Journal of Experimental Psychology:
Human Perception and Performance
2011, Vol. 37, No. 3, 669–684

0096-1523/11/$12.00 DOI: 10.1037/a0020735

Chasing vs. Stalking: Interrupting the Perception of Animacy

Tao Gao and Brian J. Scholl
Yale University

Visual experience involves not only physical features such as color and shape, but also higher-level properties such as animacy and goal-directedness. Perceiving animacy is an inherently dynamic experience, in part because agents' goal-directed behavior may be frequently in flux—unlike many of their physical properties. How does the visual system maintain and update representations of agents' animate and goal-directed behavior over time and motion? The present study explored this question in the context of a particularly salient form of perceived animacy: *chasing*, in which one shape (the "wolf") pursues another shape (the "sheep"). Here the participants themselves controlled the movement of the sheep, and the perception of chasing was assessed in terms of their ability to avoid being caught by the wolf—which looked identical to many moving distractors, and so could be identified only by its motion. The wolf's pursuit was frequently interrupted by periods in which it was static, jiggling in place, or moving randomly (amidst distractors that behaved similarly). Only the latter condition greatly impaired the detection of chasing—and only when the random motion was grouped into temporally extended chunks. These results reveal (1) how the detection of chasing is determined by the character and temporal grouping (rather than just the brute amount) of "pursuit" over time; and (2) how these temporal dynamics can lead the visual system to either construct or actively reject interpretations of chasing.

Keywords: event perception, perception of animacy, intention, chasing, social perception

What do we see? One common answer to this basic question is that we see objects and scenes that are composed of many different visual features, including colors, shapes, orientations, textures, et al. Beyond these basic visual properties, though, our percepts can also involve seemingly higher-level features: we can also see properties such as causality and animacy. The ability to see a property such as animacy (and related properties of agency, intentionality, and goal-directedness) is clear when we view other people or animals, but perhaps the most fascinating aspect of this ability is that we can also readily perceive animacy in displays that contain only simple moving geometric shapes. We may perceive such shapes as alive, as having goals, and as interacting in social relationships such as chasing or fighting. When animated, for example, the motion patterns of the shapes in Figure 1 may lead observers to perceive that the large triangle is trying to catch the circle and small triangle, which are trying to flee (Heider & Simmel, 1944).

Such phenomena have attracted the interest of a diverse array of cognitive scientists (from psychologists and neuroscientists to anthropologists and computer scientists), and several lines of evidence suggest that this ability may reflect a primitive sort of visual processing rather than only higher-level cognitive interpretations. In the first place, percepts of animacy—like percepts of other visual features—are tightly controlled by subtle aspects of the displays themselves, but can be surprisingly resistant to explicit beliefs and intentions (for a review, see Scholl & Tremoulet, 2000). They also appear to be robustly expressed in many different populations—appearing cross-culturally (e.g. Barrett, Todd, Miller, & Blythe, 2005), and in young infants (e.g. Csibra, 2008; Gergely, Nádasdy, Csibra, & Bíró, 1995; Rochat, Morgan, & Carpenter, 1997; Rochat, Striano, & Morgan, 2004). At the same time, however, the reflexive ability to perceive animacy can be disrupted by particular sorts of neuropsychological impairments, including autism spectrum disorder (e.g. Abell, Happé, & Frith, 2000; Klin, 2000; Rutherford, Pennington, & Rogers, 2006) and conditions such as amygdala damage (e.g. Heberlein & Adolphs, 2004).

Perceiving Animacy and Goal-Directedness

Research in this domain is variously referred to in terms of the perception of animacy (e.g. Gelman, Durgin, & Kaufman, 1995; Tremoulet & Feldman, 2000), intentionality (e.g. Dasser, Ulbaek, & Premack, 1989; Dittrich & Lea, 1994), goal-directedness (e.g. Csibra, 2008; Opfer, 2002), social causality (e.g. Rochat et al., 1997, 2004), or social meaning (Tavares, Lawrence, & Barnard, 2008). Sometimes these distinctions are important, as when researchers wish to attribute one property but not another (e.g. goal-directedness without other aspects of mental-state reasoning; e.g. Gergely & Csibra, 2003). However, while there are some clear differences between these constructs, it remains unclear precisely how each of these notions is to be rigorously defined.

Tao Gao and Brian J. Scholl, Department of Psychology, Yale University.

For helpful conversation and/or comments on earlier drafts, we thank Andrea Heberlein, Greg McCarthy, Mowei Shen, Jinjin Sun, and Duje Tadin. We are also indebted to Xiang Huang and Jifan Zhou for assistance with data collection.

Correspondence concerning this article should be addressed to Tao Gao or Brian J. Scholl, Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205. E-mail: tao.gao@yale.edu or brian.scholl@yale.edu

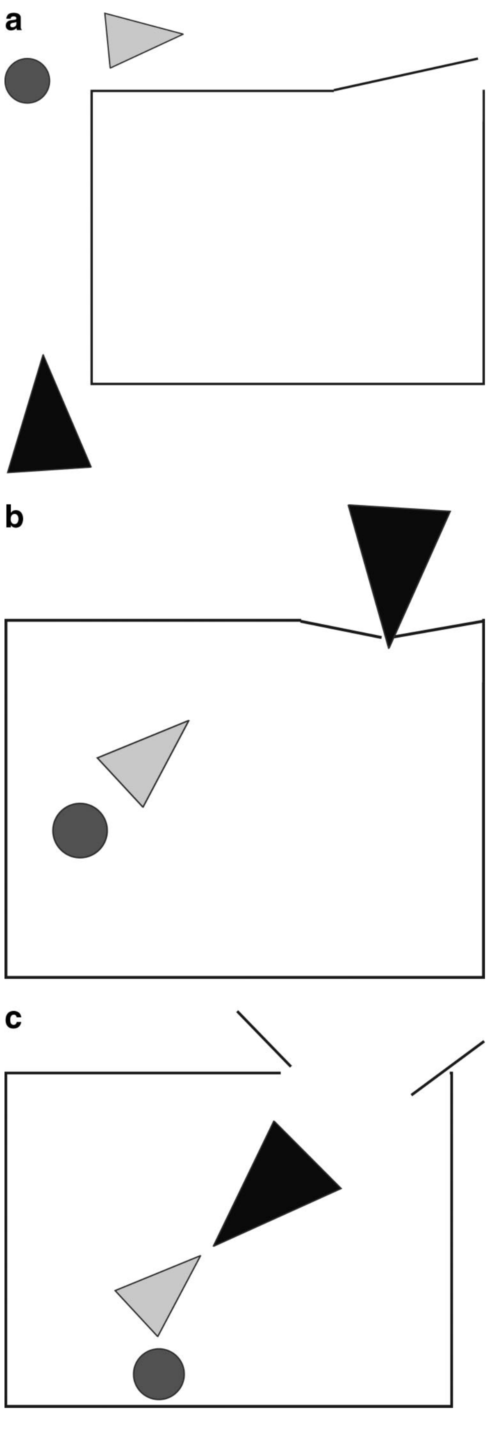

Figure 1. Sample "snapshots" of stimuli modeled on Heider and Simmel's (1944) displays, with simple shapes that have been used evoke the perception of animacy. (a) A static frame in which the large triangle is perceived as chasing the two smaller shapes around the box. (b) A static frame in which the large triangle is perceived as attempting to break into the box in which the two smaller shapes are hiding. (c) A static frame in which the large triangle is perceived as fighting with the smaller triangle, which is protecting the disc.

One of the difficulties surrounding such definitional issues is that there are so many different cues to animacy (and to goal-directedness, et al.), and each of these cues may operate to some degree independently, and may differentially activate these various theoretical notions. Researchers have sometimes seemed to assume that there is a single monolithic "animacy detector," but in fact we think that there may often be many such cues which are often mixed together in common displays. (Certainly this was true in the original Heider and Simmel studies.) We suspect, however, that these individual cues must be studied in isolation in order to start clarifying which of them trigger which important categories. We think of this by analogy to depth perception, where several very different cues operate largely independently, and are combined only later in visual processing. Thus, in order to figure out how any one cue actually operates, it must be studied in relative isolation (e.g. disparity without motion parallax, and vice versa).

The present paper reports a case study of one potentially powerful cue to animacy—chasing—and the two most relevant theoretical notions for this project appear to be animacy and goal-directedness. Note that the latter appears to be a subset of the former: you can perceive animacy without goal-directedness (it seems alive, but aimless), but not the reverse. (Of course, that may be true only with regard to the particular spatiotemporal cues that constitute chasing.) As such, we will consider our displays throughout this paper in the context of the perception of *animacy*, since this seems to us a more conservative route. In particular, though the type of animacy implicated by our chasing displays clearly involves the perception of goal-directedness, that may be only a particular type of goal-directedness, distinct from those implicated by other cues (such as perceived eye gaze).

Dynamic Goals vs. Static Physical Properties

On their surface, mental states (such as intentions or goals) and physical properties (such as colors or textures) seem like categorically different domains. The perception of properties in each category, however, may be similar in several important ways. In the first place, the perception of both sorts of properties seems to be relatively efficient, automatic and even irresistible. Just as we perceive colors without instruction and without exerting effort to do so, so too we may effortlessly see animacy in simple displays without being instructed to develop such interpretations. Moreover, just as higher-level knowledge about colors cannot readily influence what colors we *see*—e.g. in the resistance of color illusions to our knowledge of how they were constructed—so too does the perception of animacy seem to be relatively encapsulated from our beliefs and intentions. This is clear to some degree whenever we view such displays: when viewing the Heider and Simmel animation, for example, we *know* with a high degree of certainty that the displays are in fact completely lifeless and void of any intentions or goals—but this will not stop us from reflexively *seeing* the shapes as alive when they move in certain ways. In short, the perception of animacy seems more akin to a reflex than to a deliberate act.

A second way in which the perception of animacy and physical properties may be similar is in terms of their underlying "logic." In both cases, the light entering the eyes is insufficient to deductively infer the actual state of the world, and so the visual system must proceed by making "unconscious inferences" about the source of the visual stimulation (e.g. Gregory, 1980; Rock, 1983). The form of these inferences may often be heuristic assumptions that certain properties of the visual image are reliable cues to the nature of the external world, even though such connections are not strictly necessary. In the physical domain, for example, the visual system may assume that collinear line segments in the image are collinear in the world (e.g. Marr, 1982) or that fuzzy luminance boundaries reflect changes in illumination (e.g. due to shadows) rather than reflectance (e.g. Adelson, 2000). In the social domain, the visual system may similarly assume that certain spatiotemporal patterns in an image are reliable cues to the existence and nature of animacy and goal-directed action. Specific cues that have been proposed to mediate the detection of animacy include sudden direction and speed changes (Tremoulet & Feldman, 2000), synchronized motions (Bassili, 1976; Dasser et al., 1989), rational interactions with spatial contexts (Tremoulet & Feldman, 2006), and other objects (Dittrich & Lea, 1994; Santos, David, Bente, & Vogeley, 2008), apparent violations of Newtonian mechanics (Gelman, Durgin, & Kaufman, 1995), and multiple objects' coordinated orientations (Gao, McCarthy, & Scholl, 2010). These are all merely heuristic assumptions, of course, in that inanimate displays can also have such properties (as in fact they do in most experiments in this domain!), but they can nevertheless yield very real percepts of animacy—just as we may have very real but incorrect percepts of physical properties such as color, and via a similar type of underlying logic.

Nevertheless, these two domains are importantly different in other ways, and these differences may have implications for the nature of the underlying perceptual processing of such properties. In this paper we focus on one such difference, involving the *persistence* of various properties. Many (but not all) physical properties are relatively unchanging, at least on a short temporal scale: from moment to moment, for example, the shape of a cheetah's head, or the color of its coat, tends to remain fixed. In contrast, a cheetah's mental states—and its degree of goal-directedness—may be frequently in flux. During a hunt, for example, the cheetah may be chasing one gazelle *now*, a different gazelle *now*, etc. This is also true for humans, of course: we can change our minds much more readily than we can change many aspects of our appearances (say, during a game of "tag"). This contrast is also well illustrated in displays such as those pioneered by Heider and Simmel. In the display depicted in Figure 1, for example, the shape and color of the larger object stay constant, but its local goals are in flux (now chasing the green disc, now breaking into the rectangle, etc.).

This dynamic character of mental states raises a special challenge for the perception of animacy: how does the visual system maintain and update representations of agents' goal-directed behavior over time and motion? In the current project, we address this challenge by investigating the temporal dynamics of a particularly salient form of animacy: *chasing*, in which one object (the "wolf") chases another object (the "sheep").

The Psychophysics of Chasing

We explore chasing here as a case study of a particular type of perceived animacy. The almost-cinematic "stories" of Heider and Simmel's (1944) iconic displays are especially rich and compelling, but they come with a cost: these displays are typically constructed haphazardly, and they may mix together many separate cues to animacy. Our strategy, in contrast, is to study the perception of animacy in simpler displays that are algorithmically generated, and that vary only a few well-controlled variables at a time—which may in turn aid our understanding of the nature of the underlying processing. Chasing, in general, is our candidate for what is perhaps the simplest type of behavior that is nevertheless perceived in richly animate terms (Gao, Newman, & Scholl, 2009). Chasing is one small part of many displays used in studies of perceived animacy, including those of Heider and Simmel (see also Barrett et al., 2005; Blythe, Todd, & Miller, 1999; Dittrich & Lea, 1994), and previous research has shown that even infants are able to perceive the special character of such displays (Rochat et al., 1997, 2004). Moreover, chasing has clear ecological importance to our lives—certainly during our evolutionary history, and perhaps even today. (Children chase each other on playgrounds. Cops chase robbers. And researchers report that one of the most commonly remembered dreams even in today's world is that one is being chased; e.g. Garfield, 2001; Revonsuo, 2000).

Recent studies of perceived chasing have involved displays of simple geometric shapes that move about a computer display—with one shape (the "wolf") pursuing another (the "sheep"; Dittrich & Lea, 1994). This research (Gao et al., 2009) has revealed an interesting dissociation between *objective* chasing (the degree to which the wolf actually pursues the sheep) and *perceived* chasing (the degree to which we can detect this behavior). When the wolf pursues the sheep in a perfectly "heat-seeking" manner—being displaced on every frame of motion in the exact direction of the (also-moving) sheep—then the wolf/sheep interaction is readily perceived as chasing. However, when the wolf's motion only mildly deviates from perfect heat-seeking (e.g. moving in the general direction of the sheep, but not directly toward it), then chasing is much more difficult to detect—even when the wolf is still quite efficient at moving closer and closer to the sheep over time (Gao et al., 2009). This phenomenon of perceived chasing was shown not to simply reflect the perception of correlated motions, since the effects disappeared in a "Phantom Wolf" condition, wherein the wolf and sheep were just as correlated, but the wolf instead chased the sheep's mirrored reflection in the display (Gao et al., 2009, Experiment 2). Similarly, this work showed that the perception of chasing in such displays does not result simply from closer overall wolf/sheep proximity (in chasing displays relative to non-chasing displays), since the effects were greatly attenuated in conditions when the identical animations were simply played in reverse (Gao et al., 2009, Experiment 3).

When these simple shapes have orientations (e.g. being drawn as "darts" rather than discs), then the perception of chasing is automatically impacted by the direction in which the wolf is oriented during its motion. When the wolf is "facing" in the same direction in which it is moving, then the detection of chasing is facilitated—but the detection of chasing is impaired when the wolf's orientation is random with respect to its actual trajectory (Gao et al., 2009), or when many shapes are consistently oriented to face a single object (Gao et al., 2010). These studies collectively provide an initial exploration of the "psychophysics of chasing," a project that we continue here in a new context.

The Current Study: Interrupted Chasing

In the current study we explore the temporal dynamics of perceived chasing, as a case study of how the perception of animacy copes with dynamic changes. We do so by introducing a novel manipulation: temporally *interrupted* chasing. Rather than chasing the sheep continually during each trial, the wolf's chasing is periodically interrupted by different types of non-chasing motion. In particular, each trial is split into a number of temporal intervals, and each interval contains both a sub-interval of Chasing and a sub-interval of Non-Chasing (Figure 2). The key variable is then the relative duration of each sub-interval, such that each interval can contain anywhere from 0% to 100% of chasing, in 20% steps. During Chasing sub-intervals, the wolf is always displaced on each frame of motion directly toward the (also-moving) sheep. During Non-Chasing sub-intervals, the wolf either moves randomly (in Experiment 1), remains static (in Experiment 2), or jiggles in place (Experiment 3). Experiment 4 then involves all three manipulations in a single within-subjects design; Experiment 5 manipulates the overall interval duration; and Experiment 6 varies spatial deviations from chasing during Non-Chasing sub-intervals.

These types of "interrupted chasing" enable us to explore how an objects' historical and current motion information can be integrated to maintain and/or update a percept involving its goal-directed motions. In particular, we test two possible conceptual and qualitative models of the temporal dynamics of perceived chasing. According the *Accumulating Positive Evidence* model, the visual system detects evidence for chasing, leading to a perceived chase if and when that evidence reaches some threshold. In our experimental design, this would involve accumulating evidence during Chasing sub-intervals (the signal to be detected), while essentially ignoring the interrupting Non-Chasing sub-intervals (which serve as noise). According to the *Detecting Negative Evidence* model, in contrast, the visual system may not only accumulate evidence *for* chasing, but may also be sensitive to evidence *against* an interpretation of chasing. In our experimental design, this would involve the Non-Chasing sub-intervals actively disconfirming a chasing interpretation (perhaps despite prior evidence in favor of chasing), beyond serving as interrupting noise. In other words, this study allows us to explore not only some of the cues which trigger the perception of chasing, but also the possibility that there are cues which actively *block* this sort of perceived animacy.

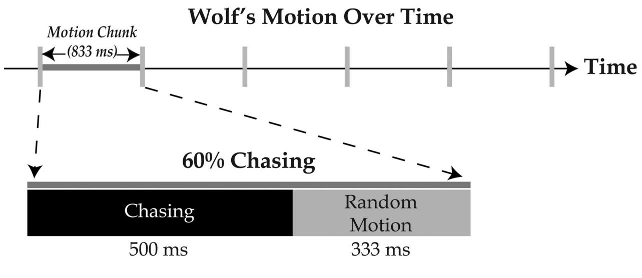

Figure 2. An illustration of the Interrupted Chasing manipulation used in Experiment 1. In each trial, the wolf's motion is split into a number of temporal intervals, containing both a sub-interval of Chasing followed by a sub-interval of Random Motion. The duration of each repeated interval is 833 ms. The percentage of each interval devoted to chasing was varied. In the 60% Chasing Percentage example depicted here, each interval consists of a 500 ms Chasing sub-interval followed by a 333 ms Random Motion sub-interval.

(Of course, the "intervals" in this study existed purely as a stimulus manipulation, and neither of these conceptual models would be sensitive to the presence of such intervals, as such. Rather, both of these qualitative models would simply accumulate evidence in favor of chasing—or the lack thereof—based on individual tokens of stimulus behavior. Thus the key difference between these models is not when they integrate evidence, but rather what type of evidence is being integrated—i.e. only approaching motions [in the Accumulating Positive Evidence conceptual model] or also random and receding motions [in the Detecting Negative Evidence conceptual model]).

General Method: The Don't-Get-Caught Task

Each of the experiments reported here employs the *Don't-Get-Caught* task, originally introduced in the context of spatial chasing deviations (Gao et al., 2009). In this task, participants control the sheep's movement themselves during each trial by moving the computer mouse. The explicit task is to move the sheep about the display in order to avoid getting "caught" (i.e. touched) by the wolf. The difficulty with this task, though, is that each display contains many different objects (see Figure 3), such that the participant cannot avoid them all. While the sheep itself is always highlighted (by being drawn in a different color), the wolf looks identical to the many other distractors. As a result, participants must detect the wolf's presence and position purely on the basis of its spatiotemporal behavior. This is thus a chasing detection task, since participants cannot escape from the wolf until they first detect its presence and position.[1] (Once they do detect the wolf, in contrast, their relative speeds make it easy to avoid being caught.) Each trial ends when either the user-controlled sheep get caught by the wolf (Failure!) or after a certain period of time has passed (Success!), and we measure performance in terms of the percentage of trials on which participants avoid being caught. (These displays are inherently dynamic and are difficult to depict in static figures. Demonstrations of the conditions used in all experiments are available online at http://www.yale.edu/perception/Interrupted-Chasing)

Note that several characteristics of the displays used in this task make it impossible in practice for participants to succeed without actually detecting the wolf. First, the relatively dense display (i.e. with many distractors) makes it impossible to simply avoid *all* of the objects in the display; rather, participants must choose which to avoid. Second, one might worry that since the initial wolf-sheep distance was always considerable, a useful "shortcut" strategy would simply be to move the sheep around to stay in the vicinity of the initial distractors (none of which could be the wolf). This strategy is not possible in practice, however, simply because that

---

[1] Sometimes "false positives" do occur, in that participants mistakenly perceive a distractor as a wolf for a short period based on its coincidental motions. However, this would happen no more frequently in any condition, so that it does not influence our results. For evidence that performance on the Don't-Get-Caught task cannot be explained by appeal only to the detection of correlation or proximity, see Gao et al. (2009).

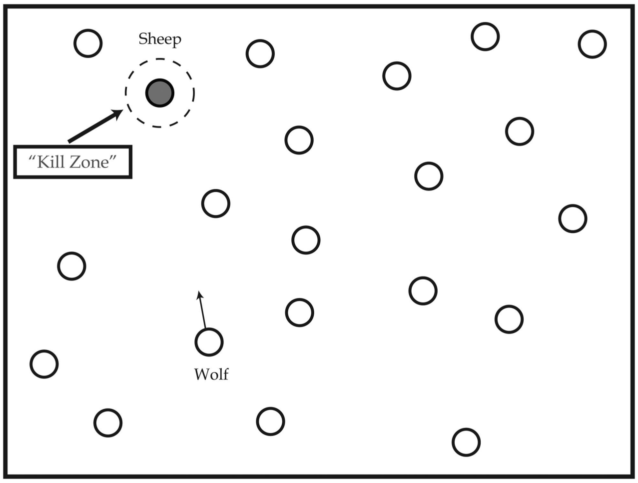

Figure 3. A screenshot from the dynamic display in the Don't-Get-Caught task from Experiment 1. The participant must use the mouse to move the sheep around the crowded display so that the moving wolf that is chasing it never enters its "kill zone." (For online demonstrations of each condition, see http://www.yale.edu/perception/Interrupted-Chasing/.)

"vicinity" does not exist: because of the relatively fast and independent motions, the objects are all shifting about the display independently at all times. Instead, the only way to reliably succeed at this task is to first detect the wolf.

We adopted the Don't-Get-Caught task in the current study for several reasons. First, unlike subjective verbal reports or ratings of perceived animacy (as used in most previous studies; see Gao et al., 2009, for discussion), this provides a measure of objective visual performance. As such, the results of these experiments are less likely to be contaminated by higher-level cognitive inferences or beliefs, since we are effectively measuring the *ability* to detect chasing. Second, the very nature of this paradigm serves as support for the ecological utility of this ability: rather than studying only a *percept*, we are also able to determine how this percept impacts actual behavior. (And, as such, the results of these experiments can tell us not only "how to chase" effectively, but how to chase—or "stalk"—in a way that will be difficult to detect!) Finally, our previous study that introduced this paradigm yielded results that were not only statistically significant, but also of a substantial magnitude (up to 30% differences in visual performance). This sensitivity and robustness may be especially useful here, allowing us to explore finer details of the temporal dynamics of perceived chasing.

Experiment 1:
Interrupting Chasing with Random Motion

We first investigated how the perception of chasing would be influenced by periodically interrupting it with sub-intervals of random motion (Figure 4), varying the relative durations of chasing vs. random motion within each interval (Figure 2).

Methods

Participants. Eight Zhejiang University undergraduates participated in a single 30-minute session for payment.

Materials and procedure. The displays were presented on a PC computer using custom software written in Matlab using the Psychophysics Toolbox libraries (Brainard, 1997; Pelli, 1997). Participants sat without head restraint approximately 50 cm from the monitor (with measurements below computed based on this viewing distance). The visible black background subtended 32°×24°. Participants directly controlled the trajectory of the sheep (drawn as a 0.87° green disc) with the mouse. To avoid making it trivially easy to keep away from all other discs, the display was crowded, containing 20 additional white discs (1.3°), one of which was the wolf (see Figure 3). The distractors all moved independently, and could pass over each other unimpeded (instead of bouncing off each other). All distractors moved haphazardly, randomly changing direction within a 240° window (centered on the most recent displacement direction) every 166.7 ms (initially jittered for each distractor, so that changes in their directions were not synchronized).

The wolf's motion consisted of two repeated sub-intervals: a *Chasing* sub-interval followed by a *Random Motion* sub-interval (see Figure 4). During Chasing sub-intervals, the wolf adjusted its direction of motion every 166.7 ms, so that it was displaced directly toward the current position of the sheep. During Random Motion sub-intervals, the wolf did not chase the sheep at all, but rather moved in a new randomly chosen direction every 166.7 ms. A pair of Chasing and Random Motion sub-intervals together lasted 833 ms (during which the wolf changed its motion direction

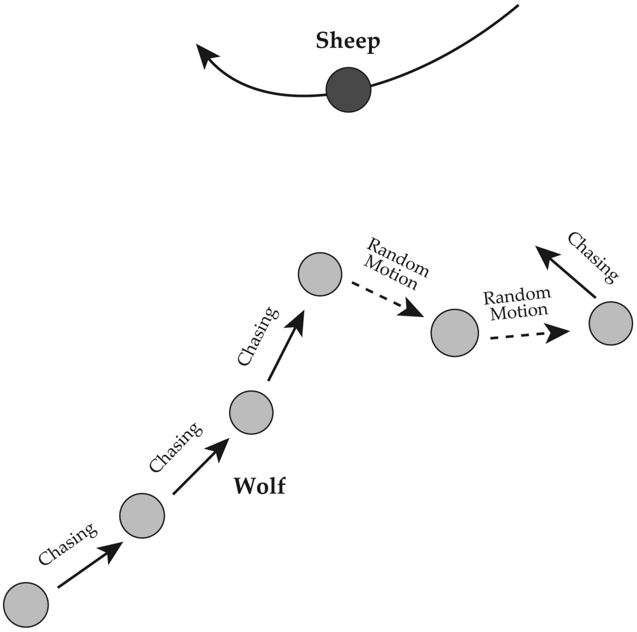

Figure 4. A cartoon depiction of a possible wolf trajectory with a Chasing Percentage of 60%. The wolf first heads directly toward the sheep's (moving) position for 500 ms (during the Chasing sub-interval), during which it updates its motion direction three times. Then it stops chasing the sheep and moves randomly for 333 ms (during the Random Motion interval), during which it updates its motion direction twice. Then this cycle repeats, with the wolf again chasing the sheep. Each disc in this cartoon depiction indicates the location at which the wolf updates its motion direction.

5 times), and this overarching interval repeated until the end of the trial. Across trials, the percentage of each 833 ms interval devoted to the Chasing sub-interval was randomly selected from one of 6 values: 0%, 20%, 40%, 60%, 80%, and 100%. For example, 40% chasing indicated that the wolf chased the sheep during each interval for 333 ms (during which the wolf updated its motion direction twice), after which it moved randomly for 500 ms (during which the wolf updated its motion direction three times). (During 100% Chasing trials, the wolf chased the sheep continuously without any interruption. During 0% Chasing trials, the wolf always moved randomly.)

The wolf-sheep distance was initially greater than 11°. The user-controlled sheep was not allowed to touch the display border during a trial—a restriction that prevented participants from circling around the display perimeter without detecting the wolf. Trials ended when (a) the sheep touched the border; (b) the wolf-sheep distance became less than 2° (*caught!*); or (c) 15 s had elapsed (*escaped!*). During the first 10 s, the maximum wolf/distractor speed gradually increased from 7.8°/s to 12°/s, while the user-controlled sheep's maximum speed gradually increased from 15.6°/s to 24°/s (i.e. twice the maximum speed of the wolf and distractors).

There were 132 randomly ordered trials, 22 for each of the six possible Chasing Percentages. Participants first completed 24 practice trials, the results of which were not recorded. During the first 12 practice trials, the wolf was always marked in a distinctive color, so that participants could become familiar with the wolves' motion patterns with different Chasing Percentages. During the final 12 practice trials, the wolf looked identical to the distractors, as in the experimental trials.

Results and Discussion

Trials on which the sheep touched the display border or got caught within 1.5 s were eliminated (14.5% of the trials). The percentage of 'successful escape' trials as a function of Chasing Percentage is depicted in Figure 5, with the relevant statistical comparisons given in Table 1. Overall levels of performance in this task are essentially meaningless, since we can impair performance to any desired degree simply by (for example) speeding up the movements of the wolves, or increasing the number of distractors (all of which are potential wolves at the start of a trial, as far as the participant is concerned). In fact, we set these parameters during pilot testing in order to bring performance to a level that was close to neither ceiling (perfect escapes) nor floor (infrequent escapes); in this range, the critical data relate to the shape of the performance curve as a function of Chasing Percentage.

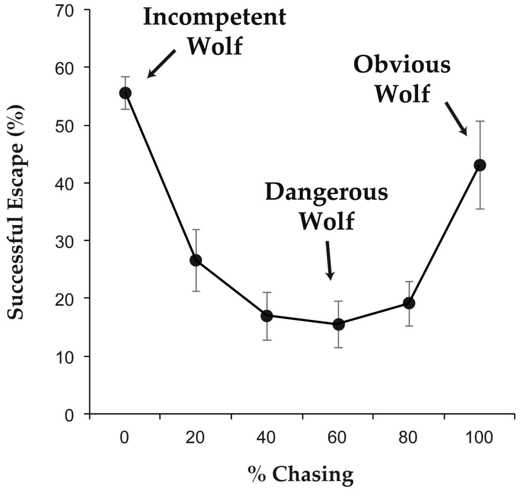

Figure 5. Results from the Random Motion manipulation in Experiment 1: The percentage of trials in which the participant successfully avoided being caught by the wolf, as a function of the Chasing Percentage.

The clear U-shaped function in Figure 5 dramatically reveals both the accuracy of *perceived* chasing and the ultimate import of *actual* (statistical) chasing. Performance was good with both the lowest and highest Chasing Percentage values—but for different reasons. With 100% Chasing (i.e. with uninterrupted chasing), it was easy for participants to escape, because they readily detected the wolf chasing them, and so could avoid it; we have thus labeled this point as the "Obvious Wolf" in Figure 5. With low Chasing Percentages, in contrast, participants were not really *avoiding* the wolf, since they could not detect it; rather, the wolf was just "chasing" them very ineffectively. Indeed, the 0% case is a baseline, with no actual wolf/sheep correlation present at all, such that participants were "caught" only by chance; we have thus labeled this point as the "Incompetent Wolf" in Figure 5.

The most intriguing results are those with moderate Chasing Percentages (between 40% and 80%), where performance was significantly worse than either continuous chasing or random motion (see Table 1). Evidently the wolf was able to catch the sheep in this range because the wolf could still effectively pursue the sheep, but its approach could not be readily detected by the participants. We have thus labeled these points as "Dangerous Wolves" in Figure 5.

The overarching pattern of these results could be explained by either the *Accumulating Positive Evidence* or *Detecting Negative Evidence* models of chasing detection, since as Chasing Percentage decreased, there was always less chasing (i.e. less positive evidence) but more random motion (i.e. more negative evidence). However, it is striking that performance declined so quickly as Chasing Percentage decreased, suggesting that only a relatively small interval of non-chasing motion is sufficient to frustrate chasing detection. In this sense, these results reveal a dissociation between perceived chasing and actual chasing. With moderate Chasing Percentages, the wolves can get efficiently closer and closer to the sheep, but they are in effect masking this pattern by periodically interrupting their pursuit. The behavior of the wolves in this special range of "unperceived chasing" might thus be likened to a new type of (imperceived) *stalking*.

Experiment 2:
Interrupting Chasing with Periods of Static Rest

The impairments of chasing detection due to random-motion interruptions in Experiment 1 could have been due to either less positive evidence for chasing, or more negative evidence against chasing. To unconfound these possibilities, the random-motion interruptions in this experiment were replaced with periods of static rest, during which the wolf simply paused momentarily on the display (as did the distractors, asynchronously). If the results of Experiment 1 were due only to less evidence in favor of chasing as

Table 1
Results of Paired t-Tests From the Random Motion Manipulation in Experiment 1, for Each Pair of Chasing Percentages

	0%	20%	40%	60%	80%	100%
0%	—	$t = 5.160$ $p = .001$	$t = 7.696$ $p < .001$	$t = 9.449$ $p < .001$	$t = 7.415$ $p < .001$	$t = 1.914$ $p = .097$
20%		—	$t = 3.219$ $p = .015$	$t = 1.857$ $p = .106$	$t = 2.306$ $p = .054$	$t = 1.694$ $p = .134$
40%			—	$t = 0.316$ $p = .761$	$t = 0.988$ $p = .356$	$t = 3.430$ $p = .011$
60%				—	$t = 1.090$ $p = .312$	$t = 4.412$ $p = .003$
80%					—	$t = 2.880$ $p = .024$

Note. ($df = 7$).

the Chasing Percentage of each interval decreased, then performance in this experiment should be identical—since the interruptions are just as long and frequent. But if the results of Experiment 1 also reflect an active *rejection* of chasing interpretations due to the random (i.e. non-chasing) motion, then impairments should disappear in this experiment.

Methods

This experiment was identical to Experiment 1 except as noted here, and involved ten new Zhejiang University undergraduates. The visible black background subtended 33°×24°. There were 24 white discs, whose motions consisted of two repeated sub-intervals: Random Motion and Static-Rest. Each white disc (both the wolf and distractors, as described below) stopped moving during its Static-Rest period. There were four discs for each of six Distractor Motion Percentage values: 0%, 20%, 40%, 60%, 80%, and 100%. A 100% Distractor Motion Percentage indicated that the disc never stopped moving, while a 0% Distractor Motion Percentage indicated that the disc never moved. Across trials, the wolf was added into the display by randomly replacing one of the distractors. The only difference between the wolf and the replaced distractor was that during its motion, the wolf chased the sheep instead of moving randomly. With this distribution of the distractors, every disc had the same probability of being the wolf, and participants could not identify the wolf merely by observing how often an item paused during its motion. Because of this, the wolves' motions were thus identical to those of Experiment 1, except that the Random Motion sub-intervals were replaced by the Static-Rest sub-intervals. There were 126 randomly-ordered trials, 21 for each Chasing Percentage (defined as in Experiment 1).

Note that this Static-Rest manipulation may make chasing relatively less objectively effective in principle, given the interruptions. However, this would not in turn lead to any better performance on our task itself, by design: With intermediate (40–60%) Chasing Percentages, the wolf typically covers the initial wolf-sheep distance in only 3 s—such that the user-controlled sheep would still be efficiently and quickly caught if the subjects did not actively avoid the wolf. However, it is of course very difficult (indeed, close to impossible) to keep away from a wolf that you do not detect in this task. As a result, our relatively long chasing period (15 s) guarantees in practice that subjects will not fail to be caught just by chance with moderate Chasing Percentages, when they do not actually detect the wolf.

Results and Discussion

Trials on which the sheep touched the display border or got caught within 1.5 s were eliminated (9.8% of the trials). The percentage of 'successful escape' trials as a function of Chasing Percentage is depicted in Figure 6, with the relevant statistical comparisons given in Table 2. Interrupting chasing with periods of static rest did not impair performance: wolves with 20% to 80% Chasing Percentages were avoided just as easily as those with 100% Chasing Percentages, indicating that this static-rest manipulation did not give rise to imperceived "stalking." Overall, these results look very different from those of Experiment 1 (compare Figures 5 and 6), and suggest that the impairment due to random motion reflects the detection of active evidence against chasing interpretations.

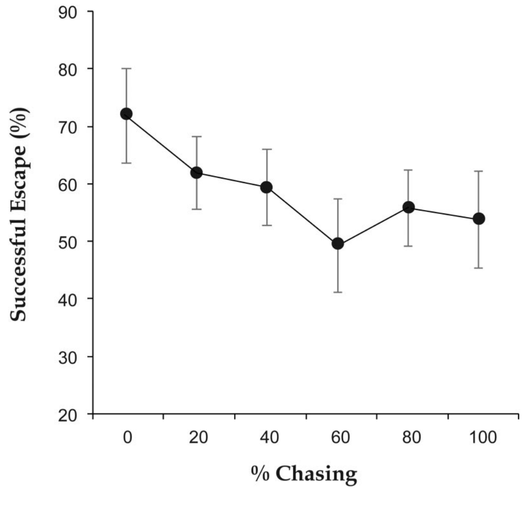

Figure 6. Results from the Static Rest manipulation in Experiment 2: The percentage of trials in which the participant successfully avoided being caught by the wolf, as a function of the Chasing Percentage.

Table 2

Results of Paired t-Tests From the Static Rest Manipulation of Experiment 2, for Each Pair of Chasing Percentages

	0%	20%	40%	60%	80%	100%
0%	—	t = 1.597 p = .145	t = 1.611 p = .142	t = 2.219 p = .054	t = 1.665 p = .130	t = 1.681 p = .127
20%		—	t = 0.663 p = .524	t = 1.539 p = .158	t = 0.742 p = .477	t = 0.955 p = .364
40%			—	t = 1.148 p = .281	t = 0.388 p = .707	t = 0.644 p = .536
60%				—	t = 1.444 p = .183	t = 1.564 p = .152
80%					—	t = 0.398 p = .700

Note. (df = 9).

Experiment 3: Interrupting Chasing with Local Motion

Random Motion intervals impaired chasing-detection performance (in Experiment 1), but static-rest intervals did not (in Experiment 2). We have interpreted this difference in terms of whether there is active "negative evidence" against chasing interpretations, but there is another possibility: rather than resulting from the specific character of the random motion, could this difference merely reflect the *existence* of non-chasing motion in one case but not the other (perhaps due to some form of attention capture)? To find out, the interruptions in this experiment were replaced with periods of local "jiggling motion," during which the wolf simply moved in place on the display (as did the distractors) during interruptions (Figure 7).

Methods

This experiment was identical to Experiment 2 except as noted here, and eleven new Zhejiang University undergraduates participated. The Static Rest sub-intervals were replaced by *Local Motion* sub-intervals, during which each disc jiggled around a fixed location: the disc was first displaced in a randomly selected direction (displaced by between 0.13° and 0.2°, depending on the current distractor speed) for 16.7 ms (i.e. one video frame) and then appeared back in its original location on the subsequent frame. This sequence was then repeated during each Local Motion sub-interval, and appeared as local apparent-motion "jiggling." (As in Experiment 2, the distractors engaged in similar behavior—alternating between random-motion and "local motion," with all local motions desynchronized across items—such that the existence and extent of the local motion could not by itself identify the wolf.)

Results and Discussion

Trials on which the sheep touched the display border or got caught within 1.5 s were eliminated (8.1% of the trials). The percentage of "successful escape" trials as a function of Chasing Percentage is depicted in Figure 8, with the relevant statistical comparisons given in Table 3. This overall pattern is a similar U-shaped function, as in Experiment 1, but is much less severe. This can be appreciated by examining Figure 9, in which the

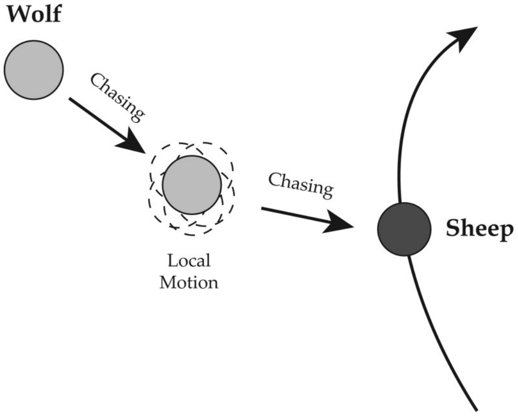

Figure 7. A cartoon depicting a 60% Chasing Percentage trial, interrupted by Local Motion in Experiment 3. After chasing the sheep for 500 ms, the wolf jiggles in place for 333 ms, after which it resumes its pursuit.

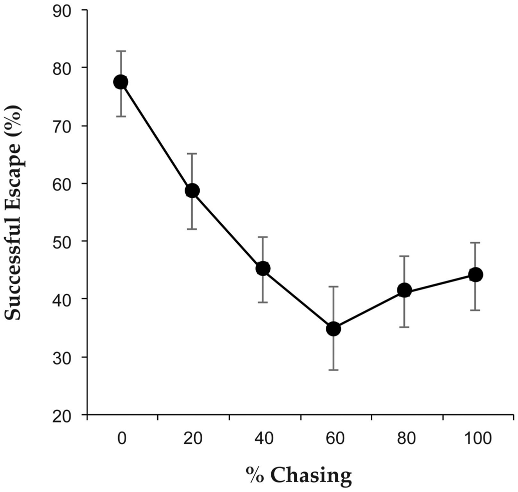

Figure 8. Results from the Local Motion manipulation in Experiment 3: The percentage of trials in which the participant successfully avoided being caught by the wolf, as a function of the Chasing Percentage.

Table 3

Results of Paired t-Tests From the Local Motion Manipulation of Experiment 3, for Each Pair of Chasing Percentages

	0%	20%	40%	60%	80%	100%
0%	—	$t = 3.451$ $p = .006$	$t = 5.662$ $p < .001$	$t = 5.879$ $p < .001$	$t = 6.545$ $p < .001$	$t = 4.842$ $p = .001$
20%		—	$t = 3.210$ $p = .009$	$t = 4.413$ $p = .001$	$t = 3.313$ $p = .008$	$t = 2.131$ $p = .059$
40%			—	$t = 2.201$ $p = .052$	$t = 0.843$ $p = .419$	$t = 0.231$ $p = .822$
60%				—	$t = 1.546$ $p = .153$	$t = 2.285$ $p = .045$
80%					—	$t = 0.617$ $p = .551$

Note. ($df = 10$).

results of Experiments 1–3 are all summarized. As can be appreciated from this figure, there was some impairment in this experiment due to the presence of motion per se, since the Local Motion curve lies under the Static Rest curve. However, the primary result of this study is the impairments due to Local Motion were not nearly as severe as those from Random Motion: the curve for Random Motion in Figure 9 shows a much steeper decline and rise—and with the greatest impairment in each condition (both at 60% Chasing) more than 20 percentage points worse with Random Motion compared to Local Motion. This supports the possibility that the specific character of the random motions impair the detection of chasing, beyond the fact that they necessitate less evidence in favor of chasing.

Although the random motion condition of the previous experiment and the local motion condition of this experiment employed the same speeds, they looked entirely different: in the random-motion condition, the items appeared to be *moving*, but in the local motion condition, the items appeared to be jiggling *in place*. This distinction is reminiscent of the contrast between "manner" and "path" in psycholinguistic treatments of motion: the random-motion condition varied the *paths* that the items took, while the local motion condition simply varied the *manner* in which the object moved (or in this case, the manner in which the object remained stationary!).

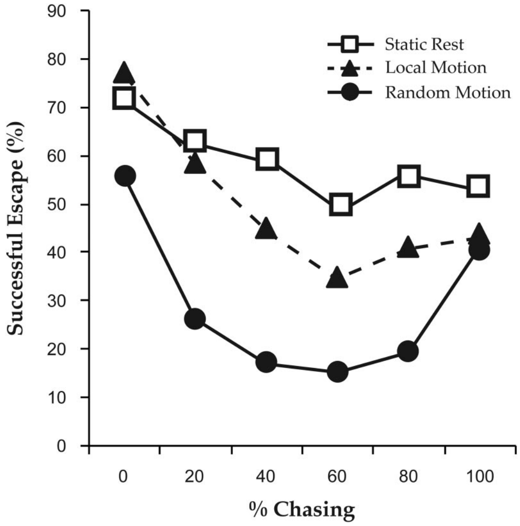

Figure 9. A summary of the results from Experiments 1–3, plotted together.

Experiment 4: Within-Subjects Interruptions

The interpretations above required comparisons between Experiments 1–3, as in Figure 9. However, it is not possible to directly analyze the differences between performance in these cases, since the experiments were run at different times, with different participants. In addition, the distactors' motions in these experiments were also different, since during the Non-Chasing sub-intervals, the distractors always mimicked the behavior of the wolves. To verify that such differences still hold when analyzed in a single group, this experiment replicated Experiments 1–3 in a within-subjects design, using a moderate Chasing Percentage (50%). The distractors in this experiment were identical in all conditions, with several distractors on each trial respectively engaging in random motion, static rest, or local motion—so that all that differed was the wolf's behavior.

Methods

This experiment was identical to Experiments 2 and 3 except as noted here, and ten new Zhejiang University undergraduates participated. The Chasing Percentage was fixed at 50% for all trials, so as to fit each condition into a single testing session, such that each Chasing and Non-Chasing sub-interval was 417 ms. During Non-Chasing sub-intervals, the wolf engaged in one of four behaviors (with that behavior remaining constant throughout each trial): it moved randomly (as in Experiment 1), remained static (as in Experiment 2), moved locally (as in Experiment 3), or continued chasing the sheep (as a baseline).

The display on every trial contained four types of distractors, corresponding to each type of possible wolf behavior during Non-Chasing sub-intervals. Distractors always moved randomly during Chasing sub-intervals, and then either stopped (five distractors/trial), moved locally (five distractors/trial), or continued moving randomly (10 distractors/trial) during Non-Chasing sub-intervals. (As in all experiments, these intervals were always desynchronized across all objects in the display.) These percentages insured that the wolf could never be identified on the basis of its type of motion during the interruptions. On each trial, there were 20 white discs (five for each type of distractor), with the wolf added to the display by randomly replacing one of the distractors, as in the previous experiments. (Thus, for example, Local Motion trials involved one wolf that periodically jiggled, along with four jiggling distractors, five periodically paused distractors, and 10 randomly moving objects that were never interrupted.) Participants completed 28 trials for each of the four possible wolf behaviors in a random order, for a total of 112 trials.

Results and Discussion

Trials on which the sheep touched the display border (9% of the trials) or got caught within 1.5 s (10.8%) were eliminated.[2] The percentage of "successful escape" trials as a function of the wolf's behavior during Non-Chasing sub-intervals is depicted in Figure 10, with the relevant statistical comparisons given in Table 4. Except for the Local Motion and Continuous Chasing conditions, performance in each condition differed significantly from the others—with a larger impairment on Random Motion trials than in any other condition. Note also that this greater impairment for Random Motion trials was again not only statistically significant but of a considerable magnitude (almost 20% worse than in the next-worst condition). Finally, this experiment makes it very clear that the impairments are not simply due to less overall chasing, since in this experiment performance was actually a bit *better* for Static Rest trials than for Continuous Chasing trials.[3] These results thus replicate the essential features of the previous experiments, and support their conclusions.

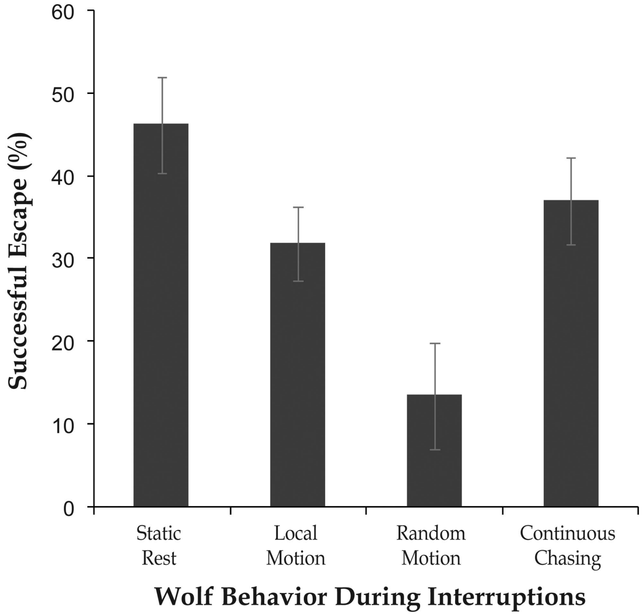

Figure 10. Results from the within-subjects manipulations in Experiment 4: The percentage of trials in which the participant successfully avoided being caught by the wolf, as a function of the wolf's behavior during the temporal interruptions.

Table 4

Results of Paired t-Tests From the Within-Subjects Manipulations of Experiment 4, for Each Pair of Possible Wolf Behaviors During Non-Chasing Sub-Intervals

	Static rest	Local motion	Random motion	Continuous chasing
Static rest	—	$t = 4.588$ $p = .001$	$t = 8.128$ $p < .001$	$t = 2.760$ $p = .022$
Local motion		—	$t = 3.304$ $p = .009$	$t = 1.598$ $p = .144$
Random motion			—	$t = 5.145$ $p = .001$

Note. ($df = 9$).

Experiment 5: The Frequency of Interrupted Chasing

The impairments in chasing detection due to random motion observed in the previous experiments could simply reflect the brute amount of evidence "against" chasing—operationalized as the summed duration of Random Motion sub-intervals across a trial. However, it is also possible that not all random motion interruptions are created equal. The process responsible for detecting the wolf on the basis of its pursuit may be able to ignore or overcome especially short Random Motion interruptions, even if there are many more of them (to equate the total amount of random motion).

Here we tested this by setting the Chasing Percentage to 50% on all trials (as in Experiment 4) and testing only Random Motion interruptions, but varying their frequency: as depicted in Figure 11, the overall interval size was varied across trials between 200 and 1000 ms. Note that the overall amount of random motion (and of chasing) was always constant; all that differed was the frequency with which these sub-intervals oscillated. In truth, we conducted this experiment without any clear hypothesis about what the outcome would be, since this frequency manipulation seemed like it

---

[2] While such a high percentage of discarded trials would be unusual in many psychophysical contexts, here it is due simply to (a) the fact that we made the display very challenging (so that border-touching was always a constant danger), and (b) the inherent noise involved in such a unconstrained task, wherein the subjects could freely move the sheep through the display. Expanding on this second reason: on roughly 10% of trials, subjects did not immediately detect the wolf, and simply had the bad luck to move immediately in its direction—thus ending the trial before they had a chance to "get their bearings," as it were. This sort of statistical event is not possible in most psychophysical experiments, which involve much more constrained displays and tasks, typically in static displays. We cannot include data from these "border touching" trials in our analyses, since the trials ended immediately when this happened, and there is thus no way to categorize them as "catch" or "escape" trials. When the "early catch" trials are included in these analyses, however, nothing qualitatively changes for any of these experiments.

[3] The reason for the improved performance in the Static Rest condition relative to baseline is not completely clear, but this pattern was also subtly apparent in the results of Experiment 2, where the performance curve (see Figure 6) had a shallow negative slope. This may be due to the fact that the periods of static rest do not impair the perception and detection of chasing, but they do impair the *actual* efficacy of chasing, by slowing it down relative to the always-moving sheep.

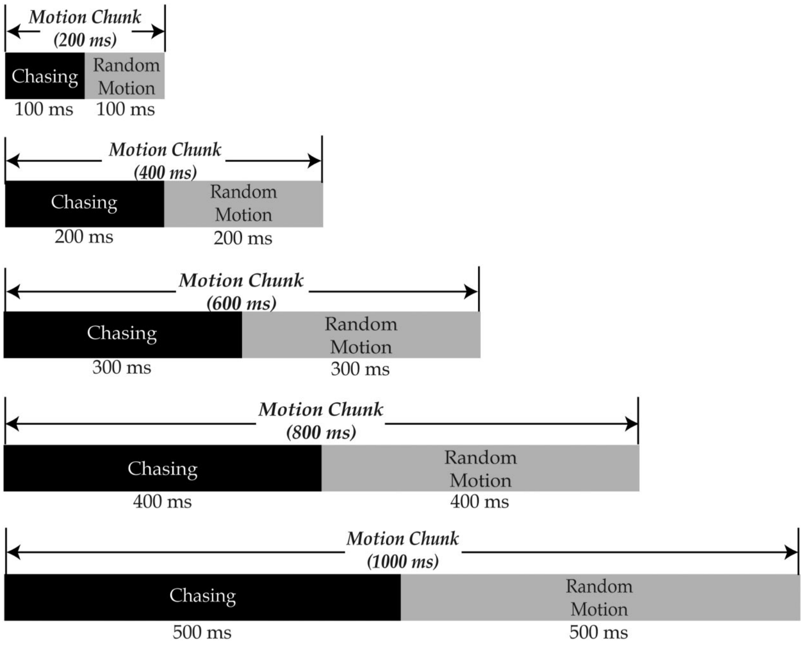

Figure 11. An illustration of the manipulation of motion interval duration and frequency from Experiment 5.

could both help and harm the detection of chasing. On one hand, the shorter Random Motion sub-intervals in this experiment (i.e. those depicted toward the top of Figure 11) may be short enough to be effectively ignored. On the other hand, though, this manipulation also necessarily resulted in much shorter Chasing sub-intervals—such that any process which required an extended period of chasing to accumulate evidence would also be frustrated. In any case, this experiment was designed to empirically settle this issue.

Methods

This experiment was identical to Experiment 4 except as noted here, and eight new Zhejiang University undergraduates participated. There were 22 discs in each display, including the wolf and 21 randomly moving distractors. All discs changed their directions every 100 ms. On $^{2}/_{7}$ of the trials, the wolf chased the sheep continuously. On the remaining $^{5}/_{7}$ of the trials, the Chasing Percentage was fixed at 50%. The interval duration was randomly selected from one of 5 values: 200, 400, 600, 800, or 1000 ms. For example, a 200-ms interval duration involved chasing for 100 ms followed by random motion for 100 ms, repeated until the trial ended. Since some of these conditions proved much harder than the previous experiments, the trial duration (for which participants had to avoid the wolf to count as an "escape") was shortened to 8 s. There were 140 randomly ordered trials presented in a random order: 40 for the continuous chasing condition, 20 for each of the other five chasing-interval durations.

Results and Discussion

Trials on which the sheep touched the display border (7.8% of trials) or got caught within 1.5 s (11.7% of trials) were eliminated. The percentage of "successful escape" trials as a function of the interval duration is depicted in Figure 12, with the relevant statistical comparisons given in Table 5. As in the previous experiments, interrupting the wolf's pursuit with periodic random motion significantly impaired chasing detection. However, the magnitude of this impairment was less for 200 ms intervals than for the other durations (all of which gave rise to equal impairments). This indicates (a) that especially short random-motion interruptions can be effectively ignored, even when there are more of them; and that

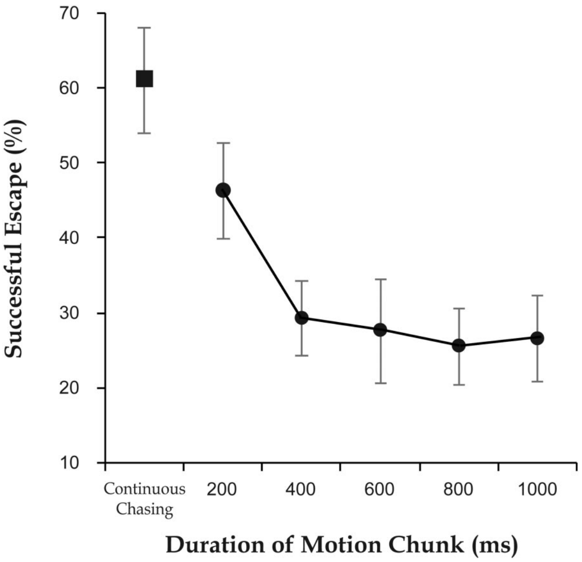

Figure 12. Results from the Motion Chunk Duration manipulation in Experiment 5: The percentage of trials in which the participant successfully avoided being caught by the wolf, as a function of the motion chunk duration.

Table 5

Results of Paired t-Tests From the Interval-Duration Manipulation of Experiment 5, for Each Pair of Interval Durations

	200 ms	400 ms	600 ms	800 ms	1000 ms	Continuous chasing
200 ms	—	$t = 3.357$ $p = .012$	$t = 3.947$ $p = .006$	$t = 3.573$ $p = .009$	$t = 2.690$ $p = .031$	$t = 6.951$ $p < .001$
400 ms		—	$t = 0.324$ $p = .755$	$t = 1.221$ $p = .261$	$t = 0.450$ $p = .667$	$t = 5.434$ $p = .001$
600 ms			—	$t = 0.497$ $p = .635$	$t = 0.203$ $p = .845$	$t = 6.387$ $p < .001$
800 ms				—	$t = 0.273$ $p = .793$	$t = 5.861$ $p = .001$
1000 ms					—	$t = 4.950$ $p = .002$

Note. ($df = 7$).

(b) evidence for chasing can be effectively accumulated across multiple chasing intervals, even when each individual chasing episode is quite short. This could occur as part of the computations involved in detecting and perceiving animacy, or it could occur at the level of initial motion integration and perception itself. It is clear, in any case, that these impairments are not just a function of the amount of random motion, but of its character and temporal grouping. For random motion to effectively impair performance, in other words, it must be distributed in the right way through the trial as a whole.

Experiment 6: The Degree of Deviation

Random motion severely disrupted the perception of chasing in the previous experiments, presumably because those random motions involve displacements that deviate from the perfect "heat-seeking" direction. But how much deviation from heat-seeking is required to "count" as evidence *against* chasing (as opposed to slightly noisy evidence *for* chasing)? This was impossible to answer in the previous experiments, since the deviations were completely random from moment to moment during Random Motion sub-intervals. In this experiment, in contrast, the deviations were fixed. There were no Random Motion sub-intervals, per se. Instead, during Non-Chasing sub-intervals, the wolf moved on a linear trajectory that was a constant angular offset to the sheep's position. We tested 7 such offsets (from 0° to 180°, in 30° steps), with the offset remaining constant within a trial, but varied randomly across trials. (A 90° offset, for example, means that during Non-Chasing sub-intervals on that trial, the wolf always moved in an orthogonal direction compared to the sheep, as depicted in Figure 13.) This experiment is thus an attempt to empirically determine the difference between evidence for and against chasing.

Methods

This experiment was identical to Experiment 5 except as noted here, and 12 new Zhejiang University undergraduates participated. The motion interval was fixed at 1000 ms, and the Chasing Percentage was fixed at 66.7%. All discs updated their motion directions every 166.7 ms. Across trials, the wolf's motion direction during Non-Chasing sub-intervals deviated from perfect heat-seeking by either 0°, 30°, 60°, 90°, 120°, 150°, or 180°. This deviation was always the same throughout each trial, but was randomly in a clockwise or counter-clockwise direction during each individual update (every 166.7 ms). There were 21 distractors in each display, three corresponding to each possible degree of wolf deviation from perfect heat-seeking (which was varied only across trials). During Non-Chasing sub-intervals, these distractors also suddenly deviated from their previous random directions to

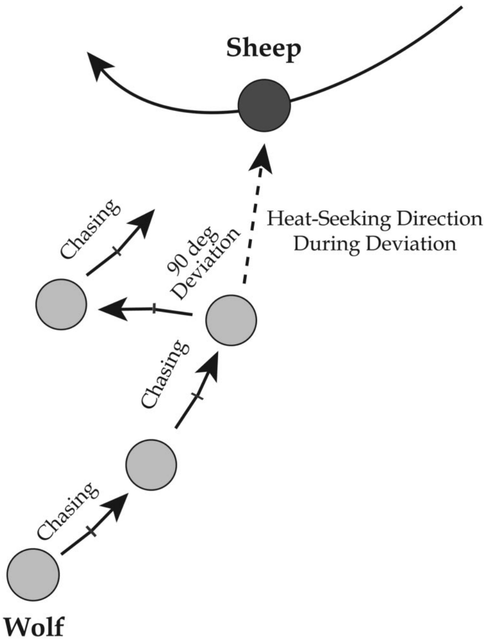

Figure 13. A cartoon depicting a 90° deviation during the wolf's Non-Chasing intervals in Experiment 6. Each individual disc (and the small markers between discs) indicate the locations at which the wolf updates its motion direction. (On each update the deviation was randomly clockwise or counterclockwise during Non-Chasing sub-intervals.) The dashed line indicates the heat-seeking direction at the moment the Non-Chasing sub-interval begins, when the wolf actually begins moving on a perpendicular trajectory.

new random directions each 166.7 ms, by their specified offset (from 0° to 180°). Because distractors corresponding to each possible offset were always included in each trial, the sudden turning radius of a disc could never be used to identify the wolf. As in the previous studies, the wolf was added into the display by randomly replacing one of its corresponding distractors. The trial duration (for which participants had to avoid the wolf to count as an "escape") was 10 s. There were 140 randomly-ordered trials, 20 for each of the seven possible deviation values.

Results and Discussion

Trials on which the sheep touched the display border (6.8% of trials) or got caught within 1.5 s (9.9% of trials) were eliminated. The percentage of "successful escape" trials as a function of the constant deviation during Non-Chasing sub-intervals is depicted in Figure 14, with the relevant statistical comparisons given in Table 6. Performance on continuous chasing trials (i.e. with 0° of spatial deviation during "Non-Chasing" sub-intervals) was significantly better than all the other conditions, except for 30°-deviation trials. The most important discovery of this experiment is thus that to impair the perception of chasing, the wolf needs only to deviate from perfect heat-seeking to a moderate degree—such that a 60° offset impairs performance just as much as a 180° offset. In other words, only 60° of deviation from heat-seeking is sufficient to count as evidence *against* chasing.[4]

General Discussion

The studies of "interrupted chasing" reported here were motivated by the observation that intentions and goal-directedness (and mental states more generally) can be in constant flux, and by the subsequent question of how the visual system updates and maintains representations of animate entities over time and motion. We explored these questions in the context of a simple case study—*chasing*—that allowed for both systematic display construction and quantitative performance measurement, but the results suggest some general lessons about the nature of perceived animacy. In particular, the results of these studies suggest: (a) that the visual system can not only detect chasing, but can also actively *reject* such interpretations; (b) that the perception of chasing diverges in interesting ways from the presence of *actual* (statistical) chasing; (c) that the perception of chasing is influenced in similar ways by deviations in space and time; and (d) that the perception of this simple form of animacy may reflect a general principle of efficiency. Below we briefly expand on each of these points.

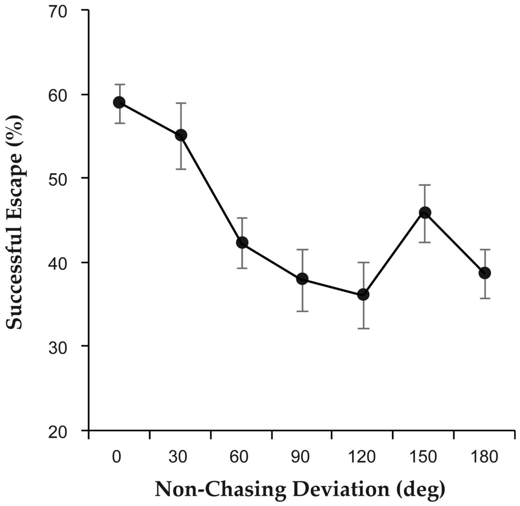

Figure 14. Results from the spatial deviation manipulation in Experiment 6: The percentage of trials in which the participant successfully avoided being caught by the wolf, as a function of the wolf's spatial deviation from perfect pursuit during Non-Chasing sub-intervals.

Constructing and Rejecting Interpretations of Chasing

Across six experiments, we varied the behavior of the "wolf" disc during haphazard but frequent brief interruptions to its pursuit of the sheep. When the wolf moved randomly during these interruptions, even a relatively small interval of random motion severely disrupted the detection of chasing (Experiment 1). This could have been for two reasons: decreased evidence in favor of chasing (consistent with the *Accumulating Positive Evidence* model), or increased evidence against chasing (consistent with the *Detecting Negative Evidence*). These interpretations were tested by having the wolf simply remain static (Experiment 2) or jiggle in place (Experiment 3) during the interruptions (among distractors that behaved similarly). In these cases, chasing detection was unimpaired (a pattern that replicated within-subjects in Experiment 4). Since the amount of actual chasing was equated in each of these cases, these results support the Detecting Negative Evidence model—suggesting that the visual system not only accumulates evidence in favor of chasing, but can also effectively reject such interpretations on the basis of input judged to be inconsistent with chasing. Experiment 6 also revealed that such "negative evidence" may be far from rare: even relatively moderate spatial deviations from perfect "heat-seeking" can serve as evidence "against" chasing, suggesting that such evidence may be plentiful in at least some situations.

At the same time, these data also provide evidence that the visual system is quite adept at accumulating positive evidence for chasing through at least some types and extents of interruptions. For example, Experiments 2 and 4 demonstrated that the detection of chasing can effectively ignore frequent static pauses in items' motions, since these potential interruptions did not impair performance. Moreover, Experiment 5 demonstrated that the detection of chasing can even ignore interruptions consisting of random motions, so long as those interruptions are especially brief (less than 100 ms)—and can do so even when the periods of consistent (uninterrupted) chasing behavior are also quite brief.

In summary, these results reveal how a type of perceived animacy is determined by the character and temporal grouping (rather than just the brute amount) of pursuit over time, and how these

---

[4] Note that these results also indicate that the impairments due to random motion in the previous experiments were not due simply to the greater lack of *correlation* between the wolf and sheep's motions during Non-Chasing sub-intervals (see also Gao et al., 2009, Experiment 2). If that had been the case, then we should not have observed any impairments in this experiment, since the wolf's motion was always correlated with the sheep's motion, albeit in different ways across the two types of sub-intervals.

Table 6

Results of Paired t-Tests From the Spatial Deviation Manipulation of Experiment 6, for Each Pair of Constant Angular Offsets of the Wolf to the Sheep During Non-Chasing Sub-Intervals

	0°	30°	60°	90°	120°	150°	180°
0	—	$t = 0.815$ $p = .433$	$t = 4.130$ $p = .002$	$t = 4.377$ $p = .001$	$t = 5.768$ $p < .001$	$t = 2.537$ $p = .028$	$t = 6.133$ $p < .001$
30°		—	$t = 2.522$ $p = .028$	$t = 3.311$ $p = .007$	$t = 3.024$ $p = .012$	$t = 1.612$ $p = .135$	$t = 2.613$ $p = .024$
60°			—	$t = 0.738$ $p = .476$	$t = 1.181$ $p = .263$	$t = 0.862$ $p = .407$	$t = 0.860$ $p = .408$
90°				—	$t = 0.274$ $p = .789$	$t = 1.829$ $p = .095$	$t = 0.139$ $p = .892$
120°					—	$t = 1.503$ $p = .161$	$t = 0.539$ $p = .601$
150°						—	$t = 1.525$ $p = .155$

Note. ($df = 11$).

temporal dynamics can lead the visual system to either construct or actively reject interpretations of chasing.

Perceiving Chasing

In the previous section we discussed the support and rejection of "interpretations" of chasing. We emphasize, however, that these interpretations are presumably the results of unconscious inferences in the visual system, rather than being a product of high-level cognition. In the Introduction we discussed the nature of such unconscious inferences, and suggested that in this respect the perception of animacy may be analogous to the perception of other visual properties such as lightness or colinearity. Having now presented the results of the "interrupted chasing" experiments, we can re-emphasize some ways in which such data seem to reflect automatic perceptual processing rather than higher-level interpretations.

First, note that our data reflect a type of visual performance, rather than perceptual reports. The latter (common in most previous studies of perceived animacy; for discussion see Gao et al., 2009) is highly subject to "contamination" by higher-level inferences; for example, observers may report that a shape "seemed alive" not because of its actual perceptual character, but because they think that it *should* be alive based on seeing how it moves. In contrast, visual performance in the current *Don't-Get-Caught!* experiments is subject to intrinsic limitations, such that participants cannot simply *decide* to perceive chasing based on any set of parameters. Instead, this ability seems confined to only a relatively narrow range of temporal and spatial deviations. This strong dependence on subtle display details, along with a relative independence from overt beliefs and intentions, has long been taken as support for the perceptual character of such phenomena (e.g. Gao et al., 2009; Michotte 1946/1963; Newman, Choi, Wynn, & Scholl, 2008; Scholl & Tremoulet, 2000).

This is especially salient in the current study, since before each experiment, participants were carefully and explicitly informed about the precise nature of the wolf's motions. Thus, for example, participants in Experiment 1 knew all about the random-motion interruptions, and had the incentive to discount them—but they could not do so. This strongly suggests that the resulting data reflect some properties and constraints of automatic perceptual processing, rather than higher-level decisions that participants are overtly making about the contents of the displays.

Spatial and Temporal Cues to Chasing vs. Stalking

Chasing detection was surprisingly poor in this study for relatively moderate percentages of random motion (as illustrated in the depth of the U-shaped curve in Figure 5, for example). This illustrates a marked difference between perceived vs. *actual* chasing: With 40–60% of random motion during each interval, the wolf was still quite effective at catching the user-controlled sheep—but this occurred precisely because participants were unable to detect the wolf in this condition. We might thus liken this to a new form of (unperceived) visual *stalking*. More generally, this phenomenon illustrates how the perception of interrupted chasing can be objectively evaluated, beyond perceptual reports.

This pattern of data is similar to that in a previous study of chasing that explored spatial rather than temporal deviations (Gao et al., 2009). This previous study employed a similar task and similar displays, but the chasing was never temporally interrupted. Instead, the efficacy of (actual, statistical) chasing was manipulated via a novel variable termed *Chasing Subtlety*—which is essentially a measure of the maximal angular deviation of the wolf's heading compared to perfect heat-seeking. When the chasing subtlety was 0° (Figure 15a), the wolf always headed directly toward the sheep, in a "heat-seeking" manner. When the chasing subtlety was 30°, for example, the wolf was always heading in the general direction of the sheep, but was not perfectly heat-seeking: instead, it could move in any direction within a 60° window, with the window always centered on the moving sheep. In other words, in this condition, the next displacement of the wolf could be in any direction within the shaded area in Figure 15b. Thus with 90° of Chasing Subtlety, the wolf could head orthogonal to the sheep (being displaced in any direction within the shaded region of Figure 15c), but still could never move directly *away* from the sheep.

Essentially, Chasing Subtlety in our previous study was a spatial analogue to the temporal manipulation of Interrupted Chasing in the current project: the former involves spatial deviations from

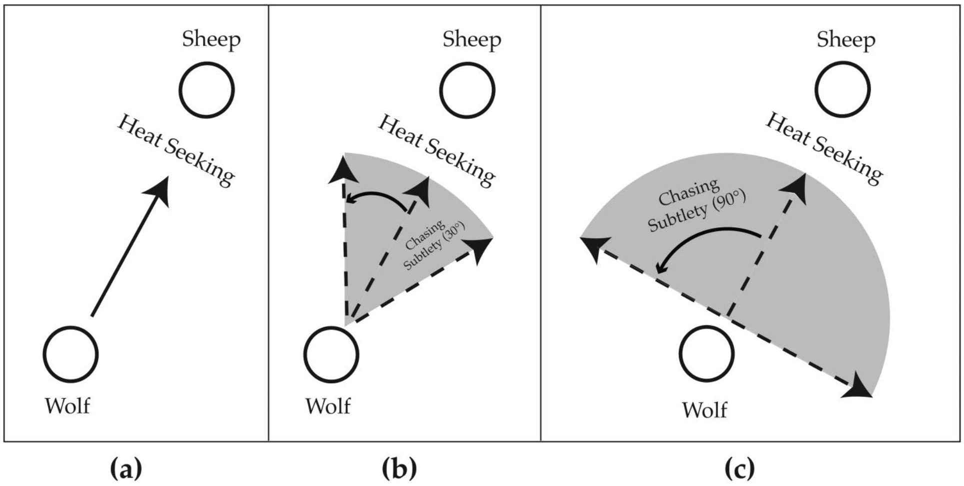

Figure 15. An illustration of the Chasing Subtlety manipulation employed by Gao, Newman, & Scholl (2009). (a) When the Chasing Subtlety is 0°, the wolf always heads directly toward the (moving) sheep, in a "heat-seeking" manner. (b) When the Chasing Subtlety is 30°, the wolf is always heading in the general direction of the sheep, but is not perfectly heat-seeking: instead, it can move in any direction within a 60° window, with the window always centered on the (moving) sheep. (c) When the Chasing Subtlety is 90°, the wolf's direction of movement is even less constrained: now the wolf may head in an orthogonal direction to the (moving) sheep, but can still never be heading away from it. The gray areas in (b) and (c) indicate the angular zone which constrains the wolf's direction of movement on that given frame of motion. (Adapted from Gao et al., 2009).

perfect heat-seeking (that were continuously present throughout each trial), while the latter involves analogous temporal deviations (from what was otherwise always perfect heat-seeking behavior). And strikingly, the results of these two studies were extremely similar: both produced a U-shaped performance curve as depicted in Figure 5. In particular, the most "dangerous" wolves were those with moderate Chasing Subtleties (in the previous study; e.g. 60°), or with moderate Chasing Percentages (in the current study; e.g. 40–60%). In both cases, the wolf could efficiently and effectively approach the sheep without being detected. Interestingly, both types of "stalking" may derive from the same underlying explanation: whatever process detects chasing has little tolerance for noise. Moderate spatial or temporal deviations from perfect chasing did little to diminish their objective effectiveness, but led to dramatic impairments in the ability to *perceive* chasing.

Beyond these basic spatial and temporal effects, perhaps the most efficient type of "stalking" we have discovered in this project as a whole comes from the current Experiment 6, which essentially investigated how mild temporal and special deviations could *collectively* disrupt the perception of chasing. The results of this experiment indicated that to effectively mask its pursuit, the wolf does not have to temporally deviate from chasing with unconstrained random motion (as in Experiment 1), or to spatially deviate from heat-seeking consistently (as in Gao et al, 2009). Instead, the perception of chasing can be severely disrupted when the wolf merely deviates from chasing in relatively brief bursts (e.g. $^1/_3$ of the time) and via a moderate angle (e.g. 60°). This "super stalker" highlights the fact that the perception of chasing is disrupted by motion with a certain character, rather than just by the brute amount of deviation from heat-seeking behavior.

Conclusions: Perceiving Animacy and the Rationality Principle

One might not intuitively expect these strict constraints on perceived chasing from earlier studies of perceived animacy that used Heider-&-Simmel-like displays with subjective verbal reports, since the perception of animacy in those displays seems so effortless and efficient. Indeed, previous discussions of the perceived animacy have commonly noted the universality of this reflex, and have suggested that this process frequently results in false alarms (cf. Miedaner, 1981). The perception of chasing in our experiments may still be "automatic," but this automaticity has two components: (a) Chasing will be perceived efficiently (and seemingly effortlessly, when viewing the relevant displays) whenever the requisite cues are present (cf. Gao et al., 2010); but (b) Chasing will *not* be perceived efficiently (despite your intentions) when those cues are not present. In this context, the current results suggest that the degree of pursuit must be especially temporally cohesive in order to trigger the perception of chasing.

Having documented these relatively strict constraints, we may still ask *why* they exist. In fact, we think that these results may reflect a more general principle that governs the perception of animacy: the *rationality principle*, which states that intentional agents will tend to choose actions that achieve their desires most efficiently, given their beliefs about the world (Dennett, 1987; see also Gergely & Csibra, 2003). The operation of this principle in perceiving animacy has been supported, for example, by infant studies: When viewing the movements of simple shapes, infants expect perceived "agents" to move toward their spatial goals via the shortest available routes, or the routes that otherwise require the least effort (Gergely et al., 1995; Southgate & Csibra, 2009).

(For models that explicate this principle in computational terms, see Baker, Saxe, & Tenenbaum, 2009.)

How does the rationality principle apply in the present study? In our displays the shortest and most efficient paths between the wolf and its goal (viz. the sheep) is perfect heat-seeking. Apparently, the wolves in our displays violate this principle when their motions deviate temporally from this perfect pursuit trajectory by even a moderate degree. Similarly, even moderate spatial deviations (in Experiment 6 or in Gao et al., 2009) may violate this principle and thus attenuate or eliminate the perception of chasing. This thus represents a possible general principle of perceived animacy that could unite many distinct cues—a possibility that we are currently exploring in the next steps of this research program.

References

Abell, F., Happé, F., & Frith, U. (2000). Do triangles play tricks? Attribution of mental states to animated shapes in normal and abnormal development. *Journal of Cognitive Development, 15,* 1–16.

Adelson, E. H. (2000). Lightness perception and lightness illusions. In M. Gazzaniga (Ed.), *The new cognitive neurosciences,* (2nd ed., pp. 339–351). Cambridge, MA: MIT Press.

Baker, C., Saxe, R., & Tenenbaum, J. (2009). Action understanding as inverse planning. *Cognition, 113,* 329–349.

Barrett, H. C., Todd, P. M., Miller, G. F., & Blythe, P. (2005). Accurate judgments of intention from motion alone: A cross-cultural study. *Evolution and Human Behavior, 26,* 313–331.

Bassili, J. (1976). Temporal and spatial contingencies in the perception of social events. *Journal of Personality and Social Psychology, 33,* 680–685.

Blythe, P. W., Todd, P. M., & Miller, G. F. (1999). How motion reveals intention: Categorizing social interactions. In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 257–286). New York: Oxford University Press.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10,* 433–436.

Csibra, G. (2008). Goal attribution to inanimate agents by 6.5-month-old infants. *Cognition, 107,* 705–717.

Dasser, V., Ulbaek, I., & Premack, D. (1989). The perception of intention. *Science, 243,* 365–367.

Dennett, D. C. (1987). *The intentional stance.* Cambridge, MA: The MIT Press.

Dittrich, W., & Lea, S. (1994). Visual perception of intentional motion. *Perception, 23,* 253–268.

Gao, T., McCarthy, G., & Scholl, B. J. (2010). The Wolfpack effect: Perception of animacy irresistibly influences interactive behavior. Psychological Science, 21, 1845–1853.

Gao, T., Newman, G. E., & Scholl, B. J. (2009). The psychophysics of chasing: A case study in the perception of animacy. *Cognitive Psychology, 59,* 154–179.

Garfield, P. (2001). *Universal dream key: The 12 most common dream themes around the world.* New York: HarperCollins Publishers.

Gelman, R., Durgin, F., & Kaufman, L. (1995). Distinguishing between animates and inanimates: Not by motion alone. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 150–184). Oxford: Clarendon Press.

Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naïve theory of rational action. *Trends in Cognitive Sciences, 7,* 287–292.

Gergely, G., Nádasdy, Z., Csibra, G., & Bíró, S. (1995). Taking the intentional stance at 12 months of age. *Cognition, 56,* 165–193.

Gregory, R. L. (1980). Perceptions as hypotheses. *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences, 290,* 181–197.

Heberlein, A. S., & Adolphs, R. (2004). Impaired spontaneous anthropomorphizing despite intact perception and social knowledge. *Proceedings of the National Academy of Sciences, 101,* 7487–7491.

Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *American Journal of Psychology, 57,* 243–259.

Klin, A. (2000). Attributing social meaning to ambiguous visual stimuli in higher functioning autism and Asperger syndrome: The social attribution task. *Journal of Child Psychology and Psychiatry, 41,* 831–846.

Marr, D. (1982). *Vision.* New York: W. H. Freeman.

Michotte, A. (1946/1963). *La perception de la causalité.* (1946) Louvain: Institut Supérior de Philosophie. English translation of updated edition by T. Miles & E. Miles, *The perception of causality* (1963) New York: Basic Books.

Miedaner, T. (1981). The soul of the Mark III Beast. In D. Hofstadter & D. Dennett (Eds.), *The mind's I* (pp. 109–113). New York: Bantam.

Newman, G. E., Choi, H., Wynn, K., & Scholl, B. J. (2008). The origins of causal perception: Evidence from postdictive processing in infancy. *Cognitive Psychology, 57,* 262–291.

Opfer, J. (2002). Identifying living and sentient kinds from dynamic information: The case of goal-directed versus aimless autonomous movement in conceptual change. *Cognition, 86,* 97–122.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10,* 437–442.

Revonsuo, A. (2000). The reinterpretation of dreams: An evolutionary hypothesis of the function of dreaming. *Behavioral and Brian Sciences, 23,* 793–1121.

Rochat, P., Morgan, R., & Carpenter, M. (1997). Young infants' sensitivity to movement information specifying social causality. *Cognitive Development, 12,* 537–561.

Rochat, P., Striano, T., & Morgan, R. (2004). Who is doing what to whom? Young infants' developing sense of social causality in animated displays. *Perception, 33,* 355–369.

Rock, I. (1983). *The logic of perception.* Cambridge, MA: MIT Press.

Rutherford, M. D., Pennington, B. F., & Rogers, S. J. (2006). The perception of animacy in young children with autism. *Journal of Autism and Developmental Disorders, 36,* 983–992.

Santos, N. S., David, N., Bente, G., & Vogeley, K. (2008). Parametric induction of animacy experience. *Consciousness and Cognition, 17,* 425–437.

Scholl, B. J., & Tremoulet, P. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences, 4,* 299–309.

Southgate, V., & Csibra, G. (2009). Inferring the outcome of an ongoing novel action at 13 months. *Developmental Psychology, 45,* 1794–1798.

Tavares, P., Lawrence, A., & Barnard, P. (2008). Paying attention to social meaning: An fMRI study. *Cerebral Cortex, 18,* 1876–1885.

Tremoulet, P. D., & Feldman, J. (2000). Perception of animacy from the motion of a single object. *Perception, 29,* 943–951.

Tremoulet, P. D., & Feldman, J. (2006). The influence of spatial context and the role of intentionality in the interpretation of animacy from motion. *Perception & Psychophysics, 68,* 1047–1058.

Received December 30, 2009
Revision received April 19, 2010
Accepted May 28, 2010 ■