Cambridge University Press
978-1-107-09472-7 — Knowing Hands
David A. Rosenbaum
Excerpt
More Information

CHAPTER 1

Introducing Hands

"The Thinker" by the great French sculptor, Auguste Rodin, is a statue you've probably seen many times. It shows a naked man with a muscular build sitting on a rock. His right elbow is perched on his left knee, and his chin is resting on the back of his right hand. The man seems to have much on his mind. He has a muscular build, so he's been active in the past and will probably be active in the future. He seems to be taking a break between rounds of physical activity. During those workouts, it's possible he didn't get much thinking done. But sitting here now, curled in his contemplation, he seems finally able to ponder. Sitting motionless on this rock, he seems at last to be able to turn to serious mental work – hence his name, "The Thinker."

As a kid growing up in Philadelphia (one of the places where "The Thinker" can be seen in 3D), I remember being struck by the appearance of this naked man as he pondered who-knew-what? In the summer, he'd sit there in the blazing heat, bronzed not just on his skin but to his very core, for bronze was what he was made of. In the winter, he'd sit there too, with snow on his frame and ice on his haunches. No matter what the weather, he'd sit there that way, impervious to his surroundings, absorbed in his meditation.

Over the years, as I thought about "The Thinker," I came to realize that what's most striking about it is what it seems to say about thinking itself or at least what Auguste Rodin may have believed about cognition. Perhaps Rodin felt that thinking is best done in stillness. He may have even thought that the best, or *only*, way to think is to sit stock-still. Why else would he have portrayed the thinker as he did, not just sitting, but with the model's right elbow perched on his *left* leg – an odd way to sit if you think about it. Rodin may have wished to communicate that the best way to ponder is to put yourself in a position that leaves you unpoised for performance. I don't know whether Rodin held these beliefs. In any case, I came to realize that the statue provides a useful foil for the idea that cognition can be divorced from action.[1]

Now is it actually true that when you think, you don't act, and that when you act, you don't think? Is acting really such a "mindless" activity? From the way I've posed the question, you can tell I think the answer is no. Why do I believe this?

Consider what you do when you write, type, or gesture. You express your thoughts; that's why you do these things. But even when you use your hands for more mundane acts like placing a bowl on a table or grabbing a knife to slice bread, you also express thoughts, albeit ones of a more mundane kind. Carrying out physical acts, even for modest reasons, reflects knowledge that makes the acts possible.

The idea that even simple acts reflect thinking is the theme of this book. It's also the theme of talks I give about my research, which I've been privileged to give in invited lectures all over the world. In these talks, I emphasize the perspective I want to offer my lecture audience by drawing attention, after starting my presentation, to some latecomer who happens to make his or her way to an empty seat in the auditorium where I'm speaking. If it's a woman – I will assume so just so I can write "she" in the example to follow – I typically call out in as kind a way as I can and with a smile on my face to convey that I mean no harm, "Do you see that person coming in now? Look how she's walking in the direction of the empty seat she's spotted. Look how she's now beginning to angle her feet toward the row with that empty seat." The latecomer usually looks up at me slightly bewildered, but I smile and nod as welcomingly as I can. After I sense that she catches what's going on, hopefully benefiting from having read the abstract and title of my talk (a typical title is "Cognition and Action"), I continue with my sports-announcer commentary:

> Look at how skillfully she's walking sideways past the others in the row she's chosen to sit in. And now, as she's arriving at her seat, watch as she removes her backpack. She's grabbing the right strap of the backpack with her left hand. Now she's looking down behind her at the seat she's about to sit in. And now she's bending down and lowering herself into the seat while simultaneously lowering her backpack.

I continue this way until the latecomer has finally settled in, at which time I thank her for being a good sport, stress how impressive the person's performance was, and invite the audience to join me in applauding her for her fine performance, which they invariably do and for which the hapless star of the narration either nods emphatically (if she is as much as a ham as I am) or looks down sheepishly, embarrassed by the unwanted publicity.

After the lectures, I always check with the latecomer to make sure she (or he) had no hard feelings or misunderstandings about my intentions. Fortunately, none of these unwitting "stars" has ever said they have. In fact, more often than not, these people or others in the audience have come up to me and said things like, "You know, I never really noticed, until you mentioned it, how incredibly smart we are when we do ordinary things like take off our backpacks." To my great satisfaction, I sometimes hear the same remark weeks, months, or even years later. "Until I heard your talk, I never really saw how tightly bound cognition and action are."

This book is all about that binding. It's about the idea that everyday physical tasks, especially ones done with the hands, take tremendous intelligence.

I'm not the first person to have appreciated this. People in fields like artificial intelligence and robotics have done so, but I may be the only person in the field of cognitive psychology who has made this theme the consistent focus of my research. The topic captivated my interest at the end of my first year of graduate school in 1974 and has stuck with me ever since. The consistency of this focus, the amount I have learned, and the amount of research my colleagues and I have done have emboldened me to write this book, an entire volume on the cognitive psychology of manual control. It is the first such book, as far as I know.

Because cognitive psychology is about knowing, it's relevant to ask what kind of knowing will be talked about here. Four kinds of knowing will comprise the "quartet" for the discussion to come. The members of the quartet are suggested, ironically, by my consideration of remarks made, by, of all people, Donald Rumsfeld, the former secretary of defense for the United States under George W. Bush. Building on remarks by Rumsfeld, you can make a little table like the one below. It has things you know you know (item 1), things you know you don't know (item 2), things you don't know you know (item 3), and things you don't know you don't know (item 4).

	First-Order Knowing	
Second-Order Knowing	**Know**	**Don't Know**
Know	1	2
Don't Know	3	4

Cambridge University Press
978-1-107-09472-7 — Knowing Hands
David A. Rosenbaum
Excerpt
More Information

The table has as its columns what I call first-order knowing, and as its rows what I call second-order knowing. By first-order knowing, I mean facts or procedures that don't refer to knowing per se; they are simply facts or procedures that directly allow for behavior. By second-order knowing, I mean facts or procedures about knowing itself.[2]

All four of these categories have clear instances. Regarding item 1, there are things you know you know – for example, that you generally hold a jar with your left hand and unscrew the top with your right hand, at least if you're right-handed. Regarding item 2, there are things you know you don't know – for example, the names of all the bones and muscles in your hands (at least if you haven't studied anatomy). Regarding item 3, there are things you don't know you know – for example, how far you spread your thumb and fingers as you reach out to grasp objects of different sizes; this sort of knowledge is often called *implicit knowledge*. Finally, regarding item 4, there are things you don't know you don't know – for example, how to play a phrase on the violin to satisfy your new violin teacher. (I'll say more about this last example a bit later to clarify why I use it here.)

This book is mainly about two of these items: item 3, things you don't know you know; and item 4, things you don't know you don't know. It's less about items 1 and 2: things you know you know (item 1); and things you know you don't know (item 2). The reason there's not much on things you know you know (item 1) is that the material would be boring: You already know those things. And the reason there's not much on things you know you don't know is because, in the case of manual control, and also assuming you're more interested in cognitive psychology than anatomy, physiology, or biomechanics, the material would delve more into those topics than makes sense for the aims of this book.

The three topics I just mentioned in passing – anatomy, physiology, and biomechanics – are fascinating and important, but I will touch on them only briefly here. There are slews of books on hand anatomy, hand surgery, and related topics, and there is, similarly, a huge amount of information about the detailed features of the neural control of the hand. I will keep these topics at arm's length. More importantly and more positively, considering what this book is really about – the cognitive *psychology* of manual control – I will focus on behavioral analyses of manual performance. From such analyses, considerable progress has been made in understanding what goes on before, during, and after manual tasks are performed.

As for item 4 – things you don't know you don't know – that item is, perhaps, the most interesting when it comes to knowing hands; it's what much of this book is about. As I said earlier, an example of not knowing what you don't know is how to satisfy your new violin teacher. If you spent some time at a music school with very demanding violin teachers, like The Juilliard School in New York City, you would know that the teachers there work with their students on high-level aspects of musical interpretation and expression, focusing on what the students are trying to achieve musically, not just on how the students hold their fiddles or grasp their bows; those are things the students learned earlier in their musical careers. Behind the teachers' coaching is the idea that the students don't yet know what they don't yet know. The freshly arrived fiddlers don't yet realize that they don't know how to approach the phrasing of selected passages, for example. Like new graduate students who think they know how to do research and then discover that they actually don't know how to design a study to maximize its chance of having significant scientific impact, some (but certainly not all) incoming conservatory students arrive feeling pretty cocky, not realizing there's much they don't know they don't know.

The great concert pianist, Van Cliburn, made this point another way. When asked what it was like to practice for hours and hours every day before competing in the Tchaikovsky Piano Competition, which he won at age 23 in Moscow in 1963 in the midst of the Cold War, Van Cliburn said years later in an interview, "[T]he horizon is always receding in art." In saying this, Van Cliburn meant that his musical goals kept changing; they kept getting more and more refined. He understood that, on any given day, he didn't yet know what he didn't yet know about how to approach a piece he was practicing.[3]

What made Van Cliburn a great pianist wasn't just that he could speed his fingers at a mighty clip but rather that his musical *goals* were so lofty. He *aspired* for things that others did not.

Regardless of whether you're an aspiring violinist, pianist, or bricklayer, part of the adventure you're on is that you don't know how much there is to aspire to in the domain you're entering. Getting better at the skill entails discovering what the skill entails, not just being faster and more accurate, the traditional measures of skill.

These comments are psychological in the sense that they're about the goals and the refinement of goals, as well as the awareness of what the goals should be. Not knowing what you don't know applies to intellectual skills as well as perceptual-motor skills.

Consider again an incoming grad student, the new kind I referred to earlier. Here is a direct quote from one of my outgoing PhD students to one of my new incoming PhD students on the day he arrived at the lab: "I had no idea how much I didn't know when I started this program," she told him.

Once newcomers enter a bigger pond than the one they swam in before, they come to see that the vast amount they thought they knew was tiny relative to what they still had to discover.[4]

It's no accident that I'm drawing this parallel between graduate school and conservatory, for I want to make the point that it's difficult to draw the line between so-called intellectual skills and so-called perceptual-motor skills. The content of the skills may be different, but at root they're more similar than dissimilar. If you try to find differences between them, the more closely you look, the harder the differences are to find.[5]

A Robot Sets the Table

If you're not convinced that it takes a lot of knowledge to act skillfully in the physical world, consider the simple act of setting a table for a meal. Suppose you have friends coming over for dinner and you want everything to look nice for the supper. As it happens, you're fortunate to have a clever robot who can help you get ready. Because you're busy with your food preparation, you turn to the robot and say, "Robby, I'd like you to set the table for tonight's meal. We'll have eight guests. I'll be serving soup for the first course, then lobster for the main course. I really splurged! Then we'll have cake and ice cream for dessert. Morris and Mindy's kids are coming over. Let's put them at the end of the table. I'll be serving that nice Riesling I got the other day."

At this point, you realize the friend you were chatting with on the phone before you started talking to Robby has been waiting for you, "Sorry, Julie. Where was I?" you ask. "Can you believe she said that?" . . . pause . . . "No, me neither!"

Robby, being intelligent, hops to. What follows is a partial list of what "he" does. (A warning: The list is boring! It goes on and on and on. Read as much of it as you care to. You'll get the idea quickly: There are loads and loads of details that the robot must know to complete the task. By extension, there are tons of things you must know to do similar tasks yourself.)

First, the robot clears all the items from the dining room table. The items lying on the table include a tennis racquet, tennis balls, an empty tennis-ball can, a violin, a violin bow lying some distance from the violin, a cup of iced coffee filled nearly to the brim, a vase with some old drooping flowers in foul-smelling water, and a bunch of thumbtacks strewn on the table beyond the mouth of an open, tipped-over jar.

The robot takes the tennis racquet and tennis balls to the garage and puts them with the other sports equipment. Then he returns and carries the violin and bow to the room you use as your music studio. There he puts the bow and violin in your violin case, which he closes, making sure not to damage the instrument or bow. Anticipating the mischief Morris and Mindy's kids could get into if they saw a violin case at eye level, he sets the case on a high shelf.

Robby then returns to the dining room and lifts the cup of iced coffee from the table. Knowing that the coffee might spill if he carries it to the kitchen, he pours the coffee into the vase. He knows the base has adequate room to receive the coffee. He then carries the vase and now-empty-coffee cup to the kitchen. There he puts the cup in the sink, lifts the flowers out of the vase, lifts the lid of the garbage can and tosses the flowers into the garbage, and then releases the garbage-can lid and lets it fall shut. Next, he pours the foul-smelling water from the vase into the sink's garbage disposal, rinses out the vase, and leaves the vase beside the cup in the sink for later washing. Finally, he returns to the table and picks up the tacks one by one, being careful not to let them slide across the table's nice wooden surface before placing them in the jar. Once he has put all the tacks in the jar, he takes the jar to the kitchen, opens the cupboard where the jar of tacks normally stands, and puts the jar inside.

Now that Robby has completed the task of *clearing* the table, he turns his attention to *setting* the table. Still in the kitchen, he decides to go to the hall closet on his way back to the dining room. He opens the hall closet door and gets chairs from the closet to round out the number of chairs needed for everyone who'll be sitting around the table at the meal. The number of chairs needed is 11, a number the robot arrived at knowing there will be eight guests, plus you (the host) and your partner, as well as your mother, whom the robot overheard you say would be visiting and would be arriving at around 6, an hour before the usual start of dinners at your home. Because there were six chairs around the table, the robot knows he needs five more. He brings the five chairs in three trips, with two chairs in the first trip, two chairs in the second trip, and one chair in the third trip. No one told him how to partition the chair-carrying into separate trips. He figured that out for himself, taking into consideration how many chairs he could carry at a time.

Once Robby has put all the chairs around the table, he puts pillows on the two chairs at the end of the table, knowing that Morris and Mindy's kids are too short to reach the table without some support. Next he gets the tablecloth – the one for dinner parties – and lays it on the bare table, squaring it so it hangs to an approximately equal degree from all edges. After that, he gets the 11 dinner plates he'll need from the kitchen cupboard and sets them down, placing them in front of each chair, approximately equidistant from one another along their respective table edges but close enough to the table's edges that the plates' rims lie about an inch from the nearest edge.

Continuing with his work, Robby gets a fork, knife, soup spoon, and dessert spoon for each of the 11 guests. He places the utensils near their respective plates as per the rules of etiquette. Then he gets nine wine glasses, realizing that only nine will be needed because Morris and Mindy's kids won't be drinking alcohol. He puts the wine glasses beyond and to the right of the plates where the adults will be sitting, again per his knowledge of the rules of decorum for fine dining.

Nearing the end of the job, he gets nine water-drinking glasses for everyone, plus two Sesame Street cups for the kids. He places the glasses and the cups where they need to go beyond each dinner plate. Last of all, Robby brings a fresh vase to the table and sets it in the middle of the table atop a special wooden stand he puts there, knowing that this vase and the wood stand are the ones used for entertainment.

What the Robot Needed to Know

I could go on with this coverage of the robot's performance, but reading still more might tax your patience. The story is boring! No one wants to read about setting cups on tables or carrying wilted flowers to sinks. It's completely uninteresting! I dragged this out to illustrate how many things need to be attended to in order to complete everyday tasks.

Let me beg your indulgence a bit more, though, and direct your attention to some other things the robot needed to know in order to do what he did. By identifying those bits of knowledge, you can appreciate how vast the knowledge needs to be. It takes extra effort to be reminded of this because, for adults, adolescents, and healthy older children, the things I'm about to mention are things you might know you know, as well as things you might not know you know.

Begin with clearing the table. The robot knew that to set the table, he first had to clear it. He knew that the tennis racquet and balls – items that don't look anything alike – go together, and he knew that those items belonged in the garage with other sports equipment. He knew as well that it would be more efficient to take those items to the garage in one trip than to make repeated back-and-forth trips between the dining room and garage. Robby could have carried the tennis racquet to the garage and then he could have carried one tennis ball, then another tennis ball, and so on, but he chose a more efficient procedure.

Next consider the violin and bow. The robot knew that these two items go together, even though, like the tennis racquet and balls, they look nothing alike. He also knew that the violin and bow belonged in the music studio, that the violin and bow needed to be placed in the case, that the case had to be opened if it was not already open, that it had to be closed once the violin and bow were inside, and that the case should be placed where kids couldn't mess with it.

Concerning the iced coffee, Robby realized that it had to be removed from the table. Being aware of the vase near the coffee cup, he understood that he could avoid carrying the coffee all the way back to the kitchen, a trip better avoided considering that the coffee cup was nearly full and was likely to spill. The vase had plenty of room for the coffee, the robot realized, and he understood as well that he could carry the vase and cup to the kitchen at the same time, that he could toss the flowers into the garbage there, that he could dispose of the water in the vase, and that he could set both of these items into the sink.

Finally, Robby remembered that the table was not yet clear and that he still had to remove the tacks from the surface. He picked the tacks up one by one and placed them in the jar, doing so in a way that minimized the chance the table would be scratched. Had he merely swiped the tacks into the jar while holding the jar below the edge of the table (another possible means of removing them), he might have damaged the surface. Having gotten all the tacks into the jar, leaving the table unscathed, he took the jar to its regular storage site.

I could next review the things Robby knew to set the table, but you get the idea. I've let the discussion bog down into the minutiae I've covered, plus more that I'll spare you, to drive home the point that there are scores of details to be worked out even for the most mundane physical tasks. To get physical tasks done, an enormous amount must be known. By invoking a robot, I also meant to convey that this enormous fund of knowledge must be acquired. A robot "at birth" knows nothing, so it must somehow be taught, or must learn, what it needs to know. By contrast, for biological agents – humans or animals – there is always a chance that some of the knowledge is available innately, though much or most of the knowledge must be acquired through experience. No one knows from birth that sharp objects swiped across a wood surface can scratch the wood. "Roboticizing" the problem brings to the fore what a vast amount of knowledge is needed. It's not sufficient to say, as one might be tempted to say, "Oh, the robot just needs to be programmed to get these things done." What the program is, how it should be structured, and what its contents are, are the problems addressed here.

What About the Robot's Hands?

At the risk of trying your patience even more, allow me to continue with the robot example, turning to another aspect of the robot's performance that I deliberately suppressed. It's an aspect you may have noticed given that this book is about hands. You may have noticed that the robot relied on hand-related knowledge to guide his behavior. That the robot had hands was something I didn't say, but if you imagined the robot as humanoid – a reasonable assumption, perhaps, given the domestic setting in which I placed him (or it) – you might have imagined him as a manually endowed mechanical device. In addition, if you imagined the robot carrying out the tasks with his hands in ways that you would, then you drew on your own hand-related knowledge.

What hand-related knowledge did you draw on, assuming you did? To pursue this question, consider the first thing the robot did. He took the tennis racquet and tennis balls to the garage and put them there with other sports equipment. If you decompose this seemingly simple act into component acts, you realize the robot made a host of choices. The act of taking the tennis racquet and tennis balls to the garage entailed carrying those items in one trip. Making that one trip entailed gathering those items and carrying them to the garage. Once inside the garage, the racquet and balls had to be brought to the "sports section" and set down sensibly. Gathering the items from the table entailed determining how to grasp the items and how to hold them so they would survive the trip. Picking up the tennis racquet required a decision about which hand to use to pick up the racquet. Other related decisions concerned where along the handle to pick up the racquet (assuming it was desirable to pick up the racquet along its handle), how much force to apply to the handle, whether the racquet would be pushed across the table until part of its handle was suspended above the table's edge, and so on.

Questions like these could keep spilling forth. The seeming endlessness of the questions points to the fact that even mundane tasks require an enormous number of choices. The choices can't just be made one at a time. They're interconnected. For example, if one of the tennis balls were too far from the side of the table where the robot was standing, the robot could either make his way to the other side of the table to get that one ball or instead coax the ball toward him with the racket. Depending on which of these strategies he pursued, the racquet would have to be picked up differently. If the racquet were lifted just to carry it to the garage, it might be picked up differently than if it were lifted to use it first to coax a ball from