

Thinking by Doing and Doing by Thinking: A Taxonomy of Actions

Hansjörg Neth (nethh@rpi.edu)

Cognitive Science Department, Rensselaer Polytechnic Institute
Carnegie 108, 110 8th Street, Troy, NY 12180, USA

Thomas Müller (Thomas.Mueller@phil.uu.nl)

Department of Philosophy, Universiteit Utrecht, Heidelberglaan 6
3584 CS Utrecht, The Netherlands

Abstract

Taking a lead from existing typologies of actions in the philosophical and cognitive science literatures, we present a novel taxonomy of actions. To promote a notion of *epistemic agency* we distinguish theoretical (mental state-directed) from practical (world-directed) actions. Our basic structural unit is that of a *teleological frame*, which spans one specific goal of an agent. Relative to a given teleological frame, actions can be classified as focal (directed towards the end) or ancillary (directed towards a means). The framework is applied to further illuminate previous attempts to distinguish between pragmatic and epistemic actions (Kirsh & Maglio, 1994). Physical actions that substitute or support mental processes are re-classified as practical ancillary actions that are strategically contingent alternatives to theoretical actions.

Keywords: Action theory; basic actions; practical vs. theoretical actions; epistemic agency

Introduction

Cognitive science researchers routinely use the terms ‘action,’ ‘act,’ ‘agent’ and ‘agency’ but rarely reflect upon this usage. Technical discourse often imports concepts from everyday language, which then develop a terminological double-life: General constructs (e. g., visual perception, memory, attention, etc.) and more specialized variants (visual short-term memory, mental model, top-down control) coexist both informally, with intended meanings identical or close to their natural-language origins, and also in more narrow, technical and operationalized meanings, typically when becoming topical in some research effort. Surprisingly, there does not yet seem to exist any precise and universally accepted definition of the notions of action and agency.

Actions are the vehicle of rationality that link theoretical ideas and insights to real-world consequences. Without the ability to act on the basis of goals and beliefs, no mathematical proof or empirical finding would ever matter to anything. Given the centrality of the concept of agency for explanations of behavior and a rich tradition of manifold discussions about the nature of human rationality, it is startling that there is little consensus about the types of actions that humans principally engage in. We believe that it is time to clarify certain ambiguities and replace some common-sense notions by a more comprehensive taxonomy of actions. Although these efforts are mainly conceptual in nature, we trust that our framework will be of value to the wider cognitive science community.

An attempt to illuminate the nature of human agency is particularly pressing in the context of recent advances in philosophy that aim to “radically reconfigure our image of ratio-

nality” (Clark, 2001, p. 121). The ubiquitous use of tools and gestures feeds philosophical arguments suggesting that many elements of the external world (e. g., body parts, interactive devices and informational artifacts) ought to be viewed as integral parts of cognition, rather than mere media for modified inputs and outputs (Clark & Chalmers, 1998). Clark (2003) even claims that we are living cyborgs, routinely wearing and relying on cognitive prostheses.

Such views are no surprise to empiricists. In fact, most real-world problem solving recruits external tools and achieves its goals through an intricate process of interaction with the physical environment. When solving arithmetic problems, people spontaneously distribute memory demands over internal and external resources (e. g., Cary & Carlson, 2001) and spontaneously employ their hands and other available resources to rearrange, add and count items (Carlson, Avraamides, Cary, & Strasberg, 2007; Neth & Payne, 2001). Continued reliance of experienced pilots on external markers to track current control states (Hutchins, 1995) shows that external aids are not just cognitive crutches for novices. On a lower level, research on so-called ‘active vision’ (e. g., Findlay & Gilchrist, 2003) supports the view that agents continually sample their environments, rather than constructing complex internal representations.

To account for these phenomena, cognitive science has seen a recent upsurge in approaches that try to cross the traditional divide between thought and action by mapping the close connections between mental processes and the environments in which they are situated (e. g., Lave, 1988; Suchman, 1987; Hollan, Hutchins, & Kirsh, 2000; Neth et al., 2007). Despite differences in emphasis and labels, their common denominator is that the embodied and embedded nature of cognition fundamentally informs our notion of human agency.

Cognition is *adaptive* in two complementary but distinct ways: On one hand, the cognitive system adapts itself to the structure of its environment to transcend its inherent limitations (e. g. of attention and memory). On the other hand, cognitive systems exhibit a pervasive tendency to adapt and structure their environments in service of their goals.

Although the world also affects the agent (mediated by various sensors), we would not regard this as a form of agency, as it would seem esoteric to regard the world as an agent pursuing an agenda. Yet there is an interesting type of action that occurs whenever agents act to change their own mental state rather than the world. We will call this a *theoretical action*

and argue that this concept is not an oxymoron, but an indispensable piece in the puzzle of predicting and explaining human behavior.

Our distinction between practical and theoretical actions bears close resemblance to the distinction between pragmatic and epistemic actions that was introduced in the context of the popular videogame Tetris (Kirsh & Maglio, 1994). To account for the fact that players often prefer to rotate a falling Tetris piece (or ‘zoid’) manually rather than mentally the authors define *epistemic actions* as physical actions that improve cognition by facilitating or reducing the need for internal computations.

While we emphatically embrace the pioneering work by Kirsh and colleagues, we believe that their chosen terminology is unfortunate. First, the series of papers that pursued and elaborated the basic distinctions shows some terminological drift, with non-pragmatic actions being labeled, in turn, ‘perceptive actions’ (Kirsh & Maglio, 1992), ‘epistemic actions’ (Kirsh & Maglio, 1994; Maglio & Kirsh, 1996), ‘complementary strategies’ (Kirsh, 1995a), or simply ‘interactive skill’ (Maglio, Matlock, Raphaely, Chernicky, & Kirsh, 1999). Given this variety of terms, it is tempting to wonder whether several possible distinctions are not being blurred.

Using the term ‘pragmatic’ to denote actions that bring an agent physically closer to its goal invites further misunderstandings. Increasing the efficiency of some process through epistemic actions can seem rather “pragmatic”. Both senses of the term used here are still different from the meaning of *pragmatics* in other disciplines, e. g., linguistics—not to mention the philosophical tradition of *pragmatism* that has influenced psychology as well (Schön, 1983).

Additional problems stem from a clash of different conceptual traditions. Kirsh and Maglio (1994) are careful to emphasize that their epistemic actions are *external* and *physical* actions. Yet the term ‘epistemic’ evokes almost the opposite connotation in philosophical or AI contexts (e. g., Baltag & Moss, 2004) that exclusively wonder about belief revisions and internal updates of an agent’s informational state.

Terminological heritage can be the source of apparent paradoxes. Pragmatic actions traverse the problem space towards a goal state. But if a problem consists in mentally computing the solution to an arithmetic problem, all steps towards reaching a solution would seem rather “epistemic.” Similarly, while epistemic actions aim to reduce cognitive complexity, their intended *object* remains unclear. Any real-world ‘problem’ emerges only through the interaction of an agent’s mental state (perceptions, beliefs, goals) and the external environment. To change the computational constraints of the task, an epistemic agent acts upon its environment (e. g., rotates a Tetris piece) and alters the environment in accordance with its goal (to recognize a piece’s shape). But why would *this* not also count as an instance of successfully traversing the physical problem space, i. e., a “pragmatic” action?

The situation that pragmatic actions can sometimes seem epistemic and epistemic actions are somewhat pragmatic cre-

ates an epistemically confusing state of affairs. We believe that the issue is conceptual in nature and rooted in the lack of a consistent theory of agency in contemporary cognitive science. The present paper attempts to remedy this situation by making and defending the following claims:

1. Actions can only be analyzed in reference to a specific goal and agent.
2. Agents can act to change their mental states. Such theoretical actions are ubiquitous and non-mysterious.
3. Explicating an action’s goal imposes a *teleological frame* upon the act of analysis and induces a hierarchical distinction between *focal* and *ancillary* actions.
4. By clarifying conceptual issues concerning the embodied and embedded nature of cognition our taxonomy of actions provides a novel and valuable contribution to recent debates in cognitive science.

Motivation

Imagine two students taking part in a cognitive science experiment that involves playing the videogame Tetris.¹ By participating in the experiment, the students earn course credit. Although the subjects’ behavior could potentially be characterized in purely physical or physiological terms (e. g., by recording muscular activity, direction of gaze, and galvanic skin responses), the Method section of an experimental report typically describes the task in terms of the subjects’ goals and actions. Thus, any description of ‘what the subjects are doing’ is already cast in the jargon of *agency*.

(A) A cooperative subject. Subject A is sitting in front of a computer screen, dealing with the falling pieces one at a time by identifying the pieces (sometimes pressing buttons to rotate them into a standard orientation) and monitoring the available positions in the growing wall at the bottom of the screen, and by moving the pieces to an appropriate place speedily by pressing buttons quickly, such as to achieve a high score. At the end of the experiment, A obtains credit for her participation, which makes up part of the requirements for a course she is taking to graduate with a college degree.

(B) An uncooperative subject. Subject B is also sitting in front of a computer screen, sometimes pressing buttons such as to move pieces in some vaguely reasonable way. He spends some of the time calling a friend to make plans for the evening, writing down the number of a cinema mentioned by his friend. Later he calls that number in order to reserve tickets. At the end of the experiment, B, like A, gets credit for his participation, which he, too, ultimately needs for his degree.

Analysis

By relating our examples to the philosophical literature on agency we can identify major challenges for a taxonomy of actions.

¹Against a recent trend in philosophical action theory, we steer clear of involved examples of, e.g., Martians zapping people with “M-rays” (Mele, 2000).

What is being done? Throughout the lab session, both subjects are participating in an experiment, playing Tetris, earning course credit, and leading their student life. Exactly half way through the session, A may be pressing the “→” key to rotate the current piece clockwise by 90° in order to identify it, while B may be scribbling down the digit “7” while paying scant attention to the current piece. The more global descriptions just mentioned remain true also at that point in time. But what are the students doing, *really*? While some lines of research point to basic units of 3 s duration (Land, Mennie, & Rusted, 1999), an activity-theoretic approach (Engeström, 2000) emphasizes larger units of analysis, viz., activities constituted by actions. The question “what is being done”, aimed at the action level, may be a red herring.

How can actions be classified? Some classifications of actions are intrinsic, whereas others are extrinsic. Consider B’s action of scribbling down a digit. By identifying this action as that of *scribbling down the digit “7,”* the action’s criterion of success is given intrinsically: The action is successful if and only if B has afterwards recorded a “7.” If he fails (e.g., because his pencil breaks), he will only have tried. Success or failure are contingent (even though success is implied by descriptions *ex post*), but the *criterion* of success is *internally related* to the action (Wright, 1963, p. 116): the action and its criterion of success form a conceptual unity.

On the other hand, many classifications of actions are *extrinsic*, relying on the identification of an ascribed goal. E.g., whether B’s action is seen as a means (e.g., for planning his evening) or as an end (setting the goal to be the recording of a digit) depends on which goal is invoked. The necessity or contingency of an action as a means is similarly goal-relative, as are classifications of an action as adaptive or non-adaptive, or as rational or irrational. Thus, it makes little sense to ask whether A or B act more rationally overall. With respect to the goal of playing Tetris, A handles the available resources better, while B diverts his attention by making a phone call. With respect to leading a life, however, B might appear to be more rational in not expending unnecessary energy on the Tetris task and using the time to cultivate his social ties.

How to identify “the right” level of description? What the subjects are doing can thus be described variously, depending on the goals attributed to them, and agents are not limited to having exactly one goal at any given time (Thompson, 2008). Agents’ goals are ordered hierarchically. The goals invoked in describing what A is doing can be *linearly* ordered: She is pressing a button in order to rotate the current piece, in order to complete a row speedily, in order to get a high score, in order to participate in the experiment, in order to obtain course credit, and so forth. For B, the corresponding hierarchical structure is only a *partial* ordering: E. g., he is not writing down the digit in order to participate in the experiment, but rather in order to plan his evening.

The fact that actions must be viewed with respect to a specific goal—and can thus be classified variously—motivates a

quest for a preferred, natural, or even metaphysically fundamental level of “basic actions” (Danto, 1965). Unfortunately, any attempt to define actions in terms of bodily movements, which might single out the movements of the limbs as the elements of “the real” action hierarchy, falls flat.

Action as bodily movement? Undoubtedly there exists bodily movement without action (e.g., in reflexes or instinctive behavior). Reversely, action does not necessarily involve bodily motion, as witnessed by actions of refraining and by our proposed type of theoretical actions.

Refraining from doing something can be a real action without bodily motion. This is most obvious in cases in which the act of refraining requires active resistance from temptation, or in which it has moral consequences, as in the debate about active vs. passive euthanasia, or killing vs. letting die (Bennett, 1989). Structurally, however, refrainings are actions no matter what is at stake (Belnap, Perloff, & Xu, 2001, Chap. 9).

A different type of actions not necessarily involving bodily motion are actions whose criterion of success is a change of an agent’s mental state. Such actions are routinely investigated in cognitive science research, e.g., when subjects are asked to complete theoretical tasks such as counting to produce cognitive load. Open issues remain, however,² as certainly not all cognitive processes are under active control. E.g., forming a belief appears not to be an action,³ and the same seems to be true generally for perception.⁴

Acting as homeostasis? The primary purpose of action may be seen in maintaining a certain homeostatic equilibrium. On such a view, whenever an agent experiences some relevant internal equilibrium as disturbed, she acts in order to reduce and overcome the discrepancy until equilibrium is reestablished. This view captures some important aspects of agency: First, an agent is typically motivated by a perceived discrepancy between the actual and a desired state of affairs (action as *problem-based*). Secondly, agency is always directed towards achieving some goal (*intentional structure*). And thirdly, both the motivational trigger for starting an action and the signal for its termination depend upon the agent’s subjective cognitive state (e.g., on his beliefs) rather than objective facts (*epistemological uncertainty*). A specific instance of the homeostasis view is reflected in the belief-desire-intention (BDI) model of agency popular in computer science and AI. Here, the emphasis on belief reflects the epistemological uncertainty of agency, the intentional structure is directly represented, and the problem-based aspect of agency is catered for by desires playing the role of perceived deviations from an equilibrium situation.⁵

²Cf. the critical discussion of the distinction between what he calls “overt” vs. “mental” actions in Mele (1997).

³Cf. the large discussion following Williams’s (anti-voluntarist) article “Deciding to believe” in Williams (1973).

⁴A useful formal criterion for agentiveness is provided by the “stit paraphrase thesis” (Belnap et al., 2001, p. 7).

⁵Cf. Georgeff, Pell, Pollack, Tambe, and Wooldridge (1999), for a critical discussion.

While this homeostatic view is useful on some general level, an adequate theory of agency must take a number of additional aspects into account. On the motivational layer, intentions, plans and policies can be distinguished (Bratman, 1999). Furthermore, the contingency of action needs to be taken into account, both metaphysically (as something the agent need not have done),⁶ with respect to an action’s success, and with respect to alternative courses of action making room for strategic rationality in the choice of means.

Which goals? Which goals are available for the description of an agent’s actions? Must the agent actually have those goals, or even be conscious of having them? This seems wrong. A change in point of view seems always possible, and even the agent herself may benefit from conceptualizing her current and future actions with respect to a hypothetical goal—e.g., in considering whether to endorse that goal.

Thus we submit that the choice of a level of description (the choice of a goal relative to which an agent’s actions are described) should in the first instance be viewed as a tool of analysis (available to the agent herself too), and not a matter of “getting things right.” That said, it can still be illuminating to describe an agent’s actions with respect to different goals.

A Taxonomy of Actions

Our taxonomical framework, employing three dimensions of categorization, leads to a taxonomical table with six entries.

Practical vs. theoretical actions. As mentioned, an action is internally related to its criterion of success specifying some change of the given situation.⁷ That change can pertain either to the agent’s internal state or to the environment. As a first, intrinsic taxonomical dimension, actions pertaining to a change in the agent’s internal state we call *theoretical actions*, while actions pertaining to the agent’s environment we call *practical actions*.

Focal vs. ancillary actions. Relative to a presupposed *teleological frame* $F = \langle \alpha, \gamma \rangle$ specifying the agent α and a goal γ , we will call the action whose criterion of success is to realize γ the *focal* action relative to F , while any subordinate action will be called *ancillary* actions. All ancillary actions are temporally contained within the given focal action, but the reverse is not necessarily true: Some actions may be performed within the temporal scope of a focal action, but may lie outside the hierarchy of actions introduced via F .

Necessary vs. contingent ancillary actions. Apart from the metaphysical and success-related contingency of actions, for ancillary actions (thus, presupposing a teleological frame) there is a classificatory dimension of necessity vs. contingency with respect to the availability of alternative means: Sometimes the means for a given goal are fixed, but usually

there is room for strategic variation. There is, as they say, more than one way to skin a cat. Non-necessary ancillary actions for which an alternative is available we call *contingent*.

Overview. Summing up, our taxonomical scheme assumes a given teleological frame F and classifies actions into one of six categories:

	focal	ancillary	
		necessary	contingent
theoretical	1	3	5
practical	2	4	6

All these categories are present in our introductory examples. First, take $F_B = \langle B, P \rangle$, where P is B ’s goal of planning the evening, which specifies a theoretical focal action (1). Relative to F_B , his getting up and having breakfast may be a necessary practical ancillary action (4), while recalling the day of the week may be a necessary theoretical ancillary action (3). Secondly, take $F_A = \langle A, T \rangle$, where T is A ’s goal of getting a high score in Tetris, which specifies a practical focal action (2). Relative to F_A , her mentally rotating a falling Tetris piece may be a contingent theoretical ancillary action (5), while pressing the \rightarrow key to physically rotate the piece may be a contingent practical ancillary action (6).

Applications

This paper focuses on the development of conceptual structure and thus contains neither empirical data nor computer models. In our view, a lucid understanding of the conceptual basis of agency is a precondition for deeper theoretical understanding. Collecting data and developing formal models to substantiate and diversify the conceptual account is a natural next step in our research agenda, but beyond the scope of this paper. Here, we will show the usefulness of our approach by shedding light on the distinction between pragmatic and epistemic actions that has been discussed in the Introduction.

Re-conceptualizing epistemic actions. The simple duality of pragmatic vs. epistemic actions is problematic and confusing because it lacks (or leaves implicit) information on each of our three levels of analysis:

- Kirsh and colleagues do not explicitly develop a hierarchical account of actions like we do in our notion of teleological frames. The potential paradoxes mentioned above (of pragmatic actions sometimes seeming epistemic, and epistemic actions appearing to be pragmatic) result from oscillating between focal and ancillary actions, i. e., viewing epistemic actions simultaneously as means to global ends (playing Tetris), means to mental subgoals (identifying a piece’s shape), and ends in themselves on an even smaller scale (rotating a piece).
- The goal of epistemic actions is to change some aspect of the world, i. e., they are *practical* actions. However, they do so in the service of a *theoretical* action (identifying a piece’s shape).

⁶At least on a libertarian account of agency. Cf. Kane (2002) for an overview of the philosophical discussion on the free will problem.

⁷More appropriately, the criterion specifies a *transition* (Belnap et al., 2001, p. 192). We will not press this technical point here.

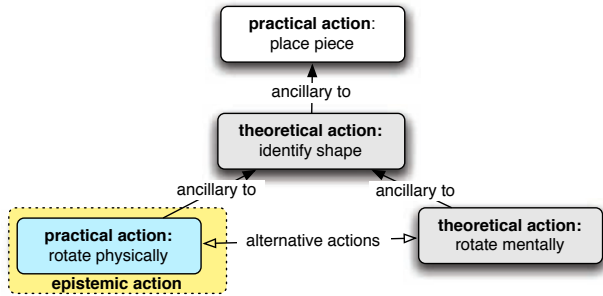


Figure 1: The complex structure of an epistemic action.

- An essential aspect of epistemic actions is that they are *non-necessary*. A practical action is only called ‘epistemic’ if its use is strategic, under the discretion of the agent, and there exists an *alternative theoretical* action (e. g., mental rotation).

The conceptual distinctions provided by our framework capture epistemic actions as *contingent ancillary practical actions* in the service of *focal theoretical actions* (i. e., actions of type 6 in the service of type 1 actions). More specifically, postulating an epistemic action presupposes a teleological frame *F* and thus, an agent with a goal. Relative to *F*, an action in a given concrete situation is called *epistemic* if and only if it fulfills the following three conditions (cf. Figure 1): (1) The action is *ancillary*, i. e., its criterion of success is not the fulfillment of goal of *F*, but of some subgoal. (2) The action is *practical*, i. e., it involves actively manipulating the agent’s environment, involving bodily motion by the agent. (3) There is an alternative course of action available in which the action would be replaced by a *theoretical* action.

It needs to be emphasized that, regardless of similarities in semantic connotations, such epistemic actions are fundamentally different from our category of theoretical actions. Epistemic actions are subordinate practical actions performed in the service of theoretical actions. Of course, any theoretical action that triggered this optional ancillary practical action may itself be ancillary with respect to some higher-level focal action, e. g., that of achieving a high score in Tetris, obtaining course credit, or leading a good life.

Our framework also reveals that Kirsh and Maglio (1994)’s pragmatic actions can be viewed as alternative practical actions within the same focal frame as the epistemic action (e. g., press a key to drop a piece to the bottom, rather than rotating it to identify its shape) or as actions on a super-ordinate level that would define a new focal frame (complete a row in the Tetris board). This ability to shift focal frames within the class of pragmatic actions illustrates that the hierarchical level must be defined as an independent dimension.

In conclusion, we agree with Kirsh and colleagues that epistemic actions are intriguing phenomena that merit close theoretical and empirical scrutiny. However, as apparent antonyms to pragmatic actions they easily give rise to conceptual confusion. We therefore propose to supplement the duality of pragmatic vs. epistemic actions by a more detailed

account of epistemic agency that incorporates the distinctions between focal vs. ancillary actions, practical vs. theoretical actions, and necessary vs. contingent actions. It would be regrettable if the notion of epistemic actions—by conflicting with traditional terminology and criss-crossing various category boundaries of our taxonomy—actually hindered further exploration and understanding of theoretical actions.

Conclusion

With respect to our introductory claims we conclude:

1. Any assessment of the rationality, behavioral adaptiveness or strategic optimization (e. g., performing a task as quickly and accurately as possible) presupposes a taxonomy of actions. Insight into the relativization of action descriptions with respect to teleological frames allows for a resolution of a number of puzzles concerning rationality, e. g., how ‘one and the same thing,’ (like making a phone call during a game of Tetris) can be both rational and irrational.
2. *Theoretical actions* are ubiquitous and non-mysterious actions through which an agent aims to change its mental state. The notion of a theoretical action is compatible with the extended mind hypothesis and an embodied and embedded view of cognition. In cases in which entities outside a person’s skin and skull make up part of her cognitive set-up, manipulation of these entities properly counts as theoretical, showing that our taxonomy is not tied to a naïvely biologicistic understanding of agency.
3. Actions can only be analyzed in reference to a *teleological frame* specifying a specific agent and goal. Such a frame induces a distinction between *focal* and *ancillary* actions, which is always relative, but non-arbitrary.
4. Our taxonomy can be fruitfully applied to recent issues in cognitive science. We have demonstrated this by re-analyzing the distinction between pragmatic vs. epistemic actions (e. g., Kirsh & Maglio, 1994) and giving a detailed account of the conceptual challenges that face any such classificatory scheme. According to our taxonomy, epistemic actions are contingent ancillary practical actions in the service of focal theoretical actions.

Future directions. Having provided an initial indication of the usefulness of our framework, here is a list of further issues that will be addressed in future research:

- Would a taxonomy of actions benefit from additional *functional* distinctions? Interesting candidates for further distinctions include actions to create and maintain spatial arrangements that simplify choice, perception or internal computation (Kirsh, 1995b); or actions that are used specifically in the service of planning or learning (Neth & Payne, 2002).
- Are all types of actions equal, or are there certain privileged types of actions? Current research in the area of immediate interactive behavior discusses conflicting hypotheses about the currency that governs behavioral trade-

offs. For instance, whereas a minimal-memory view suggests that memory resources are used more sparingly than perceptual-motor processes (Ballard, Hayhoe, Pook, & Rao, 1997), the soft-constraints hypothesis (Gray, Sims, Fu, & Schoelles, 2006) argues that adaptive organisms aim to optimize the overall time-on-task, rather than any particular resource.

- Similarly, is there a privileged level of action, e. g., at the so-called *unit-task* level (Card, Moran, & Newell, 1983)?
- How does a theoretical taxonomy of actions relate to practical approaches of cognitive tasks analysis? An applied goal of our conceptual framework consists in informing current approaches of cognitive task analysis. For instance, it would be conceivable to inform the choice of operator types when reconstructing tasks within a CPM-GOMS framework (Gray, John, & Atwood, 1993).

In addition to conceptual investigations, striving for a unifying theory of human agency will require the combination of empirical studies and computer models. Given the length of our list and the current status of the field—and perhaps fortunately for the future prospects of both philosophers and empirical scientists—there appears to be no shortage of theoretical or practical problems anytime soon.

Acknowledgments

The authors acknowledge support by the Deutsche Forschungsgemeinschaft (DFG).

References

- Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. N. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, 20(4), 723–767.
- Baltag, A., & Moss, L. (2004). Logics for epistemic programs. *Synthese*, 139(2), 165–224.
- Belnap, N., Perloff, M., & Xu, M. (2001). *Facing the Future: Agents and Choices in Our Indeterminist World*. Oxford: Oxford University Press.
- Bennett, J. (1989). Two departures from consequentialism. *Ethics*, 100(1), 54–66.
- Bratman, M. E. (1999). *Intentions, plans and practical reason*. Stanford CA: CSLI Publications.
- Card, S., Moran, T., & Newell, A. (1983). *The psychology of human-computer interaction*. Hillsdale, N.J.: Lawrence Erlbaum.
- Carlson, R. A., Avraamides, M. N., Cary, M., & Strasberg, S. (2007). What do the hands externalize in simple arithmetic? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(4), 747–756.
- Cary, M., & Carlson, R. A. (2001). Distributing working memory resources during problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27(3), 836–848.
- Clark, A. (2001). Reasons, robots and the extended mind. *Mind and Language*, 16(2), 121–145.
- Clark, A. (2003). *Natural-born cyborgs. Minds, technologies, and the future of human intelligence*. Oxford, UK: Oxford University Press.
- Clark, A., & Chalmers, D. (1998). The Extended Mind. *Analysis*, 58(1), 7–19.
- Danto, A. (1965). Basic Actions. *American Philosophical Quarterly*, 2, 141–148.
- Engeström, Y. (2000). Activity theory as a framework for analyzing and redesigning work. *Ergonomics*, 43(7), 960–974.
- Findlay, J. M., & Gilchrist, I. D. (2003). *Active Vision: The Psychology of Looking and Seeing*. Oxford University Press.
- Georgeff, M., Pell, B., Pollack, M., Tambe, M., & Wooldridge, M. (1999). The belief-desire-intention model of agency. *Proceedings of the 5th International Workshop on Intelligent Agents V: Agent Theories, Architectures, and Languages (ATAL-98)*, 1555, 1–10.
- Gray, W. D., John, B. E., & Atwood, M. E. (1993). Project Ernestine: Validating a GOMS Analysis for Predicting and Explaining Real-World Task Performance. *Human-Computer Interaction*, 8(3), 237–309.
- Gray, W. D., Sims, C. R., Fu, W.-T., & Schoelles, M. J. (2006). The soft constraints hypothesis: A rational analysis approach to resource allocation for interactive behavior. *Psychological Review*, 113(3), 461–482.
- Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7(2), 174–196.
- Hutchins, E. (1995). How a cockpit remembers its speeds. *Cognitive Science*, 19(3), 265–288.
- Kane, R. (Ed.). (2002). *The Oxford handbook of free will*. Oxford: Oxford University Press.
- Kirsh, D. (1995a). Complementary Strategies: Why we use our hands when we think. In *Proceedings of the 17th Annual Conference of the Cognitive Science Society* (pp. 212–217). Hillsdale, NJ: Lawrence Erlbaum.
- Kirsh, D. (1995b). The intelligent use of space. *Artificial Intelligence*, 73(1), 31–68.
- Kirsh, D., & Maglio, P. P. (1992). *Perceptive actions in Tetris*. Palo Alto, CA.
- Kirsh, D., & Maglio, P. P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18(4), 513–549.
- Land, M., Mennie, N., & Rusted, J. (1999). The roles of vision and eye movements in the control of activities of daily living. *Perception*, 28(11), 1311–1328.
- Lave, J. (1988). *Cognition in practice. mind, mathematics and culture in everyday life*. Cambridge, UK: Cambridge University Press.
- Maglio, P. P., & Kirsh, D. (1996). Epistemic action increases with skill. In C. Cottrell (Ed.), *Proceedings of the 18th Annual Conference of the Cognitive Science Society* (pp. 391–396). Mahwah, NJ: Lawrence Erlbaum.
- Maglio, P. P., Matlock, T., Raphaely, D., Chernicky, B., & Kirsh, D. (1999). Interactive Skill in Scrabble. In S. Hahn (Ed.), *Proceedings of the 21st Annual Conference of the Cognitive Science Society* (pp. 326–330). Mahwah, NJ: Lawrence Erlbaum.
- Mele, A. (1997). Agency and Mental Action. *Noûs*, 31, 231–249.
- Mele, A. (2000). Goal-directed action: Teleological explanations, causal theories, and deviance. *Philosophical Perspectives*, 14, 279–300.
- Neth, H., Carlson, R. A., Gray, W. D., Kirlik, A., Kirsh, D., & Payne, S. J. (2007). Immediate Interactive Behavior: How embodied and embedded cognition uses and changes the world to achieve its goals. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Conference of the Cognitive Science Society* (pp. 33–34). Austin, TX: Cognitive Science Society.
- Neth, H., & Payne, S. J. (2001). Addition as Interactive Problem Solving. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 698–703). Mahwah, NJ: Lawrence Erlbaum.
- Neth, H., & Payne, S. J. (2002). Thinking by Doing? Epistemic Actions in the Tower of Hanoi. In W. Gray & C. Schunn (Eds.), *Proceedings of the 24th Annual Conference of the Cognitive Science Society* (pp. 691–696). Mahwah, NJ: Lawrence Erlbaum.
- Schön, D. (1983). *The reflective practitioner: How professionals think in action*. New York: Basic Books.
- Suchman, L. A. (1987). *Plans and situated actions: The problem of human-machine communication*. Cambridge, UK: Cambridge University Press.
- Thompson, M. (2008). *Life and action*. Cambridge MA: Harvard University Press.
- Williams, B. (1973). *Problems of the self*. Cambridge: Cambridge University Press.
- Wright, G. von. (1963). *The varieties of goodness*. London: Routledge.