for epistemic and "cognitive" purposes— specifically as a means to physically model things. The same may sometimes be true for gestures (Goldin-Meadow 2005; Goldin-Meadow Beilock 2010) and for many of the things we manipulate. We think with them. Manipulating a physical thing is, at times, a method for driving thought forward. In this part we provide empirical support for this claim and speculate on why it is true.

3.1. An Experiment with Superexpert Dancers

The data to be reported comes from a single experiment undertaken in 2010 to test the effectiveness of different ways of practicing a new dance phrase. It is part of a much more comprehensive cognitive ethnographic study exploring embodied and distributed cognition in dance creation. See Kirsh et al. (2009), Kirsh (2012a, 2012b), and Kirsh et al. (2012) for a description of that larger project. In this experiment we found that partially modeling a dance phrase by marking the phrase, as it is called in the dance world, is a better method of practicing than working on the complete phrase, that is, practicing full-out. We also found that both marking and full-out practice are better methods of practicing than repeated mental simulation, a process found effective in other activities. (see Kossylyn and Moulton 2009). This last result is intuitive: it is better to practice physically than solely in one's head. But the first result, that partial modeling—a form of practicing a dance phrase aspect-by-aspect—can at times be better than trying to duplicate the perfect dance phraseS is a surprising result. Its explanation brings us closer to appreciating how physical activity—with body or tools—can help drive thought. Our results also suggest that prior work on learning by observation and learning by mental practice may not scale up to complex movements. Externalizing thought processes improves or reshapes inner processes.
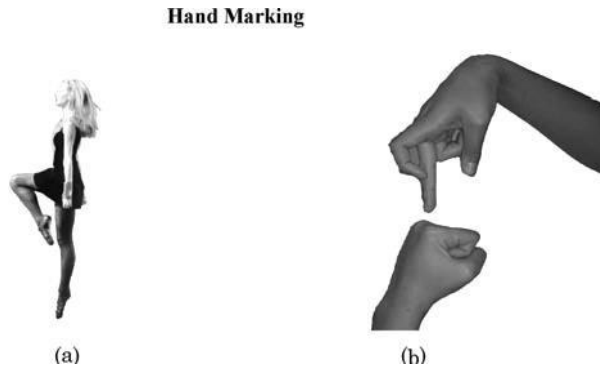
Hand Marking

Fig. 1. (a) An Irish river dancer is caught in mid move; (b) the same move is marked using just the hands. River dancing is a type of step dancing where the arms are kept still. Typically, river dancers mark steps and positions using one hand for the movement and the other for the floor. Most marking involves modeling phrases with the whole body, and not just the hands.

3.2. What Is Marking?

As discussed briefly in the Introduction, marking refers to dancing a phrase in a less than complete manner. See Figure 1 for an example of hand marking, a form that is far smaller than the more typical method of marking that involves modeling a phrase with the whole body. Marking is part of the practice of dance, pervasive in all phases: whether creation, practice, rehearsal, or reflection. Virtually all English-speaking dancers know the term, though few, if any, scholarly articles exist that describe the process or give instructions on how to do it.[23]

When dancers mark a phrase, they use their body's movement and structural form as a support structure for imagining the real thing, or perhaps as a representational vehicle pointing to the real thing or some aspect of it. The key feature is that they do not recreate the full dance phrase they normally perform; instead, they create a simplified or abstracted version—a model, a 3D sketch.

The received wisdom is that dancers mark to save energy, to avoid strenuous movement such as jumps, and to practice without exhausting themselves emotionally. But when they mark they often report that they are working in a special way, such as reviewing or exploring specific aspects of a phrase, its tempo, movement sequence, or underlying intention, and that by marking they can do this review without the mental complexity involved in creating the phrase "full-out."[24]

---

[23] Search by professional librarians of dance in the U.K. and U.S. has yet to turn up scholarly articles on the practice of marking.

[24] These reports were gathered by the author during interviews with dancers in the Random Dance company, as part of this study on marking.

Marking, or the practice of creating a simplified version of a process—a personal model to work and think with—is found in countless activities beyond dance. Adults who play tennis, golf, or basketball can be seen running through a "practice" swing or shot for themselves, as if to prepare themselves for the real thing. Sometimes they even do this without a racket, club or ball. Cellists will sometimes practice passages on their arm, running through finger positions on their "right forearm held upright in front of the chest, as a substitute for the neck of the cello" (Potter 1980, page 109) in a manner reminiscent of an Irish river dancer hand marking a jig. No sound emerges. Theatrical performers, too, can often be seen muttering their lines, or executing "practice" moves before stepping out on stage. It is a standard activity in theater to do an "Italian runthrough"—a slang phrase for saying one's lines and moving about the stage extra fast when staging a play to clarify the timing and relative positions of the actors. All these cases are related to marking. The common element throughout is that people seem to prefer working with a simplified version of a procedure to practicing the full-out version. In a slightly different way, playing tennis or ping-pong on the Wii is substantially like marking the real thing.

3.3. Why this Matters to Designers

Much learning and training is based on full-out practice. Why is this the most efficient way to teach everything? If our results generalize, then procedures and skills, in particular, might be better taught by a process akin to marking, where we create little models of things, or use our own bodies to pantomime what we must do. This is a highly general idea that can become a focus of good design for the learning component of any device. Moreover, as an example of an understudied way that humans think, it opens up new approaches to designing things as different as tools for problem solving, recipes for cooking, or resources for smarter collaboration. We return to this shortly.

3.4. Why this Matters more Generally

The finding that marking is the best method of practicing challenges common sense and previous work on complex motor learning. It is common sense that practicing something the way it should be performed ought to be more effective than practicing it with intentional distortions, or with essential components missing. If that were not so then repeatedly drawing a face in caricature, or perhaps in some other distorted fashion, rather than drawing it realistically ought to lead eventually to drawing the face more realistically than doing one's best to draw it correctly each trial. Similarly, practicing tennis stokes without a ball, or by ignoring one's body position during impact, ought to lead to better shots at times than always practicing in proper form. Future experiments may show that both these marking-like methods are, in fact,

better forms of practice than always practicing in an undistorted, full way. There are well-known precedents. In music performance, for example, using exaggeration in rehearsal is thought to be a helpful method of practicing, delivering results that surpass repeated full-out play (Hinz 2008). Players often practice one aspect of a passage—its fingering, rhythm, or bowing, while neglecting intonation or tonality (Stern and Patok 2001). Evidently, marking may already have a valued place in training.[25] But as a general method, practicing only distorted versions of the real thing, or versions that leave out essential components, is a counterintuitive method of rehearsal. Our unanticipated result is that this counterintuitive method can be effective.

Our findings also challenge recent work on dance learning. In several experiments (Cross et al. 2009), found that repeated exposure to a target phrase—and hence "practice" by mental simulation in the motor resonance system—leads to comparable performance to full-out physical practice. This unexpected result was found to hold for learning the rhythm and steps for pieces in a game like Dance Dance Revolution (DDR), where subjects must stamp their right or left foot onto footprints on a mat in time with music. Subjects watched the video repeatedly and may have played covertly. In our experiment, the phrases to be mastered were far more complex than DDR, involving movement of the entire body, with dynamics and feeling. When confronted with these more complex phrases we found that dancers benefited far more from marking and full-out practice than simulation. This suggests that moving the body in a controlled manner, even if not close in form to the target movement, can facilitate performance.

If our results about marking are true then marking during dance practice should not be seen as a sign of fatigue or laziness, as it so often is in dance studios. Rather, it may be a strategic method for selective training. This opens the door to developing more effective methods of selectively working on "aspects" of a phrase. We speculate that the success of marking also tells us something about how the body itself can be used to help manage attention, improve focus, and even facilitate simulation in a selective way. The body may well draw attention to what is important in an activity in the way a hand in speed-reading drags the eyes to help reading.

---

[25] Marking does not have an acknowledged value as a form of practice in dance despite its universality. Choreographers and dancers recognize that they cannot always practice the full form or a movement. But marking is thought to be a distant second best method. (oral communication by Wayne McGregor, and other professional dancers)

3.5. Conjecture and Method

When designing the experiment, our conjecture was the following.

(1) Practicing a dance phrase full-out would be better than mental simulation,

(2) Marking would lie somewhere in the middle: better than mental simulation but worse than full-out.

Owing to the power of the motor resonance system we wanted to see if anything would be gained by adding body activity to the mental simulation and projection we thought occurred during marking anyway. Our belief was that dancers would learn something from marking, just not as much as from practicing full-out. To test this idea we used the dancers from Random Dance, the contemporary company we have been studying (Kirsh et al. 2009). All these dancers are superexperts, chosen from an audition pool of 800 professional dancers throughout Europe and the States.

3.6. Procedure

The design required dividing the ten dancers in Random Dance into three groups: A, B, All three groups were brought into the studio and taught a new dance phrase lasting about 55 seconds. The teaching phase lasted 10 minutes. At the end of it, the group left the studio and the dancers returned, one by one, to the studio and performed the dance in front of the teacher, who graded them to set their baseline. As shown in Figure 2 each group, A, B, C practiced in one of three conditions: full-out, marking, and lying on their back using mental simulation. They were then individually regarded. After the first round the dancers swapped practice conditions and were taught a second phrase of about the same duration and complexity as the first.

Each dancer's performance was graded according to established criteria— technicality, memory, timing, and dynamics—first by the teacher in real time and later by two independent expert observers who reviewed the video frame by frame. Once all dancers were graded, the group returned to the same large studio and practiced the dance for 10 minutes. When practicing they faced in different directions and told not look at each other. Once this 10-minute practice period was over they left the studio and, as before, returned one by one to be graded by the same criteria as before. See Figure 3.

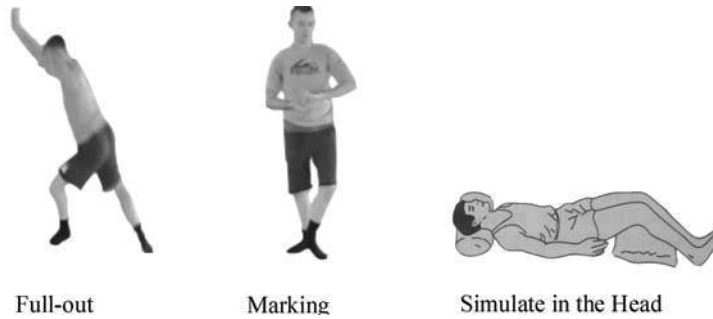Full-out   Marking   Simulate in the Head

Fig. 2. Experimental conditions. Subjects practiced mastering a phrase in one of three conditions. They marked the phrase, practiced it full out, or lay on their back and mentally simulated dancing the phrase. After being evaluated they had a five-minute rest, changed conditions, and were then taught a new phrase. In this way all subjects practiced in each condition.

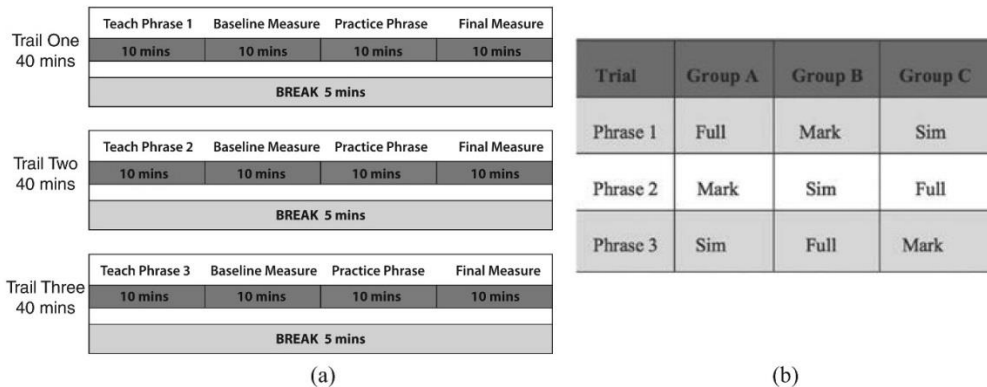| Trial	Group A	Group B	Group C
Phrase 1	Full	Mark	Sim
Phrase 2	Mark	Sim	Full
Phrase 3	Sim	Full	Mark

(a)   (b)

Fig. 3. (a) The temporal structure of the experiment is displayed. After a 10-min. teaching phase subjects are evaluated, then they practice, then they are evaluated again. Learning is understood as the change in grade acquired during the 10-min. practice phase. (b) the experimental design, a 3 by 3 Latin Square, is shown.

3.7. Measures

Technicality. This means the level of precision found in positions and transitions on a five-point scale, in increments of .5. How structurally correct is the position? When a transition is the object of interest, its structural aspect can be assessed along a technicality dimension too. Other elements of accuracy, such as the phrase's dynamic fidelity, are evaluated in the measure on dynamics.

Memory. Memory, or level of detail, refers to the completeness of the movement. Does the movement display all the elements at each level in the hierarchy of detail?

Timing. This refers to the level of precision in the duration of individual steps and the duration of transitions. To code timing, coders used frame-by-frame measures for great precision in comparing test conditions to their normative standard.

Table II.
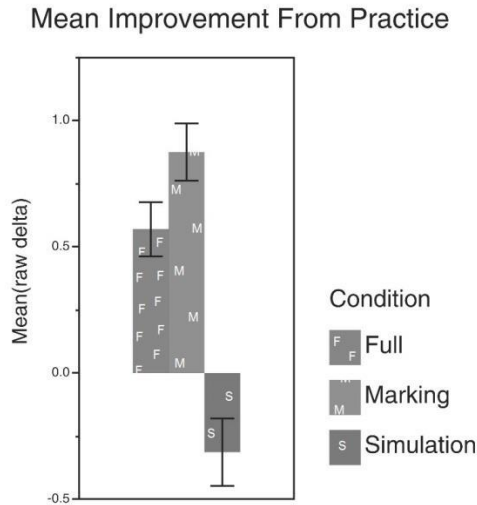
Dynamics. This refer to the force, speed, and acceleration of movements. Various qualities of motion such as resistance, juiciness, roundness, emotionality, and intentionality are also included in the category of dynamics.

3.8. Results

Our analysis of results showed the following.

(1) Marking is the most effective overall method of practicing, being slightly more learning efficient than practicing full-out across the key dimensions of Memory, Technique, and Timing (mean difference = .31; p = .0189). In dynamics, however, full-out is better.

(2) Both marking and full-out lead to substantially more learning than mental simulation across all dimensions (mean difference = 1.19; p = .0001).

(3) Mental simulation is not a strong form of practice; there was negligible learning and in many cases practice by mental simulation led to a decrease in performance.

Table II shows the mean improvement from practice (the learning delta) as measured on a 5-point scale. Improvement was best for marking, less for full-out and negative for mental simulation. The absolute difference in delta be-

tween marking and full-out is 0.31, which is significant when measured by the z-score for Technicality, Memory, and Timing (p = .0189). Full is better for Dynamics but not significantly so (p = .145). All p values were computed over z-scores to reduce noise caused by variability in dancers, measure types and graders.

Table III shows that marking is significantly better than full-out for learning the aspects of a phrase related to technicality and memory. Not surprisingly it is less effective at learning dynamics, which are rarely practiced in marking. Mental simulation was most effective for thinking about technical elements (precision in movement). It led to decreased performance, that is, negative learning, for movement details.

To compute these values we first performed one-way ANOVAs on all measures in all conditions and found highly significant differences throughout. We then ran pairwise post hoc comparisons (Tukey's HSD) and computed p values as shown in Table IV.

4. THEORETICAL IDEAS THAT MIGHT EXPLAIN WHY MARKING IS SO EFFECTIVE

What might explain why marking facilitates mental simulation? And what might explain why marking is better than full-out practice? The explanation I offer highlights a general process that, I believe, applies more broadly than just to dancing, to practicing skills and to thinking with the body. The explanatory principle proposed applies also to the process by which we think with tools and everyday objects. Marking is merely an activity where it is especially easy to notice physical thinking.

When dancers mark they project beyond what they are actually doing to a more ideal movement. We found qualitative support for this idea from interviews with the dancers. When asked what they think about when marking they reported that they have in mind the full-out movement, though with fewer dynamics. They do not "see" themselves dancing in a distorted way, as they would if observing themselves in the studio mirror while marking, but rather kinesthetically from the inside, feeling the key aspects of the movement. They seem to be projecting kinesthetically, and to some degree visually, from their marking movement to the "correct" or normative movement. This correct movement is what they have in their mind's eye. Marking is just the external support or scaffold that helps them have the right mental imagery in mind.
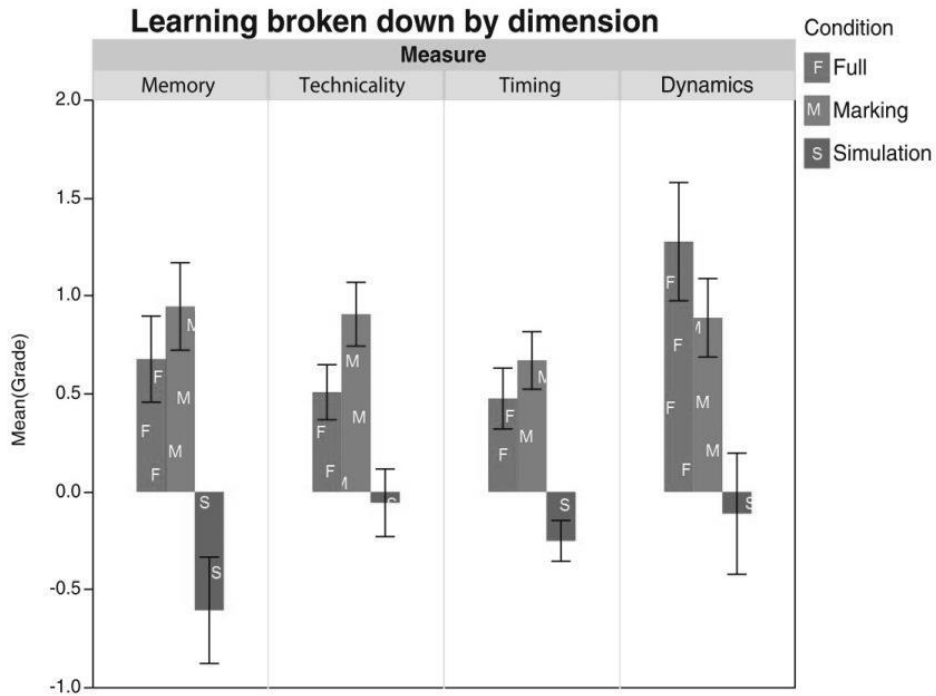
Table III.

Table IV. P Values Showing the Significance of Findings.

Measure	Mark	Full>Mar	Mark>Si	Full>Si
Memory	.7334		<.0001	<.0001
Technicality	.0029		<.0001	.0005
Timing	.0194		<.0001	<.0001
Dynamics	-	.145	.0003	<.0001
Mem, Tech,	.0189	-	<.0001	<.0001

To explain how an imperfect model of a movement—which is what marking literally is—can behave as a physical support we need to introduce a few ideas. We begin with projection and anchoring.

Projection is a mental process akin to attaching a mental image to a physical structure. When we project onto an object, whether kinesthetically or visually, we experience ourselves intentionally augmenting the object. The object *anchors* our mental image, and successful projection requires spatially or temporally locking the projected image onto the anchoring structure. In the case of visual projection, the easiest form of projection for sighted people to understand, the image to be attached must be the right size and be connected to a specific location on the external structure.

When we imagine an object, we again are dealing with mental images but we do not attach them to anything in the external world. Imagination has no physical anchor, and imagined images need have no specific size or location.

By contrast, when we perceive an object, we are not imagining or projecting anything. Our experience is of an external object or scene that is supposed to be really there. In veridical perception it is. Perception produces the highest-resolution experience and varies far less among fully sighted subjects than does the vividness of imagination and projection, which varies greatly.
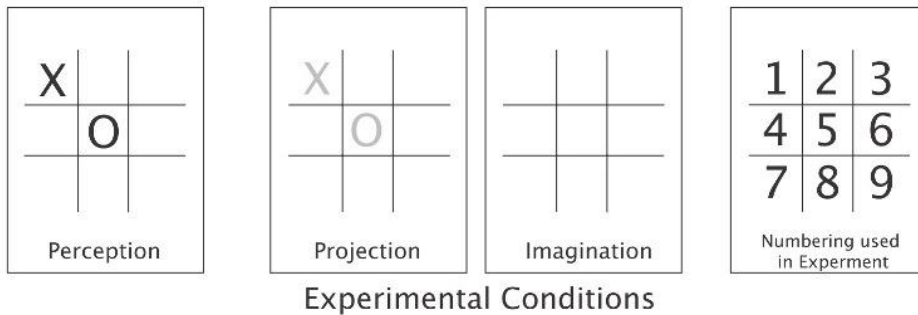

Fig. 4. The differences between perception, projection, and imagination can be understood as that between seeing the X and O marks on a grid, projecting an image of X and O onto a blank grid, and imagining X and O on a blank sheet of paper, or forming a mental image of the board and marks while blindfolded. Subjects in the experiment described shortly were tested in the projection and imagination conditions only. They played by calling out moves using the cell numbering system shown.

As shown in Figure 4, the difference between projection, imagination, and perception can be represented by three conditions in which subjects might play a tic-tac-toe game, the domain in which we experimentally explored the idea of projection and imagination.

—*Perception*. Subjects see the actual inscriptions of X and O on a board. Games are played in the ordinary way, by making marks. The complete state of the game is explicitly represented by the placement of X's and O's and visually evolves with play. No memory of past moves is required since it is on display.

—*Projection*. Subjects see only a blank tic-tac-toe grid and have to mentally augment the grid with moves. The grid never changes no matter how many moves are taken. Everything has to be remembered by the subject, but the grid might help structure or support visual recall. Projection is like augmenting reality.

—*Imagination*. Subjects see a blank page. To play the game they must imagine all moves. There is no grid to help support or scaffold their visual recall.

They can play the game in this condition either blindfolded or looking at a blank piece of paper. Imagination, at its best, is like creating a virtual reality game.

In Kirsh (2009b), the results of running 24 subjects playing tic-tac-toe in the projection and imagination conditions were reported. There was no perception condition since it was assumed that if subjects recorded their moves they would perform at ceiling. To play the game, all subjects first learned to name cells using a 1–to–9 numbering system, as shown in Figure 4 for the 3× 3 board, and a 1–to–16 system for the 4× 4 board. To play the game they called out their move after hearing their opponent's. To test for visual imagination capacity we administered the standard vividness of visual imagery test (VVIQ II) beforehand.

The results were not simple. Subjects did not play tic-tac-toe with better speed accuracy in any 3× 3 condition. This was not what we predicted. Having a grid to anchor projection did nothing in the 3× 3 game where one rarely needed to recall more than 5 or 6 moves.

Table IV.
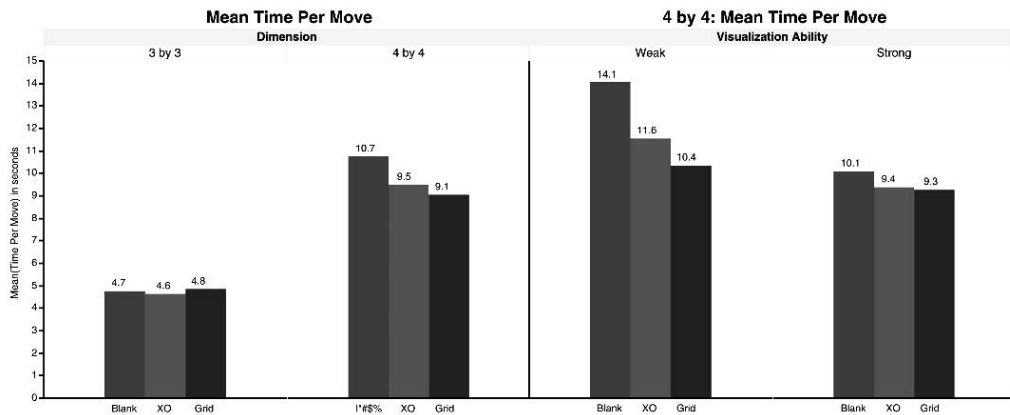
In 3× 3 games there were no significant differences between conditions. In 4× 4 games mean performance was significantly better in the grid condition than in blank (mean difference = 1.6s, p = .002). When subjects were binned into weak and strong visualizers it was clear that strong visualizers benefited less from a grid to project to. Strong visualizers grid > blank = .8s, p = .625; weak visualizers grid > blank = 3.7s, p = .014.

To challenge the subjects, we then taught them to play 4× 4 games. In this condition the visual memory load is much greater and we found that having a grid appears to facilitate all subjects, but the effect is strongest among subjects with lower visualization capacity. As predicted, the grid now appears to serve as an understructure or scaffold for projecting moves. The effect was divided, though. For weak visualizers the grid strongly enhanced performance.

By contrast, strong visualizers only trended to perform better with the grid. As shown in Table IV, in 4 × 4 games there is an interaction between mental imagery ability and the usefulness of the grid. Evidently, if a subject can play the game well in her imagination she gains little from projecting to a blank grid. Although our experiment should only be seen as a pilot study (n = 24 overall, with n = 6 strong visualizers, n = 6 weak visualizers), the implication, we believe, is that at some point, when a spatial memory task becomes hard enough, everyone benefits from external structure. Hence we predict that even good visualizers will benefit from a grid if they play tic-tac-toe on 5× 5 boards.

The relevance of this finding to marking is that if a movement is easy to learn then marking may not help dancers "visualize" the movement any more accurately than mental simulation. Only if a movement is hard to learn would we predict that marking facilitates projection by enabling dancers to bring to mind movements that are more detailed or precise than they can mentally simulate while lying down. This is consistent with the findings of Cross et al. (2009) and Williams and Gribble (2012). Both found that subjects who simply observed other people performing a dance movement were able to learn that movement comparably to those practicing the movement fullout themselves. Notably, the movements they studied were simple steps forward or sideways, unlike the dance phrases we studied, and the movements were learned in response to an on screen prompt. This suggests their learning task was easy and the memory requirements much weaker than either tic-tac-toe or learning whole dance phrases.

4.1. Marking as a Mechanism to Support Projection

The conjecture we have offered is that marking is a better way for dancers to practice than mental simulation because the act of creating an external movement provides a physical anchor for the dancers to project their full movement onto. This physical anchor carries some of the weight of imagination and helps dancers to think more effectively about aspects of their movement that they are trying to improve, recall, or practice. When dancers mark they sketch a movement schematically but have in mind the real or ideal thing. They accomplish this feat either by projecting on top of marked movements or by using marking as a mediating structure to facilitate mental simulation in some as yet unexplained way.

There is little hard proof we can provide that establishes that marking supports projection more than by analogy with our tic-tac-toe experiment at present. There are some further analogies to be found occurring in natural contexts, however, that may show the prevalence of this sort of phenomena.

In music practice keeping the beat by tapping one's foot is a plausible analog to using a blank tic-tac-toe board. When musicians tap, arguably, their tapping serves as an anchor for projecting the musical rhythm to be performed. Musicians report that when they tap it is to provide them with a stable pulse to help them stay on track while thinking about the musical rhythm they must play. Tapping serves as a support structure because it is thought to be caused by an internal oscillator (Eck et al. 2000). that is sufficiently automatic to liberate higher motor planning centers to work on different, but coordinated, sorts of actions: the rhythmically more complicated musical rhythms that are played "on top of the beat". Because the oscillator is autonomous it behaves like an external resource, like a tic-tac-toe grid, that can be leaned on. For complex rhythms it may help a performer to stay in time. Foot tapping may be a dynamic analog in the sound domain to visual projection in the visible domain.

Orchestral practice provides another example of dynamic projection, this time tied more directly to marking. When a conductor gestures to his ensemble, carrying the beat with the dynamics of arm and baton, he is doing more than simply embodying or displaying beat. He is marking the piece. His arm and gestural dynamics direct each musician to attend to specific musical features: crescendo, counterpart, entry moments, etc. These movements cue performance. But they obviously are not the same full-out thing that performers do. The conductor does not play any notes. He does not use a bow or blow or strike drums. All the information encoded in his gestures is sparse and aspectival. But by marking the music he provides performers something they can work off of. The conductor's animation anchors projection; it anchors performance. Conducting is so natural and useful we can readily believe that players themselves might run through their own part conductor-like, stressing aspects of their music, using gesture and voice rather than through full-out performance on their instrument. A final case of gesture anchoring projection, this time a purely internal form of projection, is found in studies of mental abacus. Frank and Barner (2012) studied elementary students in Gujarat, India. The students were taught to add and multiply using an abacus and, once proficient, were then asked to perform calculations without using an abacus. The practice is known as mental abacus because the students still use an abacus to calculate with, though it is an imagined abacus and their manipulations of it are also imagined. When students work on their mental abacus, however, they almost always flick their fingers, marking-like, partially miming the action of moving the beads. When they are not permitted to use their hands, their performance suffers (Frank and Barner 2012, Hatano et al. 1977). This suggests that gesturing is not simply an epiphenomenon, that is, an unnecessary accompaniment to the mental operations that do the real work. Hand motions seem to improve mental simulation. The mechanism is possibly

much like marking in dance. By imperfectly simulating moving beads on a physical abacus, the students help to create and transform mental structures. They project from their own gestures.

4.2. Marking as a Trick for Directing Attention to Aspects of a Movement

Although our dominant hypothesis about marking is that it supports projection, it may also perform a further function: it may help subjects manage their attention practice. Marking "keeps the dancers honest". It helps them to focus on aspects[26] of their movement one by one, methodically practicing aspects that would otherwise be easy to overlook when simulating the movement mentally.

For example, we found that in the mental simulation condition dancers were especially bad at remembering details, that is, knowing what to do at a detailed level with fingers, head, and feet. Marking might help because if it functions as a type of interactive sketching with the body, then whenever they mark the dancers will have to keep in mind the target body part(s) they are sketching. By focusing on specific parts they may avoid ignoring details.

This tendency to fixate attention on details of a movement may also explain why marking is better than full-out practice. When dancers practice full-out they execute all aspects of a movement at once. It is not possible to work on timing while ignoring the shape of the movements, or work on spatial extension while ignoring dynamics unless these other components can be run on auto-pilot while the dancer thinks about a different aspect. This is bound to be difficult because there are interactions between aspects that make it hard to ignore most aspects when danced at once. When marking, by contrast, dancers are allowed to practice piecemeal, working on this part or that, this aspect or that aspect, in a manner that is guaranteed to be unaffected by other elements precisely because the other elements are not being executed at the same time. The idea that performing an external action may be a cognitive strategy for helping to manage attention has been discussed before. In Kirsh (1995), and later in Carlson et al. (2007) experiments

---

[26] A dancer may fixate on any part or attribute of a movement. Laban movement analysis codifies these aspects into the major categories of body, effort, shape, space. Body – which parts are moving, connected or influenced by others, total-body organization; Space – motion in connection with the environment, and with spatial patterns, pathways, and lines of spatial tension; Effort – dynamics, qualitative use of energy, texture, color, emotions, inner attitude, often reduced to float, thrust, glide, slash, dab, wring, flick, and press; Shape has static forms: pin-like, ball-like, wall-like, pyramid-like, screw-like, it has flow forms depicting how the body changes shape during movement, and it has shape qualities: rising/sinking, spreading/enclosing, advancing/retreating. More macro relations include phrasing and relationships. See Konie, (2011). All these can be objects of attention, focal points of thought.

were discussed in which subjects were asked to count images of items laid out on a page. In Kirsh (1995) the items were nickels, dimes and quarters, and the task was to provide the total dollar value present; in Carlson et al. (2007) the items were asterisks and the task was to count the total number present. Both studies found that finger pointing leads to improved performance. In Carlson et al.'s study, it was observed further that head nodding occurred and led to similar improvement. One reason pointing is useful is that it might help a subject keep track of what was last counted. Our fingers, unlike our eyes, stay put unless intentionally moved. Eyes saccade relentlessly. Multiple fingers, moreover, can keep track of multiple locations, the last nickel, dime, and quarter counted. The relevance of such studies is that they remind us that bodies have different properties than minds. They move differently, they obey a different principle of inertia, and so they are a resource that can be harnessed to help solve complex problems. In the case of dance, pointing and dance they are useful for helping to manage attention.

5. WHY INTERACTING IS BETTER THAN OBSERVING

A few years ago it would have seemed unnecessary to ask whether practice is necessary for mastery: who would think that watching can substitute for doing? But three theories taken together now make that question worth asking: (a) the common coding theory (Prinz 1997), that is, the idea that motor and visual perception share a common worldoriented code; (b) motor resonance theory (Agnew et al., 2007), that is, the idea that during observation we activate a motor mirror system where we covertly do what we see; and (c) enactive perception (O'Regan and Noe 2001), that is, the idea that we cocreate our perceived environments, so perception is itself a form of action. Today the question is: When is observing as good as doing? Is it possible that a couch potato might learn as much about an activity by simply watching it on TV as doing it himself?

It is important for designers to know whether first-hand motor involvement matters as much as common sense claims. If, for instance, it turns out that we can learn to cook as well by watching the cooking channel as by real practice in the kitchen then designers thinking of interactive kitchens of the future will want to use video for teaching and direction rather than have instructions emanate directly from tools and surfaces. Alternatively, if working with imperfect models of things is often a better way to learn complex activity than is regular practice, and hence better than observation too, then it follows that embedding recipes in tangible things, providing cues for slicing in knives, hints in spatulas, will be the design of choice, because then tools will focus attention on aspects of a recipe in a timely way. It may even be worthwhile to embody recipes in toy models of pots, pans, and eggplants for a quick run through before full-out cooking. What applies to

kitchens applies to hundreds of environments where pervasive, context-aware computation will become the norm.

The main phenomena that we observed that calls into question the thesis that "observing is as good as acting" relates to the obvious idea that each of our senses picks up different information. Even if there is a common code and motor resonance is correct, it does not follow that visual perception drives as much covert motor activity as actual movement coupled with kinesthetic perception. Motor planning may be more responsive to kinesthetic factors. Vision provides only a fraction of the information needed to adapt interactively to rapidly changing forces in our environment or bodies. These need to be picked up kinesthetically.

To explore this question empirically we studied videos of how our choreographer works with dancers when creating new movements. One of his favorite methods is to watch as they solve a choreographic problem and then, when he sees something interesting, he physically sketches their movement himself. He doesn't just watch, he imitates. Then he modifies his sketch a few times and gives them back his own version. They engage in a physical dialog. But not always. It is quite clear the choreographer can make his ideas known using words, gestures, or sounds, rather than displays. So, presumably, on those occasions when he imitates movement it is because it gives him insight in a way that vision alone does not. See Kirsh (2012b), where this process of "riffing-off-of-others" is discussed.

Intuitively, it is clear why the choreographer physically appropriates a dancer's action. By performing the movement he can better appreciate the creative possibilities of a movement. The continuation structure he develops about how the movement might be carried forward—about where it might go creatively—is different when he is the author than when he is the observer. And for good reason.

In dance, some phenomenological attributes emerge only when the body is genuinely involved. Consider the experience of internal force, resistance, stretching-to-thepoint-where-it-hurts, or rotating-so-quickly-you-are-just-about-to-lose-control. These are phenomenologically prominent features that arise when we interact with objects or work with our bodies. But they are mostly invisible visually. We feel more than we show. And when as observers we see something subtle we may be misled. For one thing, people are different in their strength, flexibility, and pain thresholds. An easy stretch or lift for a dancer may be painful for me. So when I see what looks like a painful stretch, my sympathetic feeling, that is, my mirror pain, may be unreliable. Moreover, some body feelings are completely invisible. The feeling a person has just before falling or losing balance, for instance, or the amount of resistance

being imposed on an arm motion[27], all these are prominent kinesthetically but invisible behaviorally.

This difference in the sensory representation between watching and dancing is, I believe, one of the primary reasons the choreographer runs the movements of others through his own body. There is only so much he can know about a movement by looking, even with his motor resonance system running all out, and with his overdeveloped expert vision. This limitation is commonplace. For instance, when a potter works on the wheel, shaping a bowl, there are sensory attributes that can be felt by the hands that cannot be seen, such as the feeling when the clay is about to tear, or the feel of the texture of the clay. The potter may feel the "continuation space" of this pot, that is, the enactive landscape, more effectively and be better positioned to know whether s(he) can recover from an about-to-tear situation. The hands know things the eyes do not (Malafouris 2008; O'Connor 2006).

At a more fundamental level each sensory system supports different pathways of sensory expectation. Events that seem "natural" or obvious in one sensory system may seem unnatural or completely unobvious in another. Things that are predictable in one sense, such as, "if I move my arm any further it will hurt" are not predictable in another. Some things are easy to infer, others are not. Different senses make different attributes explicit, obvious, such as "I am in pain".

This is a universal property of representational systems: some properties are encoded explicitly; others are more implicit and must be computationally extracted. For instance, to decide whether the number 30,163 is divisible-by-7 takes some computation. The attribute divisible-by-7 is not as explicit in the base 10 as divisible-by-10 or being odd. We can instantly tell that 30,163 is not divisible-by-10, and that it is odd. In the base 7, however, 30,163 is represented as 153,640$_7$, and it is completely obvious that it is divisible-by-7. It is transparent and explicit. It is less explicit that it is odd. See Kirsh 1992, 2009d) for an analysis of explicit-implicit representation. Each sensory system has a coding language that represents some attributes explicitly and downplays others.

Returning to dance, a dancer may immediately recognize through his somato-sensory system what level of control is needed to execute a particularly complex move. Visually this may be unclear. In extreme cases, a movement that the motor system deems impossible may not seem impossible according to the visual system. The upshot is that when a choreographer considers how a movement might be continued, his conception derived from vision may be different than his conception derived from riffing. By riffing the movement, he

---

[27] Resistance is a technical term in dance to refer to the antagonism of muscles when they pull in opposite directions. Isometric exercises are extreme cases of muscle antagonism.

physically appropriates it, thereby activating a system of motor intuitions that are different than his vision-based motor intuitions.

One especially clear example of this is found in Wayne McGregor's own oeuvre. In the piece Ataxia, McGregor explored some of the movement space of ataxics, people who have imperfect motor feedback. The kinesthetic phenomenology of ataxics is unlike that of normal subjects. Dancers and McGregor himself worked with ataxics to get a sense of what it is like to have ataxia. This was not done by watching them alone. It was vital that they learned to move like ataxics, a process that took some time. By simulating an ataxic's body motion, however, both dancer and choreographer gained access to a different aesthetic. The narrative of the dance,that is, the movement vocabulary the dance was based on, was derived from familiarity with the kinesthetics of ataxia. Again, watching was not enough. Indeed, it is quite possible that much of the dance was "written" more in motor feelings than visual form. Body sense offered a different basis for aesthetics.

One consequence of our research on dance, then, is that bodies are involved in cognition in more ways than motor resonance and common coding. Riffing, in particular, shows that working across modalities can reshape the conceptualization of something (e.g., a dance phrase) beyond its origins in one modality, the seen modality. It shows how the body can figure in extending the range of thought; how our bodies can lead us to new ideas that are far from the sensory-specific ways we encountered these ideas originally.

This echoes a widely appreciated feature of human thought, namely, that when people interact with external artifacts—representations, instruments, toys—they are able to learn new things, or understand things more deeply. The physical act of sketching a familiar object can help us realize aspects of it that we never noticed or thought of before. Playing with a physical model of something first seen in the movies (a transformer toy, for example) can open our minds to some of its deeper properties. It is still the same object, and we still have the same basic concept of it, but we know more about it, or see it in a new way. Engaging with it physically helped us learn far more about it than just seeing someone else play with it. Making and working with models is good for advancing cognition.

6. EXTENDING EMBODIED COGNITION TO INCLUDE THINKING WITH THINGS

In earlier sections I discussed how tools can be absorbed into the body, how bodies can be used to model and simulate, and how running ideas through the body that were first encountered visually can lead to perceiving creative possibilities that otherwise were hidden from sight and mirror cognition. The next step is to show how things other than the body and tools can be har-

nessed and incorporated into the thinking process. It is to this I now turn. The burden of this section is to expand on the idea of thinking with things beyond its usual interpretation as computational off-loading, a topic that I believe has been well covered in discussions of situated cognition (Clark 1997; Kirsh 2009a) and external cognition (Clark 2008; Scaife and Rogers 1996).

My basic line is that thinking, in the sense of drawing inferences, can be done partly in the perceptuo-motor system and partly by manipulating external things in a manner that is tightly coordinated with inner processes. Manipulating external things, even when we do not appreciate it, is a form of simulation. When we interact with an object our interaction drives our perceptuo-motor system into a state of expectation, and we tacitly assign probabilities to outcomes conditional on what we might do ourselves. Thus, in the simplest case, if a subject were to start to upend an object, his action may begin to unocclude part of it, and he may infer something about the object's bottom, or its overall shape, as well as anticipate its appearance just before actually seeing the bottom. Few theorists would call this form of expectation and pattern completion a form of reasoning. But formally extrapolation is a type of inductive inference, and completion is a form of extrapolation or interpolation.

Suppose now that our subject were to do something manipulatively more interesting. He might, for example, hit the object on the edge of a counter or pour water into it. These actions were improbable a moment before, and the perceptual input they produce will have a major impact on the shape of the continuation tree, that is, on the lattice of possibilities representing what the subject expects could happen. Because of the nature of some actions, these expectations may have more to do with the causal mechanics of the object than its shape. An example of causal mechanics might be the effects of twisting a bottle cap. On the account I am recommending a subject who begins to twist a cap starts to simulate not just cap twisting but a twisting-off process. As he proceeds in twisting, his perceptuo-motor system predicts future outcomes, many of which represent a world where the cap falls completely off the bottle.

My hypothesis here is that when people think with things they rely on the world to simulate itself and in so doing they stimulate themselves. They indirectly control their epistemic state. The world executes part of the reasoning process, therefore, by carrying them to a new state that is reasoning relevant. Thus, as the thing they manipulate undergoes change, people revise their continuation system (and their enactive landscapes). This revision to their continuation space is equivalent to a change in their enactive perception of what is happening. But mathematically, it is identical to a form of induction, that is, of learning and reasoning. The result is that if we were to ask our cap twister whether he thinks that if he continues twisting he will re-

move the cap, his answer (yes) looks like it is the product of thought. And it is. But it is not propositional thought in a classical sense.

On a classical account (Harman 1986) every use of the cap is an experiment, a piece of data to be assimilated through induction, deduction, or abduction. A person has an internal hypothesis concerning how the world is and the effect that twisting will have on the cap. He tests the hypothesis by experimenting with the cap. Beliefs are revised and new opinions come to the fore. This whole process makes no reference to changes in perception, conception, or skill. Perception serves as a pass-through of the world state. The result is that a propositional representation of the world suffices, and reasoning becomes the same old classical idea of rules operating on propositional representations[28], something not far from the language of thought (Fodor 1975).

By contrast, when our thinking subject infers something he will often be reporting an expectation encoded in his continuation system. A lattice of continuations is not a propositional representation. Inner simulations of scenarios can reshape this continuation system. On those occasions where he cannot simulate the future well internally, or if there is too much uncertainty in how he thinks things will unfold, he can reach out and begin twisting the cap, that is, perform an outer simulation. Unlike Harman (1986), this is not equivalent to performing an experiment. It is to cause a revision in his continuation set. The new input alters the enactive landscape he cocreates.

The bottom line is that physical thinking is an external version of the idea in embodied cognition that much of our cognitive life depends on internal simulation of things. In the embodied literature, simulation is internal. I have argued here and elsewhere that we can expand this notion. If internal simulation counts as thinking why not also count external simulation as thinking (Kirsh 2010). And if that seems too extreme then at least we should allow the coordination of internal and external processes to be thinking. It is simply a matter of cost whether to simulate inside with images and ideas or to simulate outside with real things, but in a controlled manner.

I believe that developing this view—that literally we think with things—will have major implications for how designers come to understand interactive objects and systems.

---

[28] However, see Johnson-Laird (1989) for a model-based approach to propositional reasoning that relies on structural elements that need not themselves be components of a proposition.

7. CONCLUSION

HCI is at a crossroads. We are entering a new world of physical, natural, and tangible interfaces. We can interact with digital elements by gesturing and body movement, by manipulating everyday objects, and even by training brain activity to control interfaces. To understand the design principles of such a world requires that we become familiar with the ongoing developments in embodied, distributed, and situated cognition, and build closer relations to their research agenda.

In this article, we explored the idea of tool absorption and how our internal representation of personal and peripersonal space adapts to manipulables in our hands and wearables on our feet and body. There are open questions about how far we can extend our action repertoire through tele-presence and remote actuation, and how far we can push our perception beyond its normal semantics. Is it possible to control actuators that have dozens of separately articulable fingers in a wholly natural way, that is, absorb them into our body system? What is the mapping function between our own action repertoire and "tool" created repertoires? I introduced the notion of enactive landscapes to define the cocreated actionable environment that we perceive, Gibsonian style, in terms of what we can do. But unlike Gibson we included tool-supported actions as part of our action repertoire and hence ways we can alter our enactive landscape. To an agent, the world is a constellation of intersecting, overlapping enactive landscapes, engendered by the tools in hand and the resources nearby. When a tool is picked up or let go there is a change in capability that leads to a change in the enactive landscapes that are active. Because tools alter our action repertoire, they shape both our perception and conception of what is present to be acted on. With further development this idea may have useful application in understanding how digital interactivity will reshape our sense of what we can do.

I think it is fair to say that the view that tools modify our perception, conception, and even our bodies is one that the HCI community has accepted in one form or another for some time, though without adequate empirical and theoretical support. This support is starting to arrive. We also need theory and empirical support for a more modern conception of our bodies and what they are capable of doing. A consequence of our research on dance is that we have evidence that human bodies can be used for all sorts of cognitive purposes. In particular, humans use their bodies not just to act on the world and enact or cocreate a personal world, they use them to represent, model, and ultimately self-teach. They use their bodies as simulation systems, as modeling systems that make it possible to project to unseen things that would otherwise be more inaccessible. These unseen things may be dance phrases that are the target of learning, or they may be aspects of those phrases that need to be attended to in order to master the phrase. Dancers also use their

bodies to perform analog computation, since they can rely on the mechanical properties of their bodies to complete trajectories that otherwise would need to be planned and computed.

We learned further that dancers make good use of their different senses. For instance, kinesthetic perception reveals different properties than visual perception, and these kinesthetic properties, because of the way they are encoded, make it easier to recognize the validity of inferences that would be near impossible to infer from vision alone, if one did not also move the body. For our choreographer, for example, we found that by recruiting his kinesthetic sense he is able to "see" aesthetic properties and narrative properties of dance phrases that are unavailable through the visual spectrum. He uses vision to observe his dancers' work on phrases, but he runs these phrases through his own body to appropriate them and appreciate their choreographic possibilities.

Given the power of tools and bodies to extend thought it is natural to make the final step to other objects as things that people can use to think with. The hypothesis presented here is that much of human thinking takes place in the perceptuo-motor system or an extension of it. We interact with the world and in so doing we physically simulate outcomes, or begin to simulate processes that shape our internal expectations of how things may turn out. This revision of perceptuo-motor expectations is not done propositionally or consciously. It is a form of implicit cognition and buried deeply in our perceptual system. But it results in changes in how we mentally simulate the future. Thus, if someone were to ask us whether sitting on a chair distributes pressure over the chair's legs, a thoughtful person might run a mental simulation of sitting down or a mental simulation of reaching under a leg and feeling the weight of the chair. Alternatively, he might begin the reasoning process by actually sitting down and using the resulting changes in his continuation system to make mental simulation intuitive.

Once we understand the complex coordination between external and internal simulation, between doing things internally and doing things externally, we will begin to reach new heights in design, and create a cognitively better world of physical-digital coordination.

Acknowledgments

graphic process of Wayne McGregor | Random Dance, and to have been given great freedom to set up cameras, give interviews, and learn about the world of professional dance.

References

Aglioti, S. M., Cesari, P., Romani, M., and Urgesi, C. 2008. Action anticipation and motor resonance in elite basketball players. *Nature Neurosci. 11*, 9: 1109–1116.

Agnew, Z. K., Bhakoo, K. K., and Puri, B. K. 2007. The human mirror system: A theory of mind reading. *Brain Res. Rev. 54*, 2: 286–293.

Barsalou, L. 1983. Ad hoc categories. *Memory Cogn. 11*: 211–227.

Barsalou,, L. W. 1999. Perceptual symbol systems. *Behav. Brain Sci. 22*: 577–660.

Barsalou,, L. W. 2008. Grounded cognition. *Annu. Rev. Psychol. 59*: 617–645.

Bisiach, E., Perani, D.,Vallar, G., And Beerti, A. 1986. Unilateral neglect: Personal and extrapersonal. *Neuropsychologia 24*: 759–767.

Blakeslee, S. 2004. When the brain says, "don't get too close". *The New York Times,* July 13.

Brain, W. R. 1941. Visual disorientation with special reference to lesions of the right hemisphere. *Brain 64*: 224–272.

Buccino, G., Binkofski, F., and Riggio, L. 2004. The mirror neuron system and action recognition. *Brain Lang. 89*: 370–376.

Buxton, B. 2007. *Sketching User Experience. Sketching User Experiences: Getting The Design Right And The Right Design (Interactive Technologies).* Morgan Kaufmann.

Carlson, R. A., Avraamides, M. N., Cary, M., and Strasberg, S. 2007. What do the hands externalize in simple arithmetic? *J. Exp. Psychol. Learn. Mem. Cogn. 33*, 4: 747–756.

Clark, A. 1997. *Being There Putting Brain, Body, and World Together Again*. MIT Press, Cambridge, MA.

Clark, A. 2008. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension.* Oxford University Press.

Colby, C. L. 1998. Action oriented spatial reference frames in cortex. *Neuron 20*, 15–24.

Coslett, H. B. 1998. Evidence for a disturbance of the body schema in neglect. *Brain Cogn. 37*: 529–544.

Cross, E., Kraemer, D. J. M., Hamilton, A. F. De C., Kelley, W. M., and Grafton, S. T. 2009. Sensitivity of the action observation network to physical and observational learning. *Cereb. Cortex 19*: 315–326.

Eck, D., Gasser, M., and Port, R. 2000. Dynamics and embodiment in beat induction. In *Rhythm Perception and Production,* P. Desain and L. Windsor, Eds., Swets and Zeitlinger, Exton, PA.

Endsley, M. 1995. Toward a theory of situation awareness in dynamic systems. *Hum. Factors J. Hum. Factors Ergon. Soc. 37*, 1: 32–64.

Fodor, J. A. 1975. *The Language of Thought*. Harvard University Press, Cambridge, MA.

Frank, M. and Barner, D. 2012. Representing exact number visually using mental abacus. *J. Exp. Psychol. General. 141*, 1: 134–149.

Gibson, J. J. 1966. *The Senses Considered as Perceptual Systems*. Houghton Mifflin, Boston.

Gibson, J. J. and Crooks, L. E. 1938. A theoretical field-analysis of automobile-driving. *Amer. J. Psychol. 51*: 453–471.

Goldin-Meadow, S. 2005. *Hearing Gestures: How Our Hands Help Us to Think*. Harvard University Press.

Goldin-Meadow, S. and Beilock, S. L. 2010. Action's influence on thought: The case of gesture. *Persepct. Psychol. Sci. 5*: 664–674.

Goodwin, C. 1994. Professional vision. *Amer. Anthropol. 96*, 3: 606–633.

Graziano, M. S. A. and Gross, C. G. 1995. The representation of extrapersonal space: A possible role for bimodal, visuo–tactile neurons. In *The Cognitive Neurosciences*, M. S. Gazzaniga, Ed., MIT Press: 1021– 1034.

Harman, G. 1986. *Change in View: Principles of Reasoning*. MIT Press.

Hatano, G., Miyake, Y., and BinkS, M. G. 1977. Performance of expert abacus operators. *Cogn. 5*: 47–55.

Hinz, B. 2008. Practice exaggeration for large intervals and leaps. http://www.creativekeyboard.com/ oct08/hinz.html

Iriki, A. 2009. Using tools: The moment when mind, language, and humanity emerged. *Frontlines Riken Res. 4*, 5.

Iriki, A., Tanaka, M., and Iwamura, Y. 1996. Coding of modified body schema during tool use by macaque postcentral neurones. *Neurorep. 7*: 2325–2330.

Jeannerod, M. 1994. The representing brain: Neural correlates of motor intension and imagery. *Behav. Brain Sci. 17*: 187–245.

Jeannerod, M. 2001. Neural simulation of action: A unifying mechanism for motor cognition. *NeuroImage 14*: S103–S109.

Johnson-Laird, P. 1989. Mental models. In *The Foundations of Cognitive Science*, M. Posner, Ed., MIT Press, Cambridge, MA, Chapter 8.

Jones, K. 2003. What is an affordance? *Ecol. Psychol. 15*, 2: 107–114.

Kaschak, M. P., Zwaan, R. A., Aveyard, M., and Yaxley, R. H. 2006. Perception of auditory motion affects language processing. *Cogn. Sci. 30*: 733–744.

Kirsh, D. 1992. When is information explicitly represented? In *The Vancouver Studies in Cognitive Science*. Oxford University Press: 340–365.

Kirsh, D. 1995. Complementary strategies: Why we use our hands when we think. In *Proceedings of the 17th Annual Conference of the Cognitive Science Society,* J. D. Moore and J. F. Lehman, Eds.: 212–217.

Kirsh, D. 2005. Multi-Tasking and cost structure: Implications for design. In *Proceedings of the 27th Annual Conference of the Cognitive Science Society*. Lawrence Erlbaum, Mahwah, NJ.

Kirsh, D. 2009a. Problem solving and situated cognition. In *The Cambridge Handbook of Situated Cognition*, P. Robbins and M. Aydede, Eds., Cambridge University Press.

Kirsh, D. 2009b. Projection, problem space and anchoring. In *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, N. A. Taatgen and H. van Rijn, Eds., Cognitive Science Society, 2310–2315. KIRSH, D. 2009c. Interaction, external representations and sense making. In *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, N. A. Taatgen and H. van Rijn, Eds., Cognitive Science Society: 1103–1108.

Kirsh, D. 2009d. Knowledge, implicit versus explicit. In *Oxford Companion to Consciousness.* Oxford University Press, Cambridge, UK.

Kirsh, D. 2010. Explaining artifact evolution. In *The Cognitive Life of Things: Recasting the Boundaries of the Mind,* L. Malafouris, and C. Renfrew, Eds., McDonald Institute for Archaeological Research.

Kirsh, D. 2012a. How marking in dance constitutes thinking with the body. In *The External Mind: Perspectives on Mediation, Distribution and Situation in Cognition and Semiosis*, R. Fusaroli, T. Granelli, and C. Paolucci, Eds.: 112–113.

Kirsh, D. 2012b. Running it through the body. In *Proceedings of the 34th Annual Cognitive Science Society*. Lawrence Erlbaum.

Kirsh, D., Caballero R., and Cuykendall, S. 2012. When doing the wrong thing is right. In *Proceedings of the 34th Annual Cognitive Science Society*. Lawrence Erlbaum.

Kirsh, D. and maglio, P. 1995. On distinguishing epistemic from pragmatic actions. *Cogn. Sci. 18*: 513–549.

Kirsh, D. Muntanyola, D., Jao, J., Lew, A., and Sugihara, M. 2009. Choreographic methods for creating novel, high quality dance. In *Proceedings of the 5th Internation al Workshop on Design and Semantics and Form (DESFORM).* Kluwer.

Kossyln, H. and Moulton, S. T. 2009. Mental imagery and implicit memory. In *Handbook of Imagination and Mental Imagery*, K. D. Markman, W. M. P. Klein, and J. A. Suhr, Eds., Psychology Press, New York.

Knoblich, G. and Flach, R. 2001. Predicting the effects of actions: Interactions of perception and action. *Psychol. Sci. 12*, 467–472.

Là Davas, E. 2002. Functional and dynamic properties of visual peripersonal space. *Trends Cogn. Sci. 6*, 1.

Ladavas, E., Di Pellegrino, G., Farne, A., and Zeloni, G. 1998. Neuropsychological evidence of an integrated visuotactile representation of peripersonal space in humans. *J. Cogn. Neurosci. 10*: 581–589.

Makin, T. R., Holmes N. P., and Zohary, E. 2007. Is that near my hand? Multisensory representation of peripersonal space in human intraparietal sulcus. *J. Neurosci. 27*: 731–740.

Malafouris, L. 2008. At the potter's wheel: An argument for material agency. In *Material Agency,* C. Knappett, and L. Malafouris, Eds., Springer: 19–36.

Maravita, A. and Iriki, A. 2004. Tools for the body (schema). *Trends Cogn. Sci. 8,* 2: 79–86.

Maravita, A., Spence, C., Kennett, S., and Driver, J. 2002. Tool-Use changes multimodal spatial interactions between vision and touch in normal humans. *Cogn. 83*: B25–B34.

McLuhan, M. 1964. *Understanding Media: The Extensions of Man* 1st Ed. McGraw Hill, New York.

Myin, E. and O'Regan, J. K. 2008. Situated perception and sensation in vision and other modalities: Form an active to a sensorimotor account. In *Cambridge Handbook of Situated Cognition*, P. Robbins and A. Aydede, Eds., Cambridge University Press: 185–200.

Noë, A. 2005. *Action in Perception*. MIT Press.

O'Connor, E. 2006. Glassblowing tools: Extending the body towards practical knowledge and informing a social world. *Qual. Sociol. 29*, 2: 177–193.

O'Regan, J. K. and Noe, A. 2001. A sensorimotor account of vision and visual consciousness. *Behav. Brain Sci. 24*: 939–1031.

Potter, L. 1980. *The Art of Cello Playing: A Complete Textbook Method for Private or Class Instruction*. Alfred Music Publishing.

Previc, F. H. 1998. The neuropsychology of 3-D space. *Psychol. Bull. 124*: 123–163.

Prinz, W. 1997. Perception and action planning. *Euro. J. Cogn. Psychol. 9*: 129–154.

Proffitt, D. 2006. Embodied perception and the economy of action. *Perspect. Psychol. Sci. 1*: 110.

Rensink, R. A. 2002. Change detection. *Annu. Rev. Psychol. 53*: 245–277.

Rizzolatti, G. and Craighero, L. 2004. The mirror-neuron system. *Annu. Rev. Neurosci. 27*: 169–192.

Rizzolatti, G. and Sinigaglia, C. 2007. *Mirrors in the Brain: How Our Minds Share Actions and Emotions.*

Oxford University Press.

Scaife, M., and Rogers, Y. 1996. External cognition: How do graphical representations work? *Int. J. Hum.Comput. Stud. 45*: 185–213.

Sebanz, N. and Shiffrar, M. 2007. Bodily bonds: Effects of social context on ideomotor movements. In *Sensorimotor Foundations of Higher Cognition*, P. Haggard, Y. Rosetti, and M. Kawato, Eds., Oxford University Press.

Shin, Y. K., Proctor, R. W., and Capaldi, E. J. 2010. A review of contemporary ideomotor theory. *Psychol. Bull. 136*, 6: 943–974.

Simons, D. J. and Chabris, C. F. 1999. Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Percept. 28*, 9: 1059–1074.

Vaishnavi, S., Calhoun, J., and Chatterjee, A. 1999. Crossmodal and sensorimotor integration in tactile awareness. *Neurol. 53*: 1596–1598.

Varela, F., Thompson, E., and Rosch, E. 1991. *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.

Viviani, P. 2002. Motor competence in the perception of dynamic events: A tutorial. In *Common Mechanisms in Perception and Action*, W. Prinz and B. Hommel, Eds., Oxford University Press: 406–442.

Wexler, M., Kosslyn, S., and Berthoz, A. 1998. Motor processes in mental rotation. *Cogn. 68*: 77–94.

Williams, A. and Gribble, P. L. 2012. Observed effector-independent motor learning by observing. *J. Neurophysiol. 107*: 1564–1570.

Wilson, M. and Knoblich, G. 2005. The case for motor involvement in perceiving conspecifics. *Psychol. Bull.* 131: 460–473.

Winter, B. and Bergen, B. 2012. Language comprehenders represent object distance both visually and auditorily. *Lang. Cogn. 4*, 1: 1–16.