

How Motion Reveals Intention

Categorizing Social Interactions

Philip W. Blythe
Peter M. Todd
Geoffrey F. Miller

Do not go where the path may lead, go instead where there
is no path and leave a trail.

Ralph Waldo Emerson

If you should encounter a mountain lion while hiking in the Sierra Nevada mountains of California, there are two things you must not do, according to the Mountain Lion Foundation: turn your back on the animal or run away. Either of these behaviors would trigger the lion's predatory chase behavior, transforming you from startled hiker into potential prey. It is possible to avoid becoming prey by denying the lion's perceptual system the cues that normally accompany being a mealtime animal. Knowing how other creatures categorize behavior based on motion cues could thus make the difference between life and death.

Humans are also very adept at making such behavioral judgments from simple motion patterns: When two children run across a field, their parents can distinguish in a brief moment whether they are playing or fighting in earnest. When a pigeon twirls and struts before another, who ignores this display and turns away, we can quickly tell that the first is trying unsuccessfully to court the second. In situations such as these, we as outside observers can often decide what is going on—who is doing what to whom—based just on the motions of the two organisms relative to each other. Moreover, the human or animal participants in such cases can also tell what kind of interaction they are having, again using motion cues. How can organisms categorize behaviors based solely on observed motion patterns? In other words, how can humans and other animals translate from the domain of pure physical movement into the domain of

animate intentions and desires such as chasing, playing, and courting? What are the most important categories of behavior? What motion cues are most useful for distinguishing them, and how do we exploit these cues to achieve fast, accurate recognition of other agents' intentions?

The answers to these questions come in three stages. First, we must specify the major behavioral functions that an organism's whole-body motions can fulfill, such as chasing, fighting, or courting. Second, we must determine the observable motion cues that allow us to categorize which of these functions an organism is performing at a particular time. And third, we must find the cognitive algorithms that can be used to make this functional categorization based on motion cues.

These stages require very different research methods. The first stage entails an evolutionary task analysis of motion itself, asking what basic reasons animals have for moving at all, given the demands of survival and reproduction. The second stage requires analyzing the information structure of the environment, determining what useful motion cues can be recovered by observing a moving animal, and which motion cues are most useful in distinguishing the animal's reasons or intentions for movement. The third stage includes comparing the performance of different possible decision algorithms that map motion cues onto the basic functional categories of animal movement, to narrow the search for algorithms that animals and humans might actually use. Because many algorithms attain similar levels of accuracy when inferring intention from motion (see chapter 7 for more on this problem of algorithm comparison), we will not aim to be conclusive on this point. Rather, our main concern will be determining the functionally important motion categories and the available motion cues that can be used to distinguish between them.

It is not at all obvious how to get from motions to intentions. The difficulties are clear in comparing the standard dictionary definitions for motion terms with those for higher-level intentional behaviors. For example, the *Oxford English Dictionary* defines "run" as "to move the legs quickly (the one foot being lifted before the other is set down) so as to go at a faster pace than walking"; whereas it defines "chase" as "to pursue with a view to catching." Thus, "run" is defined by reference to observable motion cues, whereas "chase" is defined by unobservable intent with respect to a future goal concerning an unstated object. Nothing in the definition of "chase" suggests how a naive observer could distinguish chasing from any other intentional category of movement. Because we cannot directly perceive movement intentions and goals, and the targets of movement may even be hidden, we must infer them all indirectly using motion cues that can be directly perceived—a clear case of inference under uncertainty as described in chapter 1.

A general description of how we can infer goals and intentions from observable behavior would constitute a rather complete understanding of the human "theory of mind" (Baron-Cohen, 1995) and much of social psychology more generally. We are not attempting to provide that general un-

derstanding here—we focus only on the very simple case of using motion cues to categorize the most typical forms of intention-driven interaction between two agents. This special case is, nonetheless, biologically important to virtually every animal species, and psychologically important as the most elementary level of social cognition. Animal species differ in body structure, modes of locomotion, environmental constraints on motion, motion perception abilities, mating systems, and positions in the local food chain. However, by keeping our analysis of the special case of intention from motion sufficiently abstract, we hope to identify motion cues and intention categorization strategies that are general enough to apply to many species despite these differences.

We must still constrain the special case explored here by limiting the types of intentional motion categories, cues, and cue-processing algorithms we will consider. First, we assume that humans and other animals have domain-specific motion perception and intention inference adaptations that are attuned to ancestrally typical categories of motion patterns and intentions. These typical patterns, associated with intentions such as pursuit, evasion, fighting, courtship, and play, arise because there are just a few basic survival and reproductive goals for animate motion. The fact that there are only a few reasons why one animal moves relative to another animal makes our job as psychologists much easier, transforming an unconstrained inference task ("Why is that animal moving?") into a simple categorization decision ("Is that animal chasing, fleeing, fighting, courting, or what?").

Second, we limit our consideration of the vast range of possible information that could be used in judging the intentions of other agents. Potential cues include not only the motion of whole bodies in relation to other bodies—as when mountain lions infer edibility when humans turn from them—but also motions of one body part in relation to another, including threat and submission postures, facial expressions (e.g., snarling dogs or laughing children; see Darwin, 1872/1965) and those micromovements of throat, tongue, and lips that result in articulate human speech. Here we focus on just the first, simplest form of information for judging intentions: the overall motion trajectories of two whole organisms in relation to each other. Third, the inference algorithms that might use this motion information are also limitless, but again we are interested in simple possibilities: fast and frugal heuristics that use as little of the available information as they can to make their decisions and operate as quickly as possible. Particularly when judging the intentions of (possibly hostile) others, it is important to be able to make decisions quickly and with just the information at hand, rather than waiting until all possible evidence has been gathered and the mountain lion has pounced.

We begin this chapter with an evolutionary analysis of how animate intentions could be inferred from motion cues. We then develop a novel experimental method for studying how humans make these inferences. The first step is to gather ecologically representative examples of the six

most typical animate motion patterns, by having participants play interactive computer games that require them to pursue, evade, fight, court, be courted, or play with each other. The participants control on-screen bugs with realistic motion physics using a mouse. The resulting motion patterns are recorded and presented to another set of participants, who try to infer what the bugs are trying to do to each other. We determine which intentions are confused most frequently with which other intentions. We then construct and test various models of how people could categorize these intentions given some simple motions cues that can be computed directly from the motion patterns. The goal is to identify both useful objective motion cues and simple heuristics that can process them to infer animate intentions. We view this intention-inference task as the foundation for more advanced forms of social cognition and attribution.

Previous Research on Inferring Intentions from Motion Cues

The question of how we infer intentions from motion cues seems fundamental for motion perception and social cognition, but it has rarely caught the attention of mainstream psychologists. The few exceptions are those who take an ecological view of perception, which motivates more direct study of the structure of the observable environment in relation to an agent's goals.

The earliest example was perhaps Fritz Heider, who set out to study the perception of social events by studying the particular stimuli that led to different attributions. His experiments in the 1940s with Simmel (Heider & Simmel, 1944) demonstrated that people spontaneously attribute intentions and personalities even to featureless geometric figures such as dots and triangles, if these figures move around in a cartoon film according to patterns reminiscent of animals courting and fighting. Heider and Simmel did not explicitly identify the motion cues that provoke these interpretations. But they had enough tacit knowledge of these cues that they could lead observers to view a roving triangle as a scheming villain, or a flitting disc as a fickle adulteress, through nothing more than relative motion in a simple environment.

Interest in this area fell dormant for a quarter of a century, aside from research in ethology specifically aimed at uncovering the motion cues that animals make to signal their intentions to each other, such as wolf pups bowing to signal playful intent (see Fagen, 1981). In the 1970s, J. J. Gibson's (1966, 1979) research on "direct" perception of ecologically important visual cues of motion inspired a number of studies on the perception of gait and other forms of biological motion by Kozlowski and Cutting (1977; Cutting & Kozlowski, 1977). These studies were aimed at identifying various dynamic and structural invariants in "point-light displays," films of people walking and acting in total darkness, with only small light

sources affixed to their limbs and bodies. Even such impoverished stimuli, consisting solely of cues to movement, contained enough information to allow the visual system to make sophisticated inferences (e.g., that is a man lifting a heavy object, or that is a woman walking).

In a separate research tradition, French psychologist Albert Michotte investigated how people perceive cause and effect based on motion cues (Michotte, 1963). His project addressed mainly the psychology of causality (inspired by the philosophy of Aristotle and Hume), but Michotte was intrigued by his finding that people tend to interpret rectangles that interact without colliding as if they were animals or humans (Michotte, 1950), similar to Heider and Simmel's findings of rampant anthropomorphism. His work influenced modern studies on the innate preparedness of infants to perceive animacy (Premack, 1990; Spelke et al., 1995), and on the perception of animate motion more generally (Freyd & Miller, 1992; Gelman et al., 1995; McBeath et al., 1992). Michotte's emphasis on causation has been followed by recent studies that focus more on how people distinguish animate from inanimate motion, rather than how people distinguish between different types of animate motion.

This preoccupation with the animate-inanimate distinction is unfortunate because it stops short of what is arguably the more interesting question: Recognizing a moving object as animate is only the first step toward recognizing its intentions, which are what really matter for deciding what to do in response to it. Zebra herds have to live with lions hanging around at their watering holes, fleeing only when the lions show a real intention to chase them. The crucial distinction for a zebra is not between animate motion (a moving lion) and inanimate motion (a breeze stirring tawny grasses), but between animate motion that is relevant (a hungry lion at full sprint heading straight for you) and irrelevant (a fat lion chasing a fertile lioness). The point of animate motion perception is to guide one's own animate motion. The animate-inanimate distinction is just one relatively weak cue for deciding what to do; much stronger, more informative cues are available and so should be used to help determine other agents' intentions and one's own reactions. With that view, this chapter examines some basic goals of animate motion, some associated motion cues that may be general across species and ecologies, and some simple heuristics for categorizing intentions based on those motion cues.

The Basic Goals of Animate Motion

Moving has energetic costs, so animals are expected to move only when the fitness benefits of movement exceed these costs. The major fitness benefits of moving, as of any other behavior, are survival and reproduction. Thus, it is possible to deduce a few paradigmatic goals of animate motion from the fundamentals of natural selection and sexual selection.

Animals evolve to interact adaptively with various "fitness affordan-

ces” in their environments—things that are likely to affect the replication of their genes (Miller & Freyd, 1993; Todd & Wilson, 1993). Positive fitness affordances, such as food and sexual mates, promote survival or reproduction. Negative fitness affordances, such as predators, pathogens, parasites, and sexual competitors, interfere with survival or reproduction. Animals evolve sensory-motor systems to approach the positives and avoid the negatives. If two animals offer the potential of mutually positive yields, mutual approach usually results (e.g., in symbiotic relationships). If they threaten mutually negative yields, then mutual avoidance results. Movement patterns in these cases are just solutions to various positive-sum coordination games. The more interesting case is when animals present affordances of opposite sign to each other, that is, when one animal wants to be near the other, but the other wants to be far away. The zebra is a positive (food) affordance for the lion, but the lion is a very negative (death) affordance indeed for the zebra. Such conflicts of interest lead to more complex interactions, transforming simple approach into relentless pursuit, and simple avoidance into desperate evasion (Miller & Cliff, 1994).

From these arguments, it follows that the fundamental categories of two-agent animate interaction are mutual approach (boring), mutual avoidance (also boring), and pursuit and evasion (interesting). In the survival domain, pursuit and evasion usually occur between predators and prey, or between fighting conspecifics. In the reproductive domain, males usually pursue and females usually evade, at least for a while (Andersson, 1994; Darwin, 1871). Thus, almost every animal will need to master some subset of five basic categories of animate motion: pursuing, evading, fighting, courting, and being courted. To these categories we also add a sixth, play, which is widely used, especially by mammals, to learn mastery of the other five movement types. Although not all species need to exhibit the full set of these behaviors, they generally need to be able to recognize each of them. We will now consider the six behavior types in turn.

Pursuit Animals move toward objects they desire. If the desired object is inanimate, we have a degenerate case of goal-directed behavior. But if the object is animate and does not want to be exploited as a fitness affordance (e.g., as food or as a mate), then it will move away (evade). The simplest pursuit strategy is to point one’s front in the direction of the desired object and charge at top speed toward it, changing one’s direction if the object deviates to the left or right of one’s current heading. However, pursuit can be more efficient by predictively taking into account any environmental constraints on motion (obstacles and boundaries) and the pursued agent’s own heading, movement, and intentions.

Evasion Animals move away from things that threaten them. Again, if the threatening object is inanimate, we have a degenerate case of obstacle avoidance, or one-step “evasion.” If the threat is animate, however, and

does not wish to be evaded, then it will pursue, and sustained evasion becomes necessary. Evasion often favors strategies of deceptive feints and lunges and unpredictable, “protean” zigzagging (Driver & Humphries, 1988; Miller & Cliff, 1994), as well as avoidance of environmental traps or cul-de-sacs.

Fighting Animals of the same species often fight over fitness affordances such as territories, resources, sexual mates, and social status. Fights can be tricky to decipher because both animals must combine pursuit and evasion, attack and defense, in a way that intimidates or overcomes the opponent, without injuring or killing themselves. Because animal bodies are heterogeneous, with some parts specialized for attack and other parts vulnerable to injury, fighting usually includes a great deal of precise, dynamic body positioning in relation to one’s opponent.

Courting Animals (usually males) move toward members of the opposite sex (usually females) with whom they wish to mate (Andersson, 1994; Darwin, 1871). But because selective mate choice is almost always imposed by the opposite sex (usually females), simple approach is almost never enough. Instead, mate-seeking animals often evolve extremely complex courtship behaviors with special features designed to display their health, strength, size, status, intelligence, or creativity (Andersson, 1994; Miller, 1998). These displays are usually produced close enough for the desired mate to perceive them, but not too close, lest the audience is scared off rather than enticed. After some display time, ranging from seconds (for some insects) to years (for some humans), if the desired mate signals her (or his) interest, the final approach and copulation can occur.

Being Courtied Animals sought after as mates (usually females) have strong incentives to select among their suitors quite carefully, because they usually have more at stake when mating than the suitors, and the genetic quality of the suitors they choose to mate with will determine half of the genetic quality of their offspring. Random mating is stupid mating (see chapter 13 for more on nonrandom mating and mate search). The task when being courtied, then, is to express enough interest to elicit informative courtship behavior from various suitors so their mate quality can be assessed, but not so much interest that the courter skips courtship altogether and tries to move straight to copulation. Thus, being courtied requires a delicate balance between interactive encouragement and coy reticence. Courtied animals usually maintain enough proximity to their suitors that they can determine the suitor’s quality, but do not get close enough to risk sexual harassment or rape.

Playing Play often comprises a variety of actions through which young animals can practice all of the above movement types, using play signs to indicate that they are pursuing, evading, courting, or fighting without real

lethal or sexual intent (Fagen, 1981). In basic play, animals repeatedly switch roles between pursuer and evader, or attacker and defender. In more complex play characteristic of large-brained primates, animals may interact in more abstract ways with imaginary partners or mutual mimicry.

These six goals for animate motion are not intended to be the building blocks in some sort of universal movement grammar. There are other movement goals that cannot be reduced to these categories, such as parental protection of vulnerable offspring from conspecifics or predators, foraging for inanimate food items or nest materials, migrating to new habitats, grooming oneself or others, flocking, mobbing predators, and so forth. We simply start with the most obvious basic cases of two-agent interaction. Future research will, we hope, analyze the motion cues that help distinguish these additional categories of behavior.

Eliciting Motion Trajectories for the Six Typical Behaviors

Following in the tradition of early work on attribution of intentional behavior to simple moving stimuli by Heider and Simmel (1944), we sought to study how people categorize the functional intentions of two interacting organisms based solely on their trajectories through space relative to each other. This required collecting some trajectories representative of each of our six functional categories that could be measured objectively to find useful motion cues, and that could be presented in a standard format as stimuli to human participants. But such trajectory data is rather difficult to come by. The literature in both biology and psychology offers many studies of long-range animal navigation, migration, and commuting on the one hand, and small-scale limb movements on the other. However, there is little publicly available data on behavioral trajectories between these extremes. So we decided to generate our own sample trajectories, using the animate agents we had in ready abundance: university students.

The question was how to obtain ecologically representative samples of the six typical behavioral categories. Despite generous research resources, it proved infeasible to attach radio tracking beacons to participants' heads, record them from satellite observations, and wait for some good examples of pursuit, fighting, courtship, and play. Though overhead video recordings from the Oktoberfest near our Munich laboratory would have captured instances of all of these behaviors, the trajectories would have appeared somewhat distorted by the potent Weissbier.

Instead, we had pairs of participants interact through a computer network, generating the trajectories by instructing them to play various movement games. Each participant sat before a computer and used a mouse to control the motions of a simple buglike creature displayed on the computer screen. Participants could see their own bug and another bug, which was controlled by another participant in another room at another com-

puter. Both bugs are displayed on-screen in a featureless rectangular environment bounded by walls, without any obstacles, viewed from an overhead perspective (see figure 12-1). We engaged 10 pairs of participants to perform the six fundamental behaviors in this simple computer-mediated interaction game.

The bug did not act like an ordinary on-screen cursor that directly reflects hand movements across the mouse pad. Rather, the bug was constrained by some simulated physics. This included momentum, which produced slow acceleration and deceleration of the bugs; collision dynamics, which made the two bugs bounce off each other and the walls; and a top speed at which the bugs could travel. The participants' mouse movements were essentially treated as targets for where the bug should be heading next (for details, see Blythe et al., 1996). These semirealistic physics made the movement games challenging, enhanced the perception of animacy, and, we hope, made the resulting motion trajectories more ecologically representative of natural animal movement.

The bug form was chosen because it looked more interesting than a dot. As discussed earlier, previous experiments on motion perception suggested that the attribution of animacy to a moving object appears to be a natural tendency for humans, whether the object resembles a human form




Figure 12-1: The two-bug view that each participant saw when generating motion trajectories, captured at a single instant in time. Each bug is a different color.

or a geometric shape (as found by Heider & Simmel, 1944; Michotte, 1963; Rimé et al., 1985). However, we thought the bug form would stimulate participants' interest without biasing their behavior in anthropomorphic ways, as a human-shaped icon might. It was also important to use a figure with a clear front and back end (as opposed to a circle, for instance) so that orientation would be unambiguous to both parties involved in the motion games. Furthermore, as biologist J. B. S. Haldane is reputed to have observed, "to a first approximation, all animals are beetles."

During each experiment, two participants unknown to each other were put into separate rooms with computers and were guided by on-screen instructions to practice and play the appropriate movement games in sequence. The participants were initially given a two-minute practice period to learn how to control their bugs using the mouse. Pilot studies showed this practice period sufficient for attaining a reasonable skill level, given that most participants had substantial experience with computer mouse controls. Following this, they participated in six trials of two minutes each.¹ The asymmetric trials of pursuit-evasion and courting-courted were duplicated with roles reversed, while fighting and play were only performed once by each participant pair.

The two participants, here A and B, were instructed to play the six movement games as follows:

1. A pursues B: Participant A was instructed to move his or her bug to intercept the other player's bug as quickly and as often as possible. Participant B was instructed to try to avoid being intercepted at all times.
2. B pursues A: This situation simply reversed the roles of pursuer and evader between participants A and B.
3. A courts B: Participant A was instructed to move his or her bug so as to court the other bug, by interacting with it in any way that it might find interesting, exciting, or enticing. Participant B was instructed to play the role of being courted, moving his or her bug to show interest or disinterest, and to elicit further displays in any way desired.
4. B courts A: This situation reversed the roles of courter and courted.
5. Fighting: Both participants were instructed to attack the other bug from behind, while at the same time avoiding being attacked. Specifically, they were instructed to try to strike the other bug's rear end with their bug's front end, at the same time avoiding the attacks of the other bug trying to do the same to them. This type of fighting resembles World War I aerial combat (bring one's front guns to bear on the enemy without his guns pointing at you) rather than boxing or wrestling.

1. The courtship trials lasted three minutes as it is a slower, more gradual behavior that required more time with some participants. In real life, while fights and pursuits may be over in seconds, we expect courtship to take rather longer.

6. Playing: Both participants were instructed to play with the other bug in whatever manner they wanted.

These movement games were obviously underspecified: We gave no feedback, offered no monetary incentives for performance, and allowed almost no practice time. We were not interested in studying the long-term equilibrium strategies for these dynamic two-person games, but rather in using the games as a quick and easy way of generating ecologically representative motion trajectories for six typical kinds of animate interaction. We expected that participants would have relatively stable and readily applied notions about what movement in each category looks like, and this is the knowledge we wanted to elicit. For pursuit, evasion, and fighting it would have been possible to specify each participant's payoffs exactly, but for courting, being courted, and playing, such payoff specification would have been difficult.

Visualizing the Resulting Trajectories

In each motion category trial, the computer recorded the movement trajectory of each participant's bug at a high temporal and spatial resolution. We then analyzed the resulting fine-grained trajectories in various ways, to see if there were any motion cues that could distinguish one behavioral category from another. First, we used a simple visualization method to look at the trajectories of two bugs interacting during one two-minute motion trial. Figures 12-2 to 12-5 show space-time plots of typical pursuit-evasion, courtship, fighting, and play trajectories generated by pairs of participants. The horizontal plane of the plot represents the horizontal and vertical positions of the creatures in the two-dimensional environment, while the vertical axis represents time during one 90-second trial segment. These plots reveal some basic information about the trajectories. Higher velocities (more spatial distance covered per unit time) result in flatter trajectories. Bugs that are not moving result in vertical line segments in these plots. Smaller distances between the bugs result in tightly intermingled trajectories.

Several features are immediately apparent in the plots that result from different movement games. In pursuit and evasion (figure 12-2), one sees very flat (very high speed) movements extending over a great area of the environment, contrasting sharply with the slower, more restricted movements during courtship (figure 12-3). Both pursuit-evasion (figure 12-2) and fighting (figure 12-4) show high speeds combined with large amounts of turning and looping. Fighting (figure 12-4) is distinguished by a smaller average distance between the two bugs, and by more tightly intertwined looping, with frequent contacts between the bugs (where their trajectories meet). In courtship (figure 12-3), the courter moves much more than the often stationary courted, sometimes circling, and occasionally engaging the courted in little bursts of pursuit and evasion. Only a few body con-

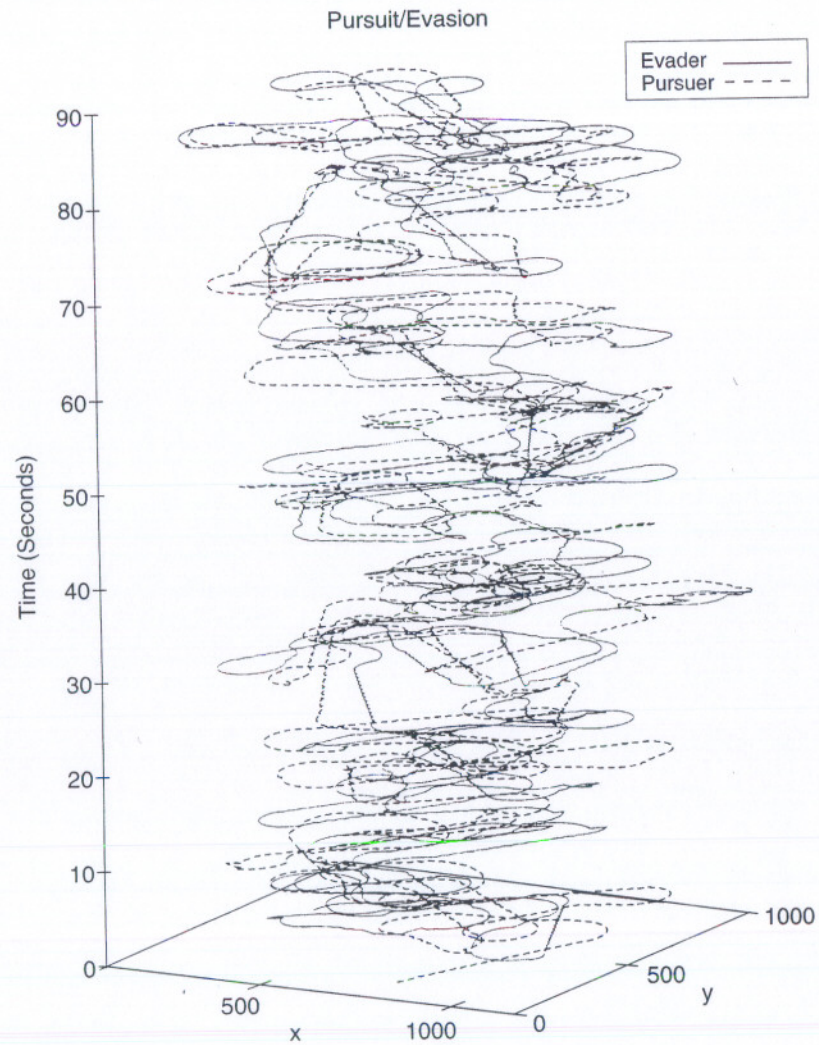


Figure 12-2: A sample trajectory generated when one participant's bug pursued the other, which evaded. Here, 90 seconds of interaction is represented, with time proceeding upward on the z-axis, and on-screen position of each bug plotted in the x-y plane. Note the generally high speed (flat segments) and large area covered.




Figure 12-3: A sample trajectory generated when one participant's bug courted the other, which responded to the overtures. Note the more elaborate motions of the courter, and occasional rapid fleeing of the courtee.




Figure 12-4: A sample trajectory generated when the two participants' bugs fought, trying to hit each other from behind. Note the high speed and high degree of looping.




Figure 12-5: A sample trajectory generated when the two participants' bugs played with each other. Here, one bug looped while the other dashed.

tacts occur in courtship. Play (figure 12-5) looks like a combination of pursuit, evasion, fighting, and courtship, combining looping, rapid dashes, and long still pauses. We will use these observations later to suggest some motion cues for distinguishing among these behavioral categories.

How Do People Categorize the Trajectories?

Our goal is to find motion cues that can distinguish exemplars of the six behavior categories and that may be used by humans in making such categorizations. But before we analyze the motion trajectories we have collected for different intentional categories, we must be sure that these trajectories do in fact contain the cues necessary to make the proper categorical distinctions. To find out, we tested whether a new set of participants could categorize the motion patterns accurately. Given example trajectories, how good are untrained people at assigning them to the behavioral task for which they were originally generated? Knowing this provides a rough check on the ecological validity of the trajectories, and makes it possible to investigate which categories are easily confused and which are easily distinguished. A model of animate motion categorization might perform well, but unless it makes roughly the same pattern of judgments, both correct and incorrect, as real human participants, it would not qualify as a good psychological model. Also, we can see what happens when we present the trajectories of both bugs together, versus only one by itself (see next section), to check the relative importance of single-bug motion cues versus interactive, relational cues for trajectory categorization. That is, by systematically removing some motion cues and recording drops in categorization performance, we can see which cues matter.

The Two-Bug Case: Relational and Individual Cues

In this categorization experiment, 10 participants saw portions of 30 motion trajectories recorded from the first experiment presented in randomized order, with one bug displayed in blue and the other in red. Participants were instructed to decide which one of the six behaviors the red bug was engaged in, as quickly as possible. As soon as they felt they could decide whether the red bug was pursuing, evading, fighting, courting, being courted, or playing with the blue bug, participants were to stop the trajectory playback and make a choice. To clarify this six-alternative forced choice, these trajectory-categorizing participants were shown the instructions given to the original trajectory-generating participants. However, these participants were not shown any examples of the motion categories beforehand, nor given any feedback on their selections. Nonetheless, participants' categorization performance changed little if at all over the course of the 30 trials, with equal accuracy in the first and second halves of the trials, suggesting that there was no significant learning effect.

Table 12-1: Participants' Categorization Confusion Matrix for Two-Bug Trajectories

Actual	Choice					
	Pursuing	Evading	Courting	Courted	Fighting	Playing
Pursuing	29	1	1	0	9	17
Evading	5	20	0	5	6	7
Courting	2	0	38	2	4	8
Courted	0	2	5	29	4	6
Fighting	5	6	3	3	10	23
Playing	2	0	9	10	7	21
Totals	43	29	56	49	40	82

Note. Each entry shows how many times participants chose a particular (column) category for a trajectory generated in some actual (row) category. The main diagonal (in bold) shows correct categorizations.

(Response times were also recorded for each judgment for later comparison among different categorization models, but we will not discuss these data in this chapter.)

Overall, participants selected the originally intended motion category in nearly half (49%) of the trials. This is well above chance (about 17%) for six-alternative choice. (In comparison, three expert participants who were familiar with the task and had previously seen many trajectories—the three authors—achieved 72% matches.) To see where participants erred, we can construct confusion matrices showing which categories were most often confused with each other. Table 12-1 shows such a confusion matrix for the 300 participant categorizations. The rows denote the actual intention-category instruction given to the trajectory generators, and the columns denote the choice decisions made by the trajectory categorizers. Thus, the leading diagonal represents correct responses. With 10 participants each making 30 categorizations distributed across the six categories, each row contains about 50 responses.²

The column totals in table 12-1 (which would equal the row totals if participants matched all of the trajectory categorizations) show that participants overestimated the base rate of play by nearly 70%, and underestimated the amount of evasion by about 30%. (Participants also overestimated courtship and underestimated pursuit and fighting by smaller amounts.) These tendencies could reflect participants' underlying beliefs about the base rates of these motions in nature, particularly that play is more common than serious conflict, especially evasion. The play overesti-

2. Most rows in tables 12-1 and 12-2 do not have exactly 50 entries due to missing data and to the way in which randomly assigned trials were split between the pursuing-evading and courting-courted pairs.

mation could also indicate that participants use play as a default “catch-all” choice when no clear distinction can be made, in keeping with the argument mentioned earlier that play provides training for the other behaviors.

Reading across rows in the table, we also see some more specific sources of confusion. Pursuit was very often miscategorized as play, and fighting was miscategorized as play more often than it was accurately categorized as fighting. Fighting was often mistaken for pursuit and evasion, as well. In addition to the general tendency to mistake other behaviors as play, actual play was often mistaken for other behaviors such as courting, being courted, and fighting. This error can arise when participants first see a chasing component of play, for instance, and make a quick (wrong) decision of pursuit-evasion from this limited time window. In nature, other nontrajectory cues (such as laughing in humans) could be used to categorize play more accurately.

On the other hand, participants did often get the categorizations right: Pursuit and evasion were rarely confused with each other, and courting was rarely confused with being courted. As might be expected, pursuit and being courted were never mistaken for each other, and courting and evading were never mistaken for each other. This is consistent with the existence of similarities between pursuit and courting, and between evasion and being courted. However, important differences are clearly present as well, because pursuit was almost never perceived as courtship, courtship was almost never perceived as pursuit, and being courted was rarely mistaken for evasion. These differences and similarities should be reflected in the set of cues we ultimately settle on for categorizing motion trajectories.

While our data showed a high percentage of systematic errors (overestimating play, underestimating evasion) and many individual confusions, it must be emphasized that even with no feedback, no practice, and minimal instruction, naive participants can categorize behaviors into biologically important classes at far above chance levels, given nothing more than the recorded trajectories of two agents interacting. We have stripped away all environmental context, all gait and posture information about the agents, all facial expression, and all communication. Nevertheless, pure whole-body motion cues are sufficient to categorize the behaviors fairly accurately, from which we conclude that the trajectories we collected do in fact contain enough information to indicate intentional categories. We can therefore proceed to analyze just how that information is reflected in the trajectories: What are the cues we can use to judge intention from motion?

The One-Bug Case: Trajectories Without Relational Cues

One way to determine what sort of cues people (in particular, our participants) use to judge intention from motion is to eliminate part of the available information from the motion trajectories and see how this affects the categorization judgments made on the basis of the modified trajectories. In

the first categorization experiment, participants saw all of the information captured in the complete trajectories of both interacting bugs. This made it possible for participants to use all of the motion cues that rely on the relative position, heading, and speed of the two creatures. To gauge the importance of this class of relational cues, we wanted to find out how well people could do at the categorization task if they were all stripped away. Such relative information can be removed by only playing back the recorded trajectory of one of the two bugs, so that participants cannot know what the other bug is doing.

To this end, we ran a second categorization experiment in the same manner as the first, but with 10 new participants, and with only one of the two bugs in an interacting pair being visible in each presented trajectory. Participants had to categorize what behavior the one visible bug was performing. Table 12-2 shows the confusion matrix from this experiment, presented in the same format as table 12-1. Now, overall percentage correct drops from the 49% level of the two-bug experiment to 30% (still well above the chance level of 17%). This large performance drop indicates that we should include relative motion information in our list of important cues in this domain. We can also learn something interesting from considering the patterns of errors made in this setting.

Inspecting the column totals in table 12-2 reveals that participants in this experiment selected all six categories more evenly; the play overestimation has disappeared, replaced by a slight tendency to overestimate the amount of courting going on. Most of the miscategorizations were due to strong confusions between certain behavior classes. Whereas pursuit was clearly distinct from evasion or courting if the second bug was visible (table 12-1), one-bug pursuit was very often confused with the latter two categories. Evasion was mistaken for pursuit as often as it was accurately categorized, perhaps because both entail high-speed, unpredictable loop-

Table 12-2: Participants' Categorization Confusion Matrix for One-Bug Trajectories

Actual	Choice					
	Pursuing	Evading	Courting	Courted	Fighting	Playing
Pursuing	15	8	12	3	7	7
Evading	13	14	9	1	5	6
Courting	2	5	16	17	5	9
Courted	5	7	4	21	2	5
Fighting	7	4	12	0	13	14
Playing	5	4	10	5	14	10
Totals	47	42	63	47	46	51

Note. Each entry shows how many times participants chose a particular (column) category for a trajectory generated in some actual (row) category, after seeing the motions of only one of the two bugs involved. The main diagonal (in bold) shows correct categorizations.

ing throughout the entire space. Courting was very often mistaken for being courted, but, strangely, being courted was still rarely mistaken for courting. Fighting was mistaken for courting and for play nearly as often as it was categorized accurately.

The differences between tables 12-1 and 12-2 indicate that relational cues are important in making some distinctions but not others. In particular, single-bug (nonrelational) information is usually sufficient to distinguish aggressive intentions (pursuit, evasion, fighting) from more passive intentions such as being courted. But relational cues seem to be necessary to decide whether the bug is the follower (pursuit or courtship) or the avoider (evasion or being courted). These two findings indicate that a different set (and number) of cues is needed to make different categorizations—in some cases, decisions can be made without using all of the available information. Our search for appropriate categorization algorithms later in this chapter will make use of this fact.

Identifying Useful Motion Cues Computable from Trajectory Information

The categorization experiments just described showed that, to categorize intentions from the motion trajectories of two interacting agents, we must consider both individually based cues computed from the motion of a single agent, and relational cues determined by comparing the motions of both agents together. To narrow in on the specific cues that could be used, we must consider the information embedded in the trajectories.

Some obvious cues might relate to the goals of the behavioral categories. For example, in our experiments the goal of pursuit was to intercept the other agent, and the goal of fighting was to strike the other agent's rear end with one's front end. Thus, one could just try to count up the number of successful interceptions and strikes to distinguish pursuit and fighting from play. However, the other agent does not want to be intercepted or struck. Successful evasion implies unsuccessful pursuit. Cues of success may be poor indicators of intention, particularly over short-term portions of an ongoing interaction, and in fact, such cues did not prove to be very diagnostic in our setting.

It is more useful to consider the simple, objective (nonintentional) cues that can be computed given an intentional motion trajectory. The trajectories recorded in our experiment took place in two spatial dimensions, so we can focus on some simple Newtonian cues. A trajectory can be measured in terms of the bug's position and velocity for each of the two linear dimensions. Because the bug figures have a head end and a tail end, they also have a rotational degree of freedom, captured in their orientation or heading, and a rotational velocity equivalent to the rate at which they are changing their orientation. Ignoring higher-order parameters such as acceleration, this yields six basic motion parameters for each agent: hori-

zontal position, horizontal velocity, vertical position, vertical velocity, orientation, and rotational velocity. A two-agent system can thus be described with 12 independent parameters.

This might be fine for a physicist interested in modeling abstract changes in position in a two-body system, but is it useful for categorizing behaviors and intentions of animate agents? Intuition, and the results of our categorization experiments presented in the previous sections, suggest not. Information about one's absolute place in the world and that of some other organism does not matter to individuals nearly as much as comparison between oneself and others—social behavior, like evolution itself, is driven by relative differences. Absolute position information is largely irrelevant in perceiving how two agents are interacting (aside from the effects of special locations in the environment such as watering holes or shelter); what matters is their position relative to each other. Likewise, their individual orientations are not as informative regarding their intentions toward each other as are their relative orientations toward each other. From any one agent's perspective, the three position and orientation parameters that matter are likely to be the distance to the other agent (what we refer to as the *relative distance* cue), the angle between one's current heading and the other agent's location (*relative angle*), and the angle between one's current heading and the other agent's own heading (*relative heading*).

We can also reduce the six velocity parameters that describe a general two-agent system into four velocity parameters that matter to any given agent whose behavior we want to categorize: *absolute velocity* (how fast the agent is going forward, rather than in the two orthogonal spatial directions separately—as we will see, we can ignore the other velocity component corresponding to sideways motion), *relative velocity* (how fast the one agent is going relative to the other agent), *absolute vorticity* (how fast the agent is changing heading), and *relative vorticity* (how fast the one agent is changing heading relative to the other agent's heading change). We can eliminate two parameters by ignoring one's own sideways velocity (which is usually zero in nature and in our bug world), and relative sideways velocity (which is also usually zero). Thus, we are left with three relevant position parameters and four relevant velocity parameters. These can be converted into motion cues by averaging them over some temporal window. In accordance with the results of our second categorization experiment, which demonstrated the importance of relational information, five of the seven cues we have ended up with are relational (uncomputable given just one visible agent).

In summary, we propose seven simple, ecologically relevant cues that can be useful in categorizing the intentions of one (focal) agent interacting with another agent.

Relative distance: the distance between the two agents. Fighting and pursuit-evasion tend to produce smaller relative distances in the bug trajectories than do courtship and play.

Relative angle: the angle between the focal agent's current heading (the direction it is facing) and the other agent. This cue indicates whether the other agent is in front of or behind the focal agent. Pursuit can be distinguished from evasion most easily by relative angle.

Relative heading: the difference in heading between one agent and the other, indicating whether they are facing the same direction (more typical of pursuit and evasion), or facing each other (more typical of fighting, courtship, and play).

Absolute velocity: the forward velocity of the focal agent with respect to the background environment. Pursuit-evasion and fighting generate higher absolute velocities than do courtship and play.

Relative velocity: the difference between the velocities of the two agents. Pursuers and evaders tend to have nearly zero relative velocity across the duration of a chase, but courtship produces large differences in velocity between courter and courtee.

Absolute vorticity: the vorticity (change in heading) of the agent with respect to the background environment. Fighting and courting produce high vorticities, while play and being courted are associated with low vorticity.

Relative vorticity: the difference between the vorticities of the two agents. Pursuers and evaders tend to turn equally often, and so have zero relative vorticity, while courtiers and courtees often have a large vorticity difference.

Before settling on this set of seven cues distilled from the motion trajectories, we need to be sure that they contain at least enough information to make reasonably accurate intention categorizations. If we discover otherwise, then we would have to search for a different set of more appropriate cues. One way to test this question is to see whether a general-purpose pattern-associating algorithm can map from the particular values of these seven cues computed from portions of motion trajectories to the correct categories for those trajectories. We trained a simple logistic-activation three-layer neural network on the same 300 examples of motion trajectories that participants saw in our categorization experiments described earlier, using the values of the seven cues as inputs and the proper category as the target output. After training, the neural network correctly categorized 247 out of the 300 examples, or 82%. (When tested for generalization ability on 300 different examples, the network still got 200 correct, or 67%—see table 12-3.) Recall that participants only correctly categorized 49% of the 300 examples. Based on this, we concluded that these seven cues, while perhaps not fully capturing all of the distinguishing trajectory information, still distill enough information to allow us to create reasonable models of human performance in this task.

Solitary animals that only encounter a single other conspecific individual at any one time need only be able to categorize the behavior of those others relative to themselves, that is, from an egocentric viewpoint. For animals in more social species, where interactions between others are commonly witnessed, it can also be important to keep track of the inten-

Table 12-3: Categorization Accuracy of Participants and Competing Algorithms for One- and Two-Bug Trajectories, Along With Number of Cues Used by Each

Criterion	Participants	Competitor			
		CBE	Dawes's Rule	Franklin's Rule	Neural Network
% correct on training set with 7 cues	49%	65%	62%	68%	82%
% matches to participants' correct	—	73%	72%	77%	90%
Average number of cues used	?	3.6	7	7	7
% correct on testing set with 2 cues	30%	34%	33%	34%	26%
% correct on testing set with 7 cues	—	57%	60%	60%	67%

Note. Here, accuracy is shown in percentage correct (out of 300 trials), except for the second row, which shows percentage of matches to participants' (147) correct categorizations.

tions of different group members relative to each other. This is particularly true in species with dominance hierarchies and kin networks. Furthermore, between-species interactions often need to be judged accurately. For example, a gazelle should be able to tell if a lioness is chasing one of its herd mates, or if she is being chased herself, perhaps by another lion. Each of the seven cues listed above should be readily computable from either an egocentric viewpoint or from a third-party observer position. Although our experiments presented trajectories from a rather ecologically implausible (but computationally simple) top-down view, the seven cues could be computed almost as easily from a more realistic ground-level view.

The seven cues just described, like the 12 Newtonian parameters, are mostly independent in principle, but rather highly correlated in practice. That is, the structure of motion trajectories produced by organisms in nature will ensure that certain cue values co-occur. These natural intercorrelations between motion cues have two implications for categorizing intentions. First, only a few cues may suffice for making an accurate decision, because there is high overlap in information content between cues. Second, if one cue is unavailable, another may take its place: Correlated cues enable *vicarious functioning* (essentially, cue substitutability) in decision making (Brunswick, 1943).

Vicarious functioning is important because in many real-world situations some or most of the possible cues may not be perceivable at any given time. Furthermore, the speed with which a cue can be registered may be another crucial constraint on the decision-making process. For high-pressure problems, such as deciding whether a mountain lion is stalk-

ing you, it may not suffice to wait for all cues to become available before making a decision. Such time pressure may favor fast and frugal, noncompensatory decision heuristics that make the most efficient use of the motion information available (see chapter 7). These are the sorts of decision heuristics for categorizing behaviors from motion cues that we will now consider.

Decision Heuristics for Categorizing Behaviors

What sort of decision mechanism might people and other animals use to process motion cues into intention category judgments? As we have indicated, judging intention from motion is a critical task often performed under time pressure. Moreover, our categorization experiments indicated that different cues are useful for deciding on different categories. These observations led us to look for a simple fast and frugal decision mechanism that uses only as much of the information as is necessary to select a single category. The Categorization by Elimination (CBE) algorithm is just such a mechanism (see chapter 11 for a full introduction). CBE categorizes a given stimulus by starting with a full set of possible categories and then using the particular values of the cues for that stimulus to eliminate more and more categories until only a single possibility remains. Thus, different stimuli may require that more or fewer cues are processed before they can be categorized. This algorithm always checks the cues in a particular predetermined order, and only uses as few cues as it can get away with to reach category decision, rendering it fast and frugal. Yet its categorization accuracy on some standard test sets is still close to that of more traditional algorithms that combine all available cues (as shown in chapter 11).

To use CBE to categorize motion intentions, we first determined the order in which to use our seven cues, based on how well they can each distinguish among the six intentional categories in our training set of 300 trajectory examples. We found that the absolute velocity cue was the most accurate at categorizing these trajectories when used alone, so it was put first in the list, followed by relative angle, relative velocity, relative heading, relative vorticity, absolute vorticity, and relative distance. (We found that the exact cue ordering did not make much difference in overall categorization performance; however, it does have a large effect on the number of cues looked up by the algorithm.) Then, to categorize a particular trajectory, CBE starts with the full set of six possible intentional categories and considers the value of the absolute velocity cue for that trajectory to determine which categories to eliminate from the set of remaining possibilities. For instance, a relatively low velocity value of only 400 pixels per second indicates that the trajectory could correspond to the more leisurely categories of courting, being courted, or play, so that the other three more frenetic categories should be eliminated from the set of possibilities. Because more than one category remains after checking velocity, the next cue in

the preordered list (relative angle) must be used to eliminate more categories. This process of checking further cues continues until a single category is left, which is the algorithm's final decision.

Categorization by Elimination has several interesting features as an algorithm for cue-based categorization. It is nonintegrative, using only those cues necessary to reach a decision. If each cue takes some time to assess, this makes CBE faster than algorithms that use all cues. It is also noncompensatory, with earlier cues eliminating category possibilities that can never be restored by later cues. Whether this is a reasonable feature depends on the type and number of errors CBE produces, and the consequences of those errors. For example, if CBE eliminates the category "fight" too readily, and stupidly mistakes fighting for playing, CBE would be maladaptive. Finally, CBE always uses cues in a particular order. Here we order the cues by their ability to make correct categorizations (cue success—see chapter 11). It is important to use cues in a good order, because this is what allows CBE to make a rapid decision. If the algorithm looked at cues starting with the *least* useful, for instance, it would take more cues (and hence more time) to reach a decision and would more often be incorrect.

We tested the performance of CBE on categorizing trajectory intentions using the seven motion cues against both the correct (originally intended) categories and the categories that participants actually chose (from table 12-1). These comparisons are shown in table 12-3. To see how this fast and frugal heuristic performs in comparison with more traditional decision-making methods, we also constructed two linear combination models (Dawes's rule, using unit weights, and Franklin's rule, using cue success as weights—see chapter 4), which process the cues in the usual integrative, compensatory fashion. As mentioned in the previous section, we also trained a three-layer neural network model on this task to see how a nonlinear, compensatory system would do.

The first row of table 12-3 shows the performance of the above four algorithms on 300 categorizations, along with participant performance on the same data set. (The cue values in each of the 300 trials were computed from the same time period of trajectory data that the participants saw in that particular trial, ensuring that no extra advantage was given to the algorithms over the participants.) Each of the algorithms outperformed the participants by a fair margin. But this is not surprising given that the algorithms were all trained with feedback on the data to make the proper categorizations, while participants were not given feedback and had to categorize each motion pattern the first time they saw it. What is more surprising, though, is that the fast and frugal CBE performs about as well as the more traditional linear information-integrating methods (see chapter 11 for more such surprises). Furthermore, CBE uses only half of the cues, on average, that the other algorithms do—and thus it achieves its good categorization performance with significantly less data.

How well do CBE and the other algorithms match the decisions made

by our participants? For all trials that participants got right (147), we compare how many times each algorithm also chose the (same) correct category and show this as a percentage of the 147 trials in row 2 of table 12-3. (The different algorithms' matches to participant mistakes appeared largely random and uninformative—there are more ways to be wrong than to be right in this task—and so are not included here.) Here again CBE does as well as the other linear-integrative algorithms, matching nearly three-quarters of the participants' correct answers. Given that CBE is also faster and more frugal than the other two linear models, we have some reason for preferring it as a psychological model worthy of further investigation. (In comparison, the baseline neural network algorithm matched 90% of the participants' correct responses.)

But just how much more frugal is CBE in this task? In the third row of table 12-3, we show how many cues the different algorithms used to make each categorization, averaged across the 300 trajectories seen. For the linear models and the neural network, this average is rather straightforward: These algorithms always use all of the available cues, so their average cue use is 7. But CBE can use a different number of cues to categorize each trajectory, and in general uses as few as possible. On average, CBE uses only 3.6 cues per categorization, little more than half of the total number of cues available.

The benefits of CBE's minimal cue usage are indicated in the fourth row of table 12-3, where we have strictly limited the number of cues that participants or algorithms can use to make categorizations. Specifically, we showed participants only one of the interacting bugs in the trajectory (this is the data from the second categorization experiment reported earlier), and we allowed the algorithms to use only the two nonrelational cues that are computable from such a restricted single-bug trajectory (absolute velocity and absolute vorticity). CBE and the two linear algorithms again do at least as well as the human participants. But the powerful neural network model, which did so well on the full training set, is now relatively crippled by the lack of information it was expecting—it is the only decision model that makes fewer correct categorizations than the participants do. In the last row of table 12-3 we see further evidence about the ability of these models to adapt to situations with new information: Here we show their performance in generalizing to a different set of 300 trajectories. All lose a fair amount of categorization accuracy, with the neural network losing the most and the linear algorithms the least; CBE lies in the middle.

These results can help us put the superior performance of the neural network model on the full-cue categorization task in perspective. Certainly the greater number of free parameters in the neural network weights can help its fitting performance. Beyond this, perhaps nonlinear cue integration has an advantage in this domain, and perhaps compensatory cue use is more appropriate here as well, when all the cues are known. However, compensatory cue use requires simultaneous access to all relevant

cues, and there may be situations where cues are naturally perceived at different times or in different orders. In such cases, standard neural network models may be at a relative disadvantage compared with fast and frugal competitors. CBE, by contrast, could use cues in the natural order in which they can be perceived, categorizing as far as possible given each piece of incoming information. For example, relative position information may often be easier to estimate than average velocity or vorticity information, and so could be used to make an initial estimate of the appropriate intentional category in a fast and frugal manner. These questions should be explored by creating on-line dynamic decision models that can appropriately categorize animate intentions under ecologically realistic time constraints.

In sum, the fast and frugal CBE heuristic performs about as accurately as more traditional linear cue-integrating mechanisms as a model of human categorization in the intention-from-motion domain.³ It achieves this surprising performance despite using on average only half of the information that the other algorithms use. This frugality allows CBE to make faster decisions, especially in realistic situations where time is required to assess the value of each cue. Using fewer cues can also lend CBE added robustness in comparison to the plethora of parameters in the competing neural network model—several cues could be unobservable in the environment without affecting CBE's performance at all. These advantages may not make much of a difference in a laboratory setting, but they can be crucial in real-world life-and-death decisions. This leads us to propose CBE as an ecologically rational algorithm for determining an organism's intention from its motion alone.

A Motion Turing Test and Other Extensions

As an adjunct to our categorization work, we devised a pilot version of a motion-based Turing test, in which a robot bug replaces one of the human participants in the trajectory-categorization experiments described earlier. After the trial, we asked the lone human participants whether they thought they were interacting with another human-controlled bug, or with a computer-controlled bug. A pilot study suggested that even the simplest robot algorithms could be surprisingly convincing, especially when their actions respond to the human-controlled bug (as opposed to acting completely independently). That is, if one bug (e.g., a pursuer) is really controlled by a human and the other (e.g., an evader) is controlled by a simple

3. As is often the case, it is difficult to make an empirical distinction between the fast and frugal algorithm and other approaches based solely on comparisons to human choices. See chapter 7 for more on this difficulty, and for some approaches to distinguishing between algorithms using process data such as reaction times in addition to choice data.

computer program, the resulting motion pattern looks almost indistinguishable from that of two human-controlled bugs interacting. This simple test further indicates the power of limited motion cues to indicate animate intention. We can similarly test other proposed cues of intention from motion by building behaviors that generate those cues into new robot bugs and investigating which combinations are most convincing to human observers.

These sorts of animate motion experiments may provide some useful new methods for investigating human theory of mind—the attribution of intentions, beliefs, and desires to others based on observable behaviors. Some researchers suggest that autistic people have deficits in their theory of mind module, such that their attributions of intention are severely impaired (Baron-Cohen, 1995). If so, it would be interesting to see how autistics do on intention-from-motion categorization tasks: The types of animate intentions that different individuals can reliably categorize from motion trajectories may indicate the presence or absence of different kinds of empathic or social-attribution impairments.

Intention-from-motion heuristics could also be used in a variety of practical applications. Many countries are increasing their use of closed-circuit cameras to detect crime. Such crime detection is basically a problem of distinguishing criminal intentions and behavior patterns from benign patterns, given motion cues. Even our simple cue-based methods of distinguishing between pursuit, evasion, fighting, and play may find uses in such systems. Automated crime-detection systems would not have to be perfect on their own; they would only have to help security guards identify which screen to pay most attention to out of the many screens they are expected to monitor. Also, pharmaceutical companies often test drugs by recording their effects on animal behavior patterns. Algorithmic systems for categorizing rat behavior as aggressive, exploratory, or playful may help in automating such evaluations. The problem of inferring intentions from motion cues is so general that there are doubtless hundreds of other related applications of simple fast and frugal decision heuristics.

Conclusions

In this chapter we have succeeded in uncovering a set of motion cues that can be used to infer some major categories of adaptively important intentions. The set of seven cues we proposed appears sufficient to capture the major regularities of motion in the six intentional categories we investigated. We introduced a new simple algorithm, Categorization by Elimination, that could accurately categorize motion patterns using a minimal number of cues, showing that these adaptively important inferences can be made in a fast and frugal manner.

Our main goal has been to provide a basis for the concept of social rationality, by finding simple motion cues and fast and frugal cue integration mechanisms that humans and other animals may use to interact effec-

tively with each other. Such mechanisms can exploit the fact that animate motion tends to fall into a few rather stereotyped categories that can be derived from basic evolutionary and ecological principles. We tested the competing decision algorithms on motion trajectories generated by human subjects in the course of playing various computer-mediated interaction games. We propose this methodological approach as a first step in examining how the basic building blocks of social cognition can be studied through a combination of evolutionary principles, ecologically representative stimuli, human experimentation, and computer simulation of how well different decision heuristics would perform against each other and compared with human data. Our guiding principle has been that natural environments offer a few key motion cues, a few typical kinds of animate motion, and a plethora of animals, including mountain lions, more than willing to eliminate others who fail to perceive the significance of their movements.

VOLUN AND COGNITION

eneraltor: Stephen Stich, Rutgers University

ublishn the Series

implauristics That Make Us Smart

erd Genzer, Peter M. Todd, and the ABC Research Group

Simple Heuristics That Make Us Smart

Gerd Gigerenzer

Peter M. Todd

and the ABC Research Group

New York Oxford

Oxford University Press

1999