Perception of intentions and mental states in autonomous virtual agents

Peter C. Pantelis, Steven Cholewiak, Paul Ringstad, Kevin Sanik, Ari Weinstein, Chia-Chien Wu, Jacob Feldman
(petercp@eden.rutgers.edu, jacob@ruccs.rutgers.edu)
Departments of Psychology, Center for Cognitive Science, Rutgers University-New Brunswick
152 Frelinghuysen Road, Piscataway, NJ 08854 USA

Abstract

Comprehension of goal-directed, intentional motion is an important but understudied visual function. To study it, we created a two-dimensional virtual environment populated by independently-programmed autonomous virtual agents, which navigate the environment, collecting food and competing with one another. Their behavior is modulated by a small number of distinct "mental states": *exploring*, *gathering food*, *attacking*, and *fleeing*. In two experiments, we studied subjects' ability to detect and classify the agents' continually changing mental states on the basis of their motions and interactions. Our analyses compared subjects' classifications to the ground truth state occupied by the observed agent's autonomous program. Although the true mental state is inherently hidden and must be inferred, subjects showed both high validity (correlation with ground truth) and high reliability (correlation with one another). The data provide intriguing evidence about the factors that influence estimates of mental state—a key step towards a true "psychophysics of intention."

Keywords: animate motion perception; theory of mind; intentionality; action understanding; goal inference.

Introduction

Comprehension of the goals and intentions of other intelligent agents is an essential aspect of cognition. Motion is an especially important cue to intention, as vividly illustrated by the famous short film by Heider and Simmel (1944). The "cast" of this film consists only of two triangles and a circle, but the motions of these simple geometrical figures are universally interpreted in terms of dramatic narrative. Indeed, it is practically impossible to understand many naturally occurring motions without comprehending the intentions that helped cause them: a person running is interpreted as trying to get somewhere; a hand lifting a Coke can is automatically understood as a person intending to raise the can, not simply as two objects moving upwards together (Mann, Jepson, & Siskind, 1997). Much of the motion in a natural environment—and certainly some of the most behaviorally important motion—is caused by other agents, and is impossible to understand except in terms of how and why they might have caused it.

Human subjects readily attribute mentality and goal-directedness to moving objects as a function of properties of their motion (Tremoulet & Feldman, 2000), and in particular on how that motion seems to relate to the motion of other agents and objects in the environment (Blythe, Todd, & Miller, 1999; Dittrich & Lea, 1994; Gao, McCarthy, & Scholl, 2010; Pantelis & Feldman, 2010; Tremoulet & Feldman, 2006; Zacks, Kumar, Abrams, & Mehta, 2009). The broad problem of attributing mentality to others has received a great deal of attention in the philosophical literature (often under the term *mindreading*), and has been most widely studied in infants and children (Gelman, Durgin, & Kaufman, 1995; Johnson, 2000). But the adult capacity to understand animate motion in terms of intelligent behavior has been researched less. Computational approaches to the problem of intention estimation are still scarce (Baker, Tenenbaum, & Saxe, 2006; Feldman & Tremoulet, 2008), in part because of the difficulty in specifying the problem in computational terms.

Almost without exception, video stimuli used in past experiments in this area have consisted of hand-crafted animations with motions chosen subjectively by the experimenters in order to achieve particular psychological impressions (Porter & Susman, 2000). This makes it difficult to investigate the way subjects estimate intentionality, because the object of the estimation procedure—the actual mental state of the agent under observation—does not actually exist. Our proposed solution to this problem is to indeed endow our stimuli agents with "minds," which our subjects then attempt to "read."

A virtual environment of autonomous agents

We developed a two-dimensional interactive virtual environment populated with autonomous virtual agents (Fig. 1). These agents (referred to as Independent Mobile Personalities, or IMPs), are simple but cognitively independent virtual robots, equipped with perception, planning, decision making, and goals. They move about in the virtual environment, interacting with each other, making intelligent though unpredictable decisions and taking steps to achieve simple goals. The IMPs are endowed with potentially distinct personalities and cognitive faculties, including variations in intelligence, memory, aggression, and strategy. The result is a complex, dynamic microcosm in which goal-directed behavior, and the perception thereof, can be studied in a controlled way. The inspiration is taken from the substantial literature on artificial life (Shao & Terzopoulos, 2007; Yaeger, 1994) in which interactions among virtual creatures have been extensively modeled. But unlike previous environments, our agents are cognitively complete, meaning that their behavior is entirely determined by autonomous decisions based on input they have received via their own senses, and are presented visually to subjects so that we may study how their intentions are interpreted by observers. Our focus is on what can be understood from the IMPs' motion alone; to this end, we depict the IMPs as triangles, so that they have clearly identifiable main axes and front ends, but otherwise minimal shape. Because we have direct access to the "actual" intentions and mental states of the agents—represented by a simple state variable in the autonomous program—we can compare this "ground truth"

to the interpretations formed by human subjects. The investigation is truly a "psychophysics of intention," because we relate a variable in the environment (the target agent's mental state) to the subject's estimate of that variable. This in turn allows us to study the inference mechanisms that underlie human judgments about intentional action.

Agent architecture

The IMPs have a simple visual system with a 1D retina (Fig. 1B), which they use to gather information about the geometric structure of their environment (Fig. 1C,D). The vision module uses a ray-casting intersection test for objects along a small number of lines of sight (see Fig. 1B). Rays that intersect with something in the environment are grouped by color to create simple object representations, which are then added to a cumulative mental map of the environment. The IMPS' principal goal is to gather food (bits of inanimate gray material randomly positioned in the environment), though that goal is occasionally interrupted by various subgoals, like the need to fend off other agents that they encounter and the need to explore in order to create their environment map. Their decisions are carried out with the help of a path planning module that allows them to select the shortest path from their current position to remembered locations drawn from the learned map.

The IMPs' behavior is modulated by four distinct action states: *exploring*, *gathering food*, *attacking*, and *fleeing* (loosely modeled on the "Four Fs" of natural ethology—feeding, fleeing, fighting and mating; Pribram, 1960). Each IMP transitions stochastically among these states, conditional upon its current state, its mental map, and its immediate perceptual input. The states are of course not directly visible to observers, but rather must be inferred. In our experiments, we refer to them as the "ground truth" of an IMP's mental state, because they represent the actual state of the IMP's dynamic autonomous program at any given time. As reported below, we asked subjects to estimate these states and detect changes between them—in effect, to guess the transitions between qualitatively different behavioral events (Kurby & Zacks, 2008). The main focus of our analyses is on subjects' ability to estimate these transitions, which directly reflects their ability to interpret the intentional structure of the IMPs' behavior.

Experiments

Our experiments tested subjects' ability to infer the agents' mental states while viewing short (60 s) scenes. In Exp. 1, we explicitly instructed subjects about the four IMP mental states, and assigned one response key per state. In Exp. 2, we allowed subjects to freely invent their own mental states after viewing some sample scenes, inducing a less constrained (but also less transparent) response set.

Experiment 1

The first experiment tested subjects ability to successfully categorize the IMPs' behaviors and detect transitions among the


Figure 1: (a) The virtual environment with its native IMPs (depicted as moving triangles). (b) The IMPs have autonomous vision, and by (c) exploring their environment they can (d) gradually develop a mental map of the obstacles it contains.

IMPs' "mental states." In this condition, the possible underlying states were known to the subjects.

Subjects. Four undergraduate students in introductory psychology classes at Rutgers University participated in the experiment, and received course credit for their participation. Each experimental session lasted approximately 30 minutes.

Stimuli. Each subject viewed the same set of 20 scenes, generated in advance. Each pre-recorded scene was 60 seconds in duration, and was presented within a 400 x 400 pixel window, horizontally subtending approximately 13.5° of visual angle. Each scene was populated with 4 identically parameterized IMPs (described above) at randomized starting positions, 15 gray food objects (divided evenly into 3 clusters, with each cluster initially placed at a random starting position), and two square red obstacles (placed at the same locations in each scene).

Procedure. Five initial training scenes were shown. Subjects were instructed to simply observe the action and try to determine what was happening within the scenes. During training, each IMP's true mental state was reflected in its color (see Fig. 2). After the subject watched these 5 scenes, they were asked what they thought the IMPs were doing, and what the colors might mean. It was then explained to them that the colors actually corresponded to the underlying mental or behavioral state of the IMP, and that an IMP could be in one of four of these states at a given time: "attacking" another agent, "exploring" its environment, "fleeing" from another agent, or attempting to "gather" food.

Figure 2: Sample scene from Exp. 1. The red IMP is in an "attack" state, the purple IMP is "exploring," the yellow IMP is "fleeing," and the green IMP is "gathering." Note that colors were only shown during training scenes. All agents maintained the same, neutral color for the remainder of the experiment.

Each subject then viewed 15 additional scenes, the first of which was treated as practice and excluded from analysis. In these scenes, IMPs did not change color; that is, the subjects' task was to infer the underlying state of an IMP solely from its behavior and context. The target IMP was colored black, and the other 3 were colored blue. Subjects were instructed to pay attention to the black agent in each scene, and indicate on the keyboard which state they thought this target agent was in at any given time. Four keys represented the 4 respective possible states; subjects were instructed to press a key as soon after a scene began as possible, and thereafter to press a key only when they thought the target IMP had transitioned into a new state.

Results. Fig. 3 illustrates how subjects responded as they observed two example scenes. The ground truth mental state of the target IMP is shown in the top horizontal bar in each panel of the figure, with the 4 subjects' concurrent responses aligned underneath.

We first examined subjects' performance by measuring the proportion of time that their classifications matched the ground truth state of the target IMP (validity; see Table 1). Mean accuracy was 47%, with each individual subject performing significantly above chance ($p < .001$).[1]

Another critical aspect of subject performance is inter-subject agreement (reliability). We calculated intersubject response agreement by tabulating the time that two subjects gave the same response, evaluating agreement using a strict criterion (Method 1) or a more relaxed one (Method 2). By Method 1, subjects were considered to disagree if their classification differed, including if one had not yet responded. By Method 2, agreement was not calculated until both had started responding. Even by the stricter Method 1 (see Table 1), average inter-subject agreement was higher than accuracy with respect to ground truth, but using the more relaxed Method 2, inter-subject agreement becomes quite high. Subject 2, for example, responded in agreement with other subjects a full 70% of the time (in this case, chance would be 25%).

A comparison of estimated mental states to actual ones shows a number of interesting patterns, as revealed by the inter-state confusion matrix (Table 2). The analysis reveals one dominant source of subject "errors." Subjects generally did not initiate responding immediately at the start of each trial; 17% of overall trial time was prior to the initial response (And, in fact, if one excludes from analysis times when subjects had not yet provided a response, mean subject accuracy rises to 54%). As IMPs were most likely to be in the *explore* state at the beginnings of trials, these errors of omission account for a large proportion of subjects' misclassifications for this action type. Otherwise, subjects' detection of the *explore* state was nearly perfect. Accuracy was lower for the other states. For example, when an IMP was in the *gather* state, subjects were about equally likely to respond *gather* or *explore*. Correct detection of *flee* was only about 8%.

The pattern of subjects' errors seems to reflect a sensitivity to the four mental states' base rates; the IMPs spend unequal amounts of time in each state, and subjects' responses reflected that imbalance. As an illustrative example, subjects misclassified *flee* as *explore* 44% of the time. But a suitably tuned ideal observer would know that *explore* is about six times more probable overall than *flee*, and could in principle use this prior information to classify more accurately.

The next analysis explored how the IMPs' objective physical parameters related to subjects' classifications of their mental state. Among the parameters we measured, some are *intrinsic* to the agent (including its speed and angular velocity), while others are *context-dependent* in that they measure the agent's relation to other elements of the scene (including the distance to the nearest other IMP, and the distance to the nearest food). In our analysis we attempted to distinguish the role of these two classes of parameters in accounting for subjects' responses.

First, we carried out a multinomial logistic regression, attempting to predict subjects' responses as a function of the two intrinsic parameters mentioned above (speed and angular velocity). This analysis echoes that of Blythe et al. (1999),

---

[1] To determine chance performance, we simulated a random performer that would begin each trial having not yet responded on the keyboard, and at subsequent time points would randomly select a response with a probability matched to the behavior of the given subject. Under these conditions, chance performance is 20%. We performed 1000 such simulations for each individual subject, and none of these random performers achieved performance as high as their human counterparts.
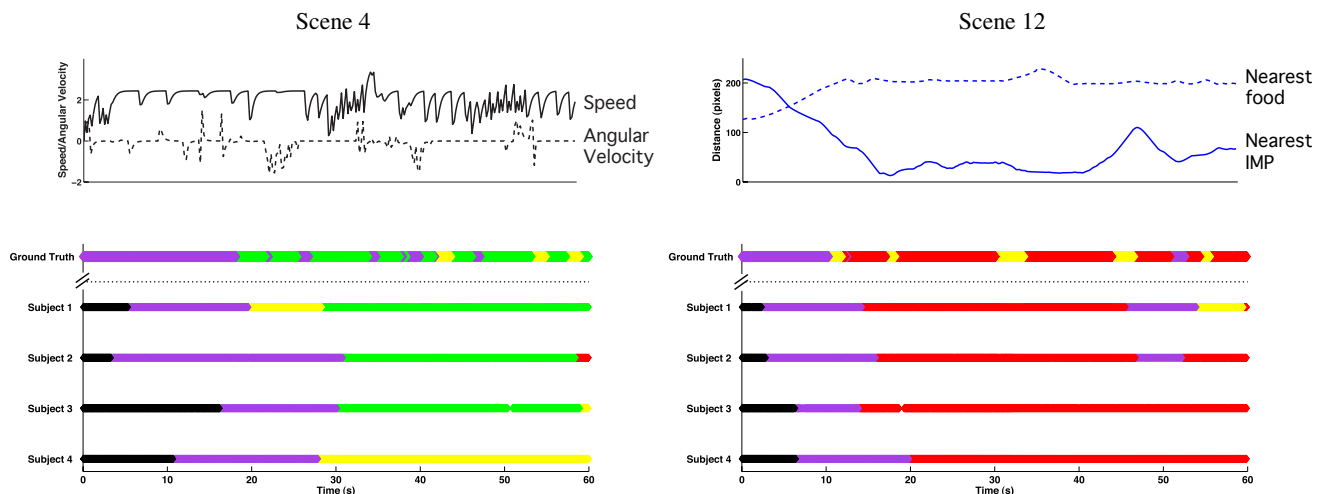
Figure 3: Two example scenes from Exp. 1 (left and right panels). Within each panel, the top figure plots physical variables related to the target IMP, over the course of the 60 s scene. In the left panel, variables intrinsic to the IMP (*solid line* = speed, *dashed line* = angular velocity) are plotted in black; in the right panel, variables related to the IMP's environmental context are plotted in blue (*solid line* = distance to nearest other IMP, *dashed line* = distance to nearest food). In the bottom figure of each panel, the top bar represents the actual "ground truth" state of the target IMP (*red* = "attack", *purple* = "explore", *yellow* = "flee", *green* = "gather"). The other bars represent the corresponding real time responses of the 4 subjects. *Black* in the subjects' bars indicates that the subject had not yet entered a response on the keyboard.

Table 1: Subject agreement with ground truth and other subjects, in Exp. 1.

	Subject			
	1	2	3	4
vs. ground truth	.50	.52	.46	.38
vs. other subjects (method 1)	.50	.56	.48	.43
vs. other subjects (method 2)	.62	.70	.67	.59

Table 2: Confusion matrix for subjects' responses in Exp. 1 (averaged across subjects). Mean proportion of IMP time spent in each state is in parentheses, and mean proportion of time subjects spent in each response category is at the bottom of each column.

	Choice				
Actual	None	Attack	Explore	Flee	Gather
Attack (.16)	.07	.39	.43	.08	.03
Explore (.42)	.29	.02	.61	.03	.05
Flee (.07)	.08	.23	.44	.08	.18
Gather (.35)	.08	.06	.35	.12	.40
	.17	.11	.48	.07	.18

who similarly classified agent motion as a function of simple, objective motion properties. Fitting the model's parameters to each subject separately yielded correct response predictions 41%, 41%, 41%, and 40% of the time (or 41% when fitting the pooled data). The same model predicts ground truth 49% of the time.

We then added the two context-dependent parameters (distance to nearest other IMP, distance to nearest food) to the model, allowing it to reflect not only the agent's motions but also its reaction to the objects around it (cf. Tremoulet & Feldman, 2006). This led to correct predictions of individual subjects' responses of respectively 47%, 62%, 46%, and 46% (43% when fitting pooled data)—somewhat better than with the intrinsic parameters alone. This model predicted ground truth 73% of the time, reflecting the fact that the ground truth state is, in fact, conditioned on context, as IMP state transitions can be triggered by nearby agents and food. These logistic models are admittedly simplistic, and clearly leave a substantial part of subjects' responses unexplained; we view them simply as baselines against which to compare a more psychologically rich future computational account.

Experiment 2

In Exp. 2, subjects performed the same task as in Exp. 1 except they classified the target agent into action categories that they themselves invented.

Subjects. Three undergraduate students participated in an approximately 30 minute experimental session in exchange for course credit.

Stimuli. Each subject viewed the same 20 pre-recorded scenes that were utilized in Exp. 1, 60 s each in duration.

Procedure. Five initial training scenes were shown, and subjects were instructed to observe the action and try to determine what was happening in the scenes. Unlike Exp. 1, the IMPs colors did not change color according to their mental states, so as not to bias the subjects' future classifications. Subjects were told that while the IMPs might have an overarching goal that they are trying to achieve, at any given time an

IMP might need to achieve a subgoal which would have a corresponding mental or intentional state. An analogy was made to basketball: the overarching goal is to score more points than the opponent, but at any given time a player might be shooting, passing, defending the goal, attempting to steal the ball, etc. Subjects were instructed that later they would be asked to classify the behavior of agents according to action categories of their own invention, and to keep this in mind as they watched. After watching the 5 scenes, each subject reported the categories they would be using. While we allowed anywhere from 2 to 6 categories, all subjects elected to use exactly 4—unlikely to be a coincidence given that this was in fact the true number of states underlying the IMP behaviors.

Subjects each then viewed 15 additional scenes, the first of which was treated as practice and excluded from analysis. The target IMP was colored black, and the other 3 were colored blue. Subjects were instructed to pay attention to the black agent in each scene, and to indicate on the keyboard which state they thought this target agent was in at any given time. Keys were assigned to each of the states previously provided by the subjects. Subjects were instructed to press a key as soon after a scene began as possible, and thereafter to press a key only when they thought the target IMP had transitioned into a new state.

Results. Because Exp. 2's subjects each used his or her own invented response categories, different subjects' responses cannot be simply aggregated as they were in Exp. 1. Instead, our analysis focuses on the overlap between each individual response pattern and ground truth, and on the overlap between the response patterns invented by the various subjects. State classification patterns among the various subjects were broadly correlated. For example, the number of state transitions perceived in each scene were correlated between subjects (Subj. 1 vs. Subj. 2: $[r(12) = .65, p < .05]$; Subj. 1 vs. Subj. 3: $[r(12) = .57, p < .05]$; Subj. 2 vs. Subj. 3 $[r(12) = .53, p = .052]$).

A more detailed analysis of the similarities among various subjects' state classifications reveals a variety of intriguing patterns (see Table 3), though the data are too rich to discuss in detail here. First, notwithstanding the different labels subjects chose for certain IMP behavior patterns, classes invented by one subject often showed consistent correlations with classes invented by others. For example, Subj. 1's labels *hesitation* and *confusion* were often associated with behavior classed by Subjs. 2 and 3 as *wandering* or *confused*, respectively—labels that were, in turn, highly correlated with each other. As another example, Subj. 1's use of *interaction* showed over 80% overlap with Subj. 2's use of *fighting*: Despite the varying lexical terms chosen (an interesting subject itself) the behavioral classes to which these terms refer are substantially identical.

Discussion

In summary, our data show that subjects are proficient at estimating our automonous agents' true mental states, both in terms of reliability (intersubjective agreement) and in terms of validity (accuracy in estimating the true IMP state). Although mental state is only implicit in the IMPs' behavior, subjects can divine it; they can "tell what the agents are thinking," and tend to concur with one another.

While the raw physical parameters of the IMP (speed, angular velocity, etc.) can usually predict the true mental state of the IMP (regression results above), the same parameters do *not* generally predict subjects' responses, suggesting that subjects' classifications are based on a different computational synthesis of features. A number of authors have argued that animacy judgments, rather than depending solely on properties of motion (e.g. the ability to expend energy and initiate motion) rest critically on an assessment of the agents' apparent mental reactions to environmental stimuli (Tremoulet & Feldman, 2006; Gelman et al., 1995). Our results reflect this phenomenon, in that subjects' classifications were more accurately predicted when IMPs' relations with their environment were taken into account. That is, a simple model reflecting only the agents' raw motions—without incorporating their context—misses some of what is driving subjects' responses.

Naturally, subjects' performance is not perfect, and we found that under-segmentation of the state trajectory was far more common than over-segmentation. That is, subjects often missed brief excursions into other states, but rarely hallucinated a transition between states when one had not occurred (see Fig. 3). Of course, most such brief excursions entail virtually no observable change in behavior. This finding merely reinforces the idea that detecting a change in intentional state is a concrete computational process that requires sufficient data or evidence in order to yield useful, robust results.

Conclusion

Our methods pave the way towards a true "psychophysics of intention," in which the subjects' perception of *psychological* characteristics of motion in the environment can be studied in the same way that perception of *physical* properties has been studied for decades. Our results confirm that subjects can indeed detect mental states systematically (though of course not perfectly) and make it possible to more directly investigate the computational mechanisms underlying this essential cognitive function—a key next step for building on these preliminary results culled from a small number of subjects. In future work, we hope to expand the range of behaviors and degree of intelligence exhibited by our IMPs, which, after all, are still extremely limited compared to human agents. Eventually, our hope is to use a future version of our environment to study comprehension of more cognitively complex phenomena—that is, to move beyond the "Four F's" and closer to the range of behavior exhibited by real human agents.

Acknowledgments

Table 3: Overlap among subjects' uses of categories in Exp. 2. *top:* Subj. 1 vs. Subjs. 2 and 3; *middle:* Subj. 2 vs. Subjs. 1 and 3; *bottom:* Subj. 3 vs. Subjs. 1 and 2. The proportion of time a subject spent in each response category is shown in parentheses. As instances when one or both of the subjects had not yet responded are not included in this analysis, many rows and columns deviate strongly from summing to 1.

Subject 1	Subject 2				Subject 3			
	"Competition"	"Fighting"	"Taking..."	"Wandering..."	"Confused"	"Collecting..."	"Eating"	"Following"
"Hesitance" (.35)	.06	.04	.30	.60	.54	.29	.03	.06
"Going for Gray Objects" (.12)	.06	.00	.59	.36	.15	.40	.42	.03
"Confusion" (.10)	.06	.03	.03	.88	.85	.02	.00	.12
"Interaction" (.10)	.02	.89	.00	.09	.59	.00	.00	.41

Subject 2	Subject 1				Subject 3			
	"Hesitance"	"Going for..."	"Confusion"	"Interaction..."	"Confused"	"Collecting..."	"Eating"	"Following"
"Competition" (.04)	.59	.19	.17	.06	.28	.27	.00	.45
"Fighting" (.11)	.17	.00	.03	.81	.59	.01	.00	.39
"Taking Food" (.18)	.58	.39	.02	.00	.04	.73	.23	.00
"Wandering/Searching" (.55)	.38	.08	.17	.02	.67	.02	.04	.05

Subject 3	Subject 1				Subject 2			
	"Hesitance"	"Going for..."	"Confusion"	"Interaction"	"Competition"	"Fighting"	"Taking..."	"Wandering..."
"Confused" (.45)	.40	.04	.19	.13	.02	.13	.01	.80
"Collecting/Grabbing" (.15)	.66	.31	.02	.00	.05	.01	.88	.06
"Eating" (.06)	.19	.81	.00	.00	.00	.00	.65	.35
"Following" (.09)	.25	.05	.14	.44	.15	.51	.01	.34

References

Baker, C. L., Tenenbaum, J. B., & Saxe, R. R. (2006). Bayesian models of human action understanding. In Y. Weiss, B. Schölkopf, & J. Platt (Eds.), *Adv. neural information processing systems 18.* Cambridge: M.I.T. Press.

Blythe, P. W., Todd, P. M., & Miller, G. F. (1999). How motion reveals intention: categorizing social interactions. In *Simple heuristics that make us smart* (pp. 257–285). New York: Oxford University Press.

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

Feldman, J., & Tremoulet, P. D. (2008). *Attribution of mental architecture from motion: towards a computational theory* (RuCCS Technical Report No. 87).

Gao, T., McCarthy, G., & Scholl, B. J. (2010). The wolf-pack effect: perception of animacy irresistibly influences interactive behavior. *Psych. Science*, *21*, 1845–1853.

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.* New York: Oxford University Press.

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

Johnson, S. C. (2000). The recognition of mentalistic agents in infancy. *Trends Cogn. Sci.*, *4*(1), 22–28.

Kurby, C. A., & Zacks, J. M. (2008). Segmentation in the perception and memory of events. *Trends Cogn. Sci.*, *12*, 72–79.

Mann, R., Jepson, A., & Siskind, J. M. (1997). The computational perception of scene dynamics. *Computer Vision and image understanding*, *65*(2), 113–128.

Pantelis, P., & Feldman, J. (2010). Exploring the mental space of autonomous intentional agents. In S. Ohlsson & R. Catrambone (Eds.), *Proc. of the 32nd Annual Conference of the Cognitive Science Society* (pp. 2554–2559). Austin, TX: Cognitive Science Society.

Porter, T., & Susman, G. (2000). Creating lifelike characters in Pixar movies. *Communications of the ACM*, *43*, 25–29.

Pribram, K. H. (1960). A review of theory in physiological psychology. *Annual Review of Psychology*, *11*(1–40).

Shao, W., & Terzopoulos, D. (2007). Autonomous pedestrians. *Graphical models*, *69*, 246–274.

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*(6), 1047–1058.

Yaeger, L. (1994). Computational genetics, physiology, metabolism, neural systems, learning, vision, and behavior; or Polyworld: Life in a new context. In C. Langton (Ed.), *Proc. artificial life III conference* (pp. 263–298). Addison-Wesley.

Zacks, J. M., Kumar, S., Abrams, R. A., & Mehta, R. (2009). Using movement and intentions to understand human activity. *Cognition*, *112*, 201–216.