17

Ways of Action Science

Wolfgang Prinz

Abstract

This chapter argues that an action-oriented view of cognition is nice to have, but not enough. The study of cognition most certainly needs to be extended to include action. However, the study of action requires more than simply understanding its cognitive foundations. This chapter discusses two functional features of action that a cognitive approach fails to capture: *top-down control* and *action alignment*. Top-down control operates within individuals and requires a framework that addresses the formation of motives, goals, and intentions as precursors of action selection and execution. Action alignment operates between individuals, necessitating a framework that addresses the common representational basis of perception and production. To accommodate these features we need to proceed from including action in cognitive science to including cognition in action science.

Introduction

The aim of this Forum was "to examine the key concepts of an emerging action-oriented view of cognition and the consequences of such a paradigm shift." Although I certainly share this aim, I address a more ambitious aim in this article; namely, to examine key concepts of an emerging action science and the consequences of such a paradigm shift for cognition. My aim is to find a place for cognition in action science, not just for action in cognitive science.

As has been elegantly pointed out in Neisser's classical foundation of modern cognitive psychology, the task of research in cognition is to trace the fate of the input and study stimulus information and its vicissitudes in attention, perception, memory, imagery, thought, etc. (Neisser 1967). This input-oriented way of viewing cognition is not surprising, since it reflects the roots of cognitive science in endeavors such as sensory physiology, psychophysics, and philosophical epistemology. Accordingly, action was not included in Neisser's list, in spite of the fact that much of the research he reported relied on reactions and reaction times. However, reactions were not regarded as targets of study in themselves. They served as mere indicators of cognitive states.

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science," Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18, series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

Since Neisser, the *zeitgeist* has changed and action has come to the fore. Today it is broadly acknowledged that cognitive processing is intimately intertwined with action processing. As a result, researchers claim that the study of cognition needs to be extended to also include the study of action. This is what we may call the action extension view of cognition.

While I certainly want to support an extension along these lines, my proposal here is to suggest a still more radical move. Numerous authors claim that the chief adaptive function for which minds/brains evolved and have been optimized pertains to smart action rather than true cognition. To put it briefly, smart actions let the animal do right things at right times, aiming at altering conditions in accord with currently given inner and outer circumstances. If this is true, cognition is a secondary, subsidiary function: true cognition helps with smart action but it is not the proper function for which mind/brain systems are optimized. If one takes this view seriously, the science that one needs to strive for is *action science*: How is smart action possible, and what can cognition contribute to it? A move in this direction takes us from action extensions of cognition to cognitive contributions to action. Moreover it allows us to view action as a target of study in its own right, including both its cognitive and noncognitive foundations.

The claim that action builds on noncognitive foundations can be understood in two different ways. One reading implies that action builds on noncognitive foundations that cannot be captured by the theoretical language of representationalism. This is what various brands of enactivism claim. The other reading invokes the notion that action goes beyond cognition in a descriptive, but not theoretical sense. This reading maintains that actions exhibit features that do not come into view from the perspective of an input-oriented approach, without, however, implying that these features cannot be captured by the theoretical language of representationalism. This view, which I am following here, considers representationalism a useful framework for all action science, powerful enough to capture both the cognitive and noncognitive foundations of action.

In what follows I sketch two basic ways of action science: individual and social. The individual way studies actions as tools of control (i.e., of altering events in accord with demands, needs, and desires). The social way studies actions as tools of interindividual alignment (i.e., of coordinating own and foreign actions).

The Individual Perspective: Action for Control

The notion of adaptation provides a convenient starting point for understanding what it means for animals to control their actions. Adaptation can be studied at two levels: structural dispositions and functional interactions. At the structural level, we may study how long-term dispositions match global, long-term conditions in the environment. Here we speak of a species being adapted to

From “The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,”
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

its environment. At the functional level, we may study local matches between fluctuating inner states and fluctuating outer conditions in individuals. Here we speak of their capacity to adapt, or adapt themselves, to changes in bodily states or environmental conditions.

When we speak about action control, we focus on particular kinds of such local adaptations—those that involve bodily movements as a means for altering environmental conditions or bodily states. The adaptive value associated with these movements does not reside in themselves but in the alterations of inner and outer states achieved through them. Accordingly, we think and talk about these movements in terms of the goals toward which they are directed. In other words, we consider them actions, not just movements. Actions are segments of bodily activity that converge on some goal state. When a lion chases a zebra, that action terminates when the lion eventually catches it. Likewise, when someone hammers a nail into the wall, that action terminates when the nail is eventually embedded in the wall.

Modes of Control

How are means and ends related to each other? Does control proceed from means to ends or from ends to means? In addition, how are means and ends related to circumstances under which the action is performed? Answers to these questions can be divided into two major camps: bottom-up and top-down control. Both camps agree that actions have the potential to achieve desirable outcomes in terms of the current needs and interests of individuals, but they disagree on the machinery involved. Bottom-up control posits that ends follow from means; that is, goals are attained as outcomes of given actions. Top-down control posits that means follow from ends: actions get selected to achieve given goals.

Bottom-Up

The notion of bottom-up control captures the idea that goal-directedness is an emergent property of the workings of control systems whose operation does itself not draw on goals or goal representations at all. Bottom-up control posits that goal-directed behavior can be explained as a consequence of currently given states of affairs (or representations thereof), with no role being played by future intended states of affairs (or anticipatory representations thereof). Explanatory strategies along this line have been advanced by several classical approaches (e.g., Skinner 1953; Thorndike 1911; Tolman 1959). These approaches have devoted much effort to explaining purposeful behavior without purposes and goal-directed action without goals.

How can this happen? A useful framework is provided by the technical metaphor of control. Engineers furnish technical systems with controllers:

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

computational devices that determine, for each configuration of current circumstances, which action to take in order to establish or maintain certain inner and outer circumstances. Likewise, we may think of animals as being furnished with controllers. In this case, controllers determine for the animal, under each configuration of current inner and outer circumstances, which actions to take to establish or maintain satisfactory or desired inner and outer circumstances.

These devices can be characterized in terms of the output they provide, the input they require, and the algorithms on which they rely. On the *output* side, controllers steer bodily movements suited to modulate inner or outer conditions. These movements may, for instance, act to alter environmental conditions (e.g., a frog catching a fly) or to move the body relative to the environment (e.g., navigating around an obstacle). To perform these computations properly, controllers need to be informed, on the *input* side, about the configuration of current outer and inner circumstances (i.e., the position of the fly or the obstacle). Third, controllers need to dispose of *algorithms* for input interpretation and output generation. The operation of these algorithms depends on event knowledge and action knowledge. Event knowledge reflects what the controller "knows" about the to-be-controlled events. Conversely, action knowledge reflects what the controller "knows" about possible actions and interactions with these events.

At first glance, this scheme looks like a linear sequence that leads from event interpretation to action generation (and, hence, from stimuli to responses). Yet, to complete the picture, we need to realize that these sequences are embedded in cyclical interactions between controllers and target events: actions alter events which in turn give rise to new cycles of interpretation and generation, and so forth. Still, within each given cycle, action generation always depends on event interpretation.

How is event and action knowledge acquired? Different kinds of learning mechanisms have been proposed, both on the ontogenetic and the phylogenetic scale. They all share the same functional logic—the logic of trial, error, and success. On the ontogenetic scale, learning applies to individual animals. Here one of the most powerful mechanisms for the acquisition of mappings between stimulus events and responding actions relies on the production of behavioral variation plus subsequent selection of those variants which prove to be successful. Eventually, the animal will home in on stimulus-response mappings that turn out to be successful in terms of its current needs under current conditions. On the phylogenetic scale, learning applies not to individuals but to populations and their gene pools. Here, through production of genetic variation and selection, populations may "learn" that some stimulus-response mappings are more advantageous than others, and eventually they will come to incorporate those mappings into their gene pools. Such phylogenetic learning may require hundreds or perhaps thousands of generations, based on variation, selection, and survival of the fittest mapping.

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

In sum, bottom-up control explains the occurrence of goal-directed actions as a causal consequence of two kinds of factors: currently given circumstances and outcomes of previous learning. Outcomes of learning are stored in the controller's machinery for event interpretation and action generation. This generates actions previously proven to be successful under similar conditions. Naturally, such actions are often meaningful and may even, in many cases, look as if they were goal-directed in a literal sense (i.e., guided by explicit intentions).

Top-Down

Bottom-up control is, in essence, a stimulus-triggered, stimulus-guided affair. Since cognitive scientists like stimuli and their consequences, this explains why generations of them have tried to push the limits of this control mode as far as possible. Still, there is no reason to believe that bottom-up control is the only game in town. Top-down control may perhaps be a late arrival in the evolution of action control, but in humans it is certainly in place and plays a strong, if not dominant, role. Top-down control posits that means follow from ends; that is, movements derive from goals. Put somewhat paradoxically, top-down control implies that movements are selected on the basis of desired or intended circumstances—circumstances which can only be attained through (and hence after) performing those movements.

Whether this sounds like magic depends on what we are considering: circumstances in the world or a controller's representations thereof. When we talk about circumstances in the world, we are in fact invoking a teleological explanation which claims that certain movements are performed at t_1 to achieve certain states of affairs at t_2. Yet, when we express the same relationship in terms of representations, the paradox goes away. At that level, movements performed at t_1 go back to desires and intentions at t_0. What we thus require at that level are goal-related representational states such as desires and intentions, which act as temporal and causal antecedents of movements.

What does top-down control require? To build a top-down controller, we need to take two major steps beyond the scheme for bottom-up control. First, we need to furnish our device with the ability to create action goals; that is, to form representations of events that are independent from the configuration of currently given circumstances. Second, we need to furnish goal representations with the power to make an impact on action selection and generation. The first step invokes dual representation, the second ideomotor control (cf. Prinz 2012, chap. 7).

The notion of *dual representation* refers to the ability to maintain two parallel streams of event representations and keep them apart: one for events that are currently present and another for events not present in the current situation. Whereas the first stream is fed from external sources (stimulus information

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

pertaining to the current situation), the second emanates from internal sources (memory information pertaining to events beyond the current situation).

When applied to action control, the implementation of dual representation provides a first step toward top-down control. Top-down controllers need to be capable of tracing perception-based and intention-based episodes. Whereas perception-based episodes address ongoing and upcoming events, intention-based episodes address to-be-attained events. One of the two streams thus refers to factual events that are actually happening, the other to fictitious events which the agent would like to make happen. Strict separation between these streams is required since mistaking fact for fiction is no less maladaptive than mistaking fiction for fact.

The notion of *ideomotor control* explains how goal representations contained in intentional episodes become functional for action control. Put in a somewhat old-fashioned language, ideomotor theory views actions as creations of the will (James 1890; Lotze 1852). Thus phrased, two conditions must be met to carry out a voluntary action: there must be a mental image of what is being willed, and conflicting ideas or images must be removed. When these two conditions are met, the mental image acquires the power to guide the movements required to realize the intention.

Ideomotor theory claims that the links between intentions and movements arise from learning. Whenever a movement is performed, it is accompanied by perceivable effects. Some are directly linked to carrying out the movement itself, such as the kinesthetic sensations that accompany each movement (resident effects). Others are linked to the movement in a more indirect way since they occur in the agent's environment at a spatial and/or temporal distance from the actual movement (remote effects).

In any case, the regularities between actual movements and their resident and remote effects are captured in associations. Thus, representations of movement outcomes become associated with representations of the movements leading to them. Once established, these associations work in two directions. One allows the anticipation of movement outcomes (i.e., to predict perceivable consequences from given movements). This is the case of forward-directed computation in the service of bottom-up control. The other allows for backward-directed computation in the service of top-down control (i.e., to select and generate movements required to achieve given intentions). Such backward computations guarantee that events which have been learned to go along with, or follow from, a particular action will hereafter exhibit the power to call that action forth.

In sum, ideomotor control relies on two principles: one for learning and one for performance. The learning principle claims that the system is capable of establishing associations between actions and their outcomes (both resident and remote). The performance principle claims that, once established, these associations can also be used in the reverse direction (i.e., from outcomes to actions effectuating them).

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

Implications for Cognition and Action

What does it mean for cognition to be tailored to the needs of action and for action to be grounded in cognition? These issues have been broadly discussed over the past decades (cf. Braitenberg 1986; Hommel et al. 2001; Jeannerod 1997, 2006; Morsella et al. 2009; Neumann and Prinz 1990; Noë 2004; O'Regan 2011; Prinz 2012; Prinz et al. 2013; Prinz and Hommel 2002).

The first implication is that *cognition is for action*. While competing approaches make use of different theoretical frameworks for addressing these issues, they all share the basic idea that cognitive functions are embedded in an architecture for control. This idea can be read genetically as well as functionally. The genetic reading maintains that cognitive functions evolve for the sake of action control, on both phylogenetic and ontogenetic scales. However, the fact that evolutionary history has shaped the cognitive machinery to serve the needs of control does not necessarily imply that its online operation must always include elements of control. Accordingly, this reading does not imply that each and every act of cognition must be associated with an act of control. This is what the functional reading maintains. The genetic reading claims that acts of cognition always entail elements of control—be it in the weak role of associated extensions or the strong role of constitutive ingredients.

Both readings have important implications for research agendas in the respective fields of study. On the phylogenetic scale, the genetic reading has inspired neuroethological programs in the study of the natural history of mental functions as they emerge from architectures for sensorimotor interaction (e.g., Dean 1990; Gallistel 1980). The same applies to the ontogenetic scale, where the genetic reading has inspired psychological programs to study the construction of an architecture for cognition from basic interactions between perception and action (Bertenthal and Longo 2008; Piaget 1954; Thelen and Smith 1996; Vygotsky 1979).

The functional reading is associated with research programs which aim at demonstrating the secret workings of action in cognition (i.e., the implicit involvement of action control in perception, memory, and thought). In recent years, these programs have made much progress, as new technologies for online recording of brain activity have provided new tools for assessing the latent involvement of action in cognition. As a result, there is now convincing support for the functional reading of the cognition-for-action claim. As has been shown in a variety of experimental paradigms, input-related cognitive processing is intimately intertwined with output-related action processing (Barsalou 2008; Hommel et al. 2001; O'Regan 2011; Prinz et al. 2009; Viviani 2002). We may therefore conclude that action makes essential contributions to cognition—not only at the level of architecture construction but at the level of information processing as well.

A further implication is, of course, that *action relies on cognition*: If cognition is for action, then action must rely on cognition. This implication is

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

not trivial vis-à-vis the classical view that action is a thing that commences only after cognition has terminated. While this view maintains that the two draw on disjunct representational resources, the new view maintains that event interpretation and action generation draw on the same pool of knowledge resources. Actions are, accordingly, represented like any other types of events, so that representations are entirely commensurate (cf., e.g., Hommel et al. 2001; Jeannerod 1997; Rosenbaum 2009, 2013).

A final implication is that we need an *architecture for voluntary action*. This architecture must capture three basic segments of voluntary action: motivation, volition, and execution. Several proposals have been made concerning underlying mechanisms and dividing lines as well as transitions (de Wit and Dickinson 2009; Gollwitzer and Moskowitz 1996; Hassin et al. 2009; Heckhausen and Heckhausen 2008). One of the crucial features that the architecture must capture pertains to the origin of motives, goals, intentions, and their underlying representations. Importantly, these representations originate in the agent, not the environment, and the architecture must take means to keep them apart from representations of ongoing external events. A further important feature that needs to be covered by the architecture pertains to the dynamic nature of volition. Unlike percepts, memories, or thoughts which act as "cool" placeholders for things that are there, motives, goals, and intentions serve as "hot" placeholders for things that are wanted and desired and ways of getting at them. Importantly, both of these features cannot be captured by an extended architecture for bottom-up control.

At the same time, it should be clear that a machinery for top-down control requires a machinery that is nested within for bottom-up control. This is because goals and intentions in the mind can only be realized if agents are in a position to shape their actions according to current circumstances in the world. Accordingly, an architecture for voluntary action must build on communications between representations of wanted and given events. Top-down controllers can only work with built-in bottom-up controllers.

The Social Perspective: Action for Alignment

A peculiar functional condition arises when we turn to control scenarios in which individuals address other individuals' actions as events of control. When two (or more) people interact, each person can be seen to become involved in controlling the other's actions through his/her own actions. Here, event interpretation turns into action interpretation, based on event knowledge that pertains to foreign action. As a result, control draws on action knowledge on either side. Action generation relies, as usual, on knowledge pertaining to one's own action, but event interpretation now relies on action knowledge as well; namely, knowledge pertaining to foreign action.

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

This peculiar condition opens a unique opportunity for action alignment across individuals. Action alignment requires that knowledge resources for own and foreign action are integrated and combined, so that perception of foreign action and production of own action draw on common resources. For social animals, like humans, this functional condition provides an invaluable asset since it offers a direct, effortless way to align own with foreign and foreign with own action.

Common Coding

The notion of common coding invokes the idea that production of one's own action and perception of foreign action draw on common representational resources. In other words, tokens of one's own action get entries in the same representational domain and on the same dimensions as tokens of foreign action. Common coding thus makes it possible to assess similarity relationships between one's own and foreign action. For instance, as concerns production, own action may replicate or continue foregoing foreign action. Likewise, as concerns perception, foreign action may be understood to replicate or continue foregoing own action. In the first case, perceived action primes production of own corresponding action, whereas in the second case, production of own action primes subsequent perception of corresponding foreign action.

There is now ample evidence that this principle is instantiated in human minds and brains. Evidence comes from various fields of study. Results from behavioral experiments on action imitation, action induction, and perception/action interference have lent support to a strong role of similarity between own and foreign action (Prinz 2012, chap. 5). Parallel to this, numerous electrophysiological studies on mirror neurons and mirror systems in the monkey brain provide what may be considered an existence proof of shared representational resources for action production and perception (Rizzolatti and Sinigaglia 2008). Similarly, brain imaging studies on humans have shown that shared brain circuits may not only be involved in processing own and foreign actions but own and foreign sensations as well as emotions (Keysers and Gazzola 2009; Schütz-Bosbach and Prinz 2015).

In spite of the overwhelming evidence, common coding is often considered a strange and somewhat mysterious notion. Two brief remarks may help to demystify this perception. First, as has been pointed out in several places, the emergence of shared resources for own production and foreign perception can easily be explained in terms of classical principles of association and connectionist models instantiating them (Cook et al. 2014; Keysers and Gazzola 2009; Keysers and Perrett 2004; Pulvermüller et al. 2014). Second, common coding is, of course, not everything. The claim that production and perception draw on shared representational resources does not imply that *all* resources on which they draw are shared. Claiming shared resources at one

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

level is entirely compatible with acknowledging unshared resources at other processing levels.

Social Mirroring

Common coding devices provide powerful tools for interindividual alignment. Since they use the same resources for representing own and foreign action, they offer themselves for both: own action resonating to foreign action and understanding foreign action as resonating to own action. To grasp what this alignment entails, let us look briefly at episodes of social mirroring (Prinz 2012, chaps. 4–6; Schütz-Bosbach and Prinz 2015).

Episodes

Social mirroring episodes have two sides: the target individual whose acting is being mirrored and the mirror individual who mirrors the target's action. The mirror individual functions for the target individual like a mirror in the target individual's environment. Here, two basic types of mirroring episodes can be discerned: reciprocal and complementary.

In episodes of reciprocal mirroring, the target individual sees her own action imitated or replicated by the mirror individual. In such a setting, the mirror individual acts as a mirror for the target individual in a more or less literal sense. Still, social mirrors are fundamentally different from physical mirrors. Even if the mirror individual tries to provide as-perfect-as-possible copies of the target's action, these copies are always delayed in time and their kinematics will never be as perfectly correlated with the target's acting as spectacular images are. We may speak of reciprocal mirroring as long as the target is in a position to understand that the mirror's acting is a delayed copy of the target's own preceding action. Hence, the constitutive feature of reciprocal mirroring is the target's *understanding* of the mirror's action as a copy of the target's foregoing own action.

In episodes of complementary mirroring, the target sees her own action continued and carried out by the mirror individual, rather than replicated. This, of course, is entirely different from what physical mirrors do. Nevertheless, what complementary mirroring has in common with reciprocal mirroring is that the mirror individual's action is strongly contingent upon the target's preceding action and that the target may perceive and understand this contingency. In this case, too, the reach of effective mirroring goes as far as the target is in a position to *understand* the mirror individual's acting as a continuation of her own acting.

Not surprisingly, episodes of action-based mirroring play an important role in interactions between young infants and their caretakers. Babies and their caretakers often find themselves involved in episodes of action-based,

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

proto-conversational interaction and communication (Trevarthen 1998). They take turns in imitating or continuing each other's action. Most of this work concentrates on the baby in the role of the mirror (i.e., mirroring the caretaker's actions), not the target (i.e., perceiving herself being mirrored by the caretaker). To understand the power of these episodes for interindividual alignment, however, we need to take both roles into account: not only that of the mirror individual but also that of the target individual who perceives her own actions being mirrored by the other.

Action-based mirroring is not limited to interactions with young infants. Mirroring episodes are likewise widespread among adults. For instance, an individual may, in a conversation, shrug his arms in response to his conversation partner doing the same (reciprocation). Likewise, an individual may take up another individual's action when the other temporarily withdraws (continuation). Such mirror episodes may often reflect automated habits rather than controlled and deliberate actions. Still, they act to align individuals through production and perception of closely related actions. Mirror episodes help individuals match their own actions to others' actions and others' actions to their own.

Practices

Social mirroring depends on functional mechanisms that instantiate common coding. At the same time, it also depends on social practices in which individuals must engage to exploit the potential that is inherent in these mechanisms. These practices can be viewed at a local and a global level. At the local level, episodes of social mirroring require that two individuals interact in a particular way. They need to engage in *mirror games*. Mirror games are designed to align actions through mutual reciprocation and continuation, deliberately or automatically. One may speculate that engaging in such mirror games is a human universal, at least as far as interactions with young babies are concerned.

At a more global level, mirror games are embedded in *mirror policies*. These policies reflect strategies that govern individuals' participation in mirror games. Individuals may be quite selective in playing these games. They may mirror some behaviors but not others. They may engage in mirror games under some circumstances but not others. Most importantly, they may be selective with respect to the target individuals to whom they grant their mirroring. For instance, they may mirror their children, family, and peers but not strangers, disabled individuals, or the elderly. Mirror policies thus act to induce both social assimilation and dissimilation, and eventually even discrimination.

Conclusion

An action-oriented view of cognition is nice to have, but, at the same time, it is not enough. Viewing action as an extension of cognition captures the

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.

functional logic of bottom-up control and the relationships between cognition and action implied in that control mode. However, the cognitive approach to action fails to capture essential features that make up the signature of human action (e.g., top-down control and action alignment). Top-down control operates within individuals and requires a framework that includes the formation of motives, goals, and intentions as precursors of action selection and execution. Action alignment operates between individuals and necessitates a framework that includes shared representations for perception and production.

To accommodate these features, we need to embark on a path that will lead us to *action science*. The study of action must address all aspects and functionalities of action, including its cognitive as well as noncognitive foundations.

From "The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science,"
Andreas K. Engel, Karl J. Friston, and Danica Kragic, eds. 2016. Strüngmann Forum Reports, vol. 18,
series ed. J. Lupp. Cambridge, MA: MIT Press. ISBN 978-0-262-03432-6.