Title
Exploring the mental space of autonomous intentional agents

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 32(32)

ISSN
1069-7977

Authors
Pantelis, Peter
Feldman, Jacob

Publication Date
2010

Peer reviewed

Exploring the mental space of autonomous intentional agents

Peter C. Pantelis and Jacob Feldman
(petercp@eden.rutgers.edu, jacob@ruccs.rutgers.edu)
Department of Psychology, Center for Cognitive Science, Rutgers University-New Brunswick
152 Frelinghuysen Road, Piscataway, NJ 08854 USA

Abstract

How do we use the motion of animate objects to make inferences about their intentions? We investigate this question using displays containing a number of autonomous, independently programmed agents moving about the screen and interacting with each other. Each agent behaves according to an independent autonomous program, controlled by a small number of parameters that define its "personality." We probe subjects' impressions of the similarities among the behaviors of the various agents, and then use multidimensional scaling to recover the subjective parameters defining the mental space of agent types. The most important variable turns out to be one that determines how the agent reacts to a nearby agent at one critical distance. A followup experiment suggests that variation along this parameter contributes to modulating a higher-level percept of how "hostile" or "friendly" the agents appear to be.

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

Introduction

Intelligent agents can and must distinguish between animate and inanimate objects that they encounter. Even infants make this distinction, and apparently possess a naïve theory of other beings' mental states and intentions (Gergely, Nádasdy, Csibra, & Bíró, 1995; Keil, 1994; Johnson, 2000). Socially intelligent agents naturally conceive of other humans as animate, mentalistic agents with independent perceptions and motivations. We further benefit from being able to infer the *intentions* of other agents in the environment. This is essential for understanding and predicting others' behavior, a prime skill both for chess players contemplating their moves, and gazelles and lions engaging in mutual scrutiny on the African plain.

This research explores how adult subjects use an observed agent's motion to make inferences about its mental architecture. For this task, motion is only one cue among many (Gelman, Durgin, & Kaufman, 1995), but it is a particularly salient one, with subjects readily ascribing intentionality even to simple moving geometric figures (Heider & Simmel, 1944). A handful of studies have shown that varying the motion of simple geometric figures along certain parameters (e.g. speed, trajectory) can influence the perception of animacy and intentions (Dittrich & Lea, 1994; Tremoulet & Feldman, 2000, 2006). But the factors determining these percepts are still very poorly understood.

Baker, Tenenbaum, and Saxe (2006) and Baker, Saxe, and Tenenbaum (2009) have proposed a Bayesian framework for "inverse planning," that is, inferring or estimating the goals or intentions of an agent assumed to be rational. Sloman, Fernbach, and Ewing (2009) use Bayesian belief networks to describe causal reasoning in the domain of morality. We too presume a Bayesian formulation of the problem, in which the goal is to assign a posterior probability to the mental state (behavioral disposition, goal set, payoff matrix, or some other representation of the other agent's mind) A on the basis of its motion:

$$p(A|\text{motion}) \propto p(\text{motion}|A)p(A). \tag{1}$$

Ultimately, such an inference maps a visual input (the motion observed) onto a distribution of possible agent types. The prior $p(A)$ is defined over the set of possible agent types, that is, the space of behavioral dispositions the observer is in principle willing to entertain as explanations for the observed motion. The nature and structure of this space have been discussed only very speculatively in the literature; Barrett, Todd, Miller, and Blythe (2005) have argued that it probably includes such natural action classes as chasing, courting, following, guarding, fighting, and playing. Some studies have presented subjects with scenes constructed to resemble these different "natural categories" of dyadic interaction, and demonstrate that subjects are reliably able to categorize these scenes, even in degraded forms for which motion is the only salient cue (Barrett et al., 2005; McAleer & Pollick, 2008).

In contrast to most previous experiments, the scenes we present to subjects have *not* been pre-constructed to convey particular categories of interaction. Our aim is to show subjects a broad array of agent interactions—from a richer and more general collection of possibilities—in an attempt to allow subjects' minds to impose *their own* structure on the agent space. The way we produce the desired scenes is also novel: We program the agents inhabiting these scenes to behave *autonomously*, which results in often chaotic multi-agent interactions that we cannot predict in advance.

In Exps. 1 and 2, we use multidimensional scaling (MDS) in an attempt to extract the natural clusters and cleavages present within this stimulus space of intentional behavior. Exp. 3 is explicitly designed to help clarify the results of Exps. 1 and 2 by unraveling the "semantics" of the features uncovered by the MDS. Displays were programmed using the breve Simulation Environment (Klein, 2002), an open-source software package freely available at `http://www.spiderland.org`.

Programming Lifelike Automata

In designing and coding the agent behaviors, we aimed to employ a simple programming scheme that would impose minimal structure on the agents' interactions but, nonetheless, would be capable of producing a rich variety of lifelike agent behaviors.[1] We programmed the triangular agents to behave autonomously, each running its own independent program. Inspired by the work of Braitenberg (1984), we aimed to create rule-governed agents which, notwithstanding the simplicity of their programs, yield vivid and lifelike behaviors that give subjects a strong impression of intentions.

Agent design Rather than presenting subjects with prefabricated animations, we populate simulations with autonomous agents and then allow these simulations to run for a predetermined length of time (15 seconds). Each agent starts off in the simulation environment with a randomly-assigned velocity and location. The agent always orients one vertex of its triangular body (that which lay on its axis of symmetry) in the direction of its movement, inducing the impression that the front end is the agent's "head" (see Tremoulet & Feldman, 2000). When an agent either collides with another agent or the edge of the scene, it "bounces off" for one iteration of the simulation.[2]

At each iteration of a simulation, an agent finds the nearest *other* agent within the scene and then accelerates toward or away from it to an extent determined by a set of six parameters contained in its program. The parameters control the direction and magnitude of the agent's acceleration—relative to the nearest other agent—at six respective distances from this other agent: 0-5 "units", 5-10, 10-20, 20-40, 40-70, or $>$ 70. A schematic of these 6 radii around an agent, along with a snapshot of Experiment 1, is shown in Figure 1.

One example agent might approach another agent from afar but then veer away as it gets to a closer radius. Others might consistently accelerate away from another agent. Depending on how this other agent is programmed, their interaction might resemble chasing/fleeing, or one pushing the other, or even one agent circling the other.

We constructed a pool of 12 agents, each with 6 randomized parameters within the programming scheme.

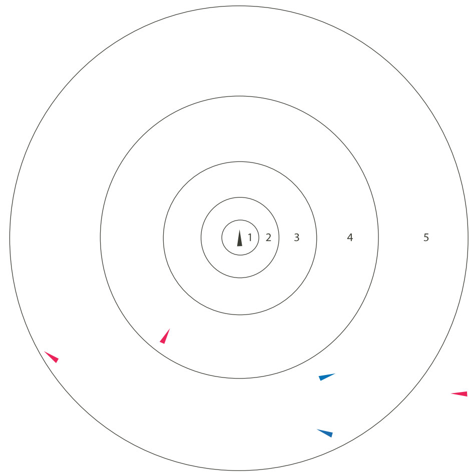

Figure 1: Screenshot from Experiment 1 (colors inverted), with black circles and numbers superimposed onto the scene to help illustrate the programming scheme for the automata. The black automaton in this scene accelerates toward or away from the nearest other agent in the scene. The direction and magnitude of this acceleration depends on the distance to this nearest other agent, with possible distances divided into six zones. Zone #5 seems to be most psychologically relevant.

Experiment 1

Method

Subjects Eight students between the ages of 18 and 24 participated in an approximately one-hour experimental session in exchange for course credit.

Stimuli Scenes were presented to subjects on a 1440 x 900 LED display, on a 15 inch MacBook Pro laptop with a 2.2 GHz dual core processor. The simulation environment itself measured 33.0 x 16.5 cm, and the viewing distance was approximately 45 cm. The programming library employed units that were equivalent to 8.7 units/cm. The triangular agents had bases of 1 unit length and heights of 4 unit length.

Procedure In each 15 s scene, the subject observed 7 agents interacting: 3 red, 3 blue, and 1 white. The reds behaved according to the same parameters as the other reds, the blues according to a different set of parameters, and the lone white according to a third set of parameters. The agents were drawn from a larger 12 agent pool; thus, there were 220 possible triads of these 12 agents.[3] For each scene, one of these 220 triads was selected at random, and each of the three programs in the selected triad was randomly assigned to either red, blue, or white. Each subject saw 220 such scenes, exhausting the possible triads.

Subjects were openly encouraged to construe the triangular agents as animate. At the end of each scene, they were asked "Is the white agent behaving more like a red, or more like a blue?" They answered by clicking on a button in a dialog box.

We constructed a 12 x 12 symmetric distance matrix for each subject, to be fed into the individual differences multidimensional scaling (MDS) algorithm (INDSCAL/ALSCAL; Takane, Young, & Leeuw, 1977). Within this matrix, an agent was assigned a distance of 0 from itself. As two different agents appeared in the same trial of an experimental session 10 times, the distance in this matrix between any two agents was initially set at 11.

[1] It is important to note that the programming scheme we employ here is only one possible choice among many. The design of lifelike agents is a complex and multifaceted problem that extends far beyond the scope of our research. For us, these simple automata are merely tools for aiding an empirical study of the perception of intention.

[2] In Experiment 1, this sometimes resulted in jerky and unnatural-looking behaviors at agent collisions, so in Experiments 2 and 3 we changed collision behavior slightly: agents in these experiments bounced off each other for a full .2 s at some random velocity vector.

[3] Strictly speaking, because the status of the white agent in each trial is special, and, as a result, during a given trial the subject cannot respond that he actually believes the blue and red agents to be most alike, 660 possible arrangements actually exist. Rather than show all 660 possibilities, we randomized the procedure so that no agent type would be more or less likely to be "white" during a trial. Nevertheless, this presents a source of noise in the data, and we altered the procedure in Experiment 2 to address this issue.

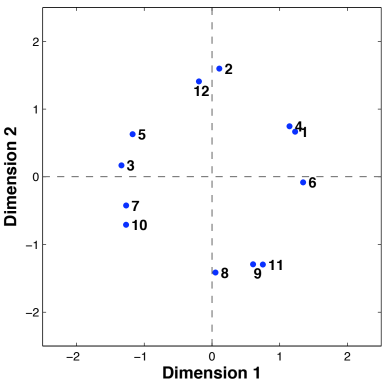

Figure 2: The 2-dimensional MDS solution for the 12 agents, fitting data from Experiment 1.

If the subject chose "red," then the agent whose programming was used for the red agents in this trial was made to be closer together (more similar) in this distance matrix with that of the white agent, and likewise for if the subject chose "blue." That is, the distance between these two agents in the matrix was reduced by 1. Previous studies have used similar methodologies to gauge subject similarity ratings of visual stimuli (e.g., Kahana & Bennett, 1994; Pantelis, van Vugt, Sekuler, Wilson, & Kahana, 2008).

Results and Discussion

We derived a 2-dimensional (2D) MDS solution in order to visualize the space of agents that subjects (on average) perceived (see Figure 2). For this amount of points in the space, the INDSCAL algorithm allows for fits of 2-5 dimensions. Deriving higher-dimension solutions will always result in better fits to the experimental data.[4] However, a higher number of dimensions would be even more difficult to interpret than the 2 condensed dimensions we present, and even a 5-dimensional fit would probably be a condensed version of the true amount of psychologically relevant dimensions in this agent space (which could hypothetically be even higher than the total number of agents in our sample).

A 2D solution allows for the easiest visualization of the inter-agent distances, an important motivation for using the MDS analysis in the first place. If interesting structure emerged only in higher-dimensional fits for these data, this might have justified using these MDS solutions. However, we actually found the clearest and most interesting structure within a 2D fit.

The most striking aspect of the space is its ring-like structure, similar to what one would observe in a 2D MDS plot of the color wheel (see Shepard, 1980). The significance of this ring structure was not immediately clear, in part because MDS dimensions are in general not self-explanatory but rather pull out subjectively primitive parameters. Exp. 3, presented below, was designed to help clarify the nature of the parameter exhibited in this ring.

[4]While we examined a scree plot of the pooled data from Experiments 1 and 2, we do not display it here due to space considerations. This scree plot does not demonstrate a clear "elbow" favoring one particular number of dimensions over another.

Table 1: Correlations ($r[10]$) between programmed parameters (rows) and MDS dimensions (columns). Bold font represents $p <$ .01

	MDS Dimension 1	MDS Dimension 2
Parameter 1	-.070	.384
Parameter 2	-.275	-.074
Parameter 3	.527	.199
Parameter 4	.411	-.375
Parameter 5	**-.801**	.093
Parameter 6	.459	.197

The goal of the present experiments was not, per se, to see how the somewhat arbitrary parameters with which we programmed the agents mapped to subjects' percepts of the agents' behaviors. Rather, we had aimed to infer the structure of the perceptual space itself. Nonetheless, relating these parameters to the MDS dimensions was a useful step in understanding the 2D MDS space.

Subjects' perception of the agents' behaviors arises from some complex interaction of its underlying programming and the chaotic interaction with other agents that arises during each unique simulation. This contributed to there being many individual differences between subjects' results; few subjects' distance matrices showed obvious correlation. However, one of the 6 parameters with which we programmed each agent was indeed strongly correlated with one of the MDS coordinates (see Table 1). This parameter controlled how an agent behaved when the nearest other agent was between 40 and 70 units (4.6 to 8.1 cm) away from it. This finding is addressed further in the Experiment 2 discussion.

Experiment 2

In Exp. 2, we adjusted the basic methodology of Exp. 1 in hopes of reducing the amount of noise in the data. The most significant change was to allow the subject to control one of the agents in each simulation via the mouse. The chance to interact with the simulated agents would, we expected, allow the subject to glean more information about the other agents' behaviors during the short 15-second display time and thus promote stronger impressions of the agents' "personalities" than was possible in Exp. 1.

Method

Subjects Seven students between the ages of 18 and 23 participated in an approximately one-hour experimental session in exchange for course credit.

Stimuli We presented scenes to subjects on an eMac with a 17 inch (16 inches viewable) monitor and a 1152 x 864 display. The monitor refresh rate was 80 Hz and the computer had a 1.25 GHz processor. The simulation environment itself measured 25.4 x 16.5 cm, and the viewing distance was approximately 45 cm.

Exp. 2's scenes were populated with triangular agents of the same size and programmed under the same scheme as in Exp. 1. We used the same pool of 12 agents from Exp. 1, each which had been created with 6 randomized parameters within the programming scheme.

Additionally, the subject controlled one agent with the mouse: a white circular agent 4 units in diameter. The automatic agents reacted to the subject-controlled agent in the same manner as any other triangular agent in the simulation.

Procedure In each 15 s scene, the subject observed 6 agents and controlled 1 agent. 2 agents were red, 2 were green, 2 were blue, and the subject-controlled agent was white. The reds would behave according to the same parameters as the other reds, the greens according to a different set of parameters, and the blues according to a third set of parameters. The agents were drawn from a larger 12 agent pool; thus, there were 220 possible triads of these 12 agent programs. For each scene, one of these 220 triads was selected at random, and then each of the three programs in the selected triad was randomly assigned to either red, green, or blue. Each subject saw 220 such scenes, exhausting the possible triads.

Subjects were openly encouraged to construe the triangular agents as animate, and were instructed that how agents of a certain color behaved during one trial would have nothing to do with how they behaved in subsequent trials. At the end of each scene, they were asked to determine which color of agent behaved *least* like the other two—that is, which was most different: red, green, or blue? They responded by key press, at which point the next trial began.

As in Exp. 1, we constructed a 12 x 12 symmetric distance matrix for each subject, to be fed into the individual differences multi-dimensional scaling (MDS) algorithm. For each trial, the two non-chosen agents in the odd-one-out procedure were made more similar within this distance matrix.

Results and Discussion

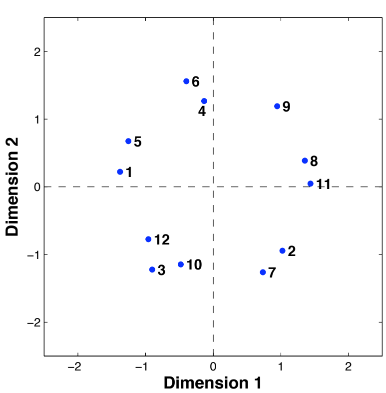

Figure 3: The 2-dimensional MDS solution for the 12 agents, fitting data from Experiment 2.

Table 2: Correlations ($r[10]$) between programmed parameters (rows) and MDS dimensions (columns). Bold font represents $p <$ .05

	MDS Dimension 1	MDS Dimension 2
Parameter 1	-.129	-.147
Parameter 2	-.529	-.050
Parameter 3	.096	.195
Parameter 4	.548	.200
Parameter 5	-.310	**-.619**
Parameter 6	.249	.233

Once again, we derived a 2D MDS solution in order to visualize the space of agents that subjects (on average) perceived, and we once again observed a ring-like structure in the space (Fig. 3).

The MDS solutions for the two experiments—processed representations of subjects' raw similarity matrices—were correlated with each other. Dim. 1 of Experiment 1's MDS was strongly correlated with Dim. 2 of Experiment 2's MDS [$r(10) = .713, p < .01$]. Dim. 2 of Experiment 1's MDS was weakly (and negatively) correlated with Dim. 1 of Experiment 2's MDS [$r(10) = -.551, p = .063$]. (The direction of these correlations is arbitrary and unimportant, but helpful in relating the 2D MDS spaces presented in Figures 2 and 3.) These correlations provide some assurance of the robustness and psychological reality of the subjective mental spaces that we have uncovered.

As shown in Table 2, Dim. 2 in Experiment 2 correlated significantly with parameter 5 of the agents' programming. This is consistent with the results of Experiment 1, where Dim. 1 had been correlated with this same parameter. Apparently, how an automaton reacted to (i.e. accelerated toward or away from the direction of) the nearest other agent in the simulation when that agent was 40-70 units away (10 to 17.5 times the length of an agent) was a psychologically important variable.

We wondered if the prominence of this parameter in subjects' judgments was actually an artifact of the frequency with which interactions at this distance actually occurred in the displays. But the data do not bear this out. Because the entire displays were recorded (10 frames/second), we could assess the proportion of the time the inter-agent distance between any automaton and its nearest other agent was within each of the six intervals corresponding to the six underlying programmed parameters. The two most common distances between an automaton and its nearest other agent during a simulation were 0-5 units (0-1.25 agent lengths) and 20-40 units (5-10 agent lengths). 40-70 units (10-17.5 agent lengths) was only the fourth most common inter-agent distance. The pivotal role of this inter-agent distance is not an artifact, but rather reflects a genuine cognitive focus on behavioral interactions at this distance.

Experiment 3

The results of the first two experiments were qualitatively similar, and we therefore choose to pool data from all 15 subjects for the following analysis and discussion. The 2D MDS solution for these pooled subjects reveals an even cleaner ring structure (see Figure 4). But what does it mean as we travel around this ring?

In the combined MDS, Dim. 1 is connected to how an agent behaves when the closest other agent is between 10-17.5 agent lengths away (i.e. programmed parameter #5). Agents low on Dim. 1 all tend to accelerate away from the nearest other agent; agents high on Dim. 1 tend to accelerate toward the nearest other agent. The meaning behind Dim. 2 is less straightforward. While this dimension is clearly not independent from Dim. 1, it is uncorrelated with any of the programmed agent parameters. Hence we turn to further psychophysics to provide evidence about its meaning.

We hypothesize that a potential "friendly" versus "hostile" dimension emerges from the interaction of these two MDS dimensions. This hypothesized dimension would be neither orthogonal nor redundant with whether an agent accelerates toward or away from another agent at a certain distance—say, the distance with which programmed parameter #5 is concerned. When an agent moves in the direction of another, it may, for instance, appear to be aggressive or merely curious.

Method

Subjects Seven students between the ages of 18 and 24 participated in an approximately half-hour session in exchange for course credit.

Stimuli and Procedure We presented scenes to subjects under the same viewing conditions as Exp. 1. We again populated the simulations with the pool of 12 agents employed in Exps. 1 and 2. During each trial, the subject watched 7 agents interacting for 15 seconds. Six of the agents were colored red and behaved under programs randomly selected from the pool of 12. The seventh, critical agent was colored blue, and the subject was instructed to attend to it. At the end of each trial, the subject was asked, "On a scale of 1-5, 1 being most hostile, and 5 being most friendly, how do you rate the blue agent?" The subject indicated his response on the keyboard. Each of the 12 agents in the pool was assigned the blue color for 8 of the session's trials, for a total of 96 trials presented in random order.

Results We first normalized each subject's responses, then calculated each subject's mean normalized response for each of the 12 agents observed over the experimental session. Then, averaging across subjects, we were able to get a sense of how friendly versus hostile subjects perceived each of the 12 autonomous agents. Figure 4 shows, on a gradient from red to green, what these perceptions were. The most hostile agents seem to be those which were high on MDS Dim. 1 and low on Dim. 2, while the friendlier agents tended to be low on Dim. 1 and high on Dim. 2. Agents low on both dimensions were quite neutral. Fig. 4 shows the space in a rotated coordinate frame so that the horizontal dimension optimally reflects the friendliness vs. hostility dimension. (All of the inter-agent distances and relationships have been preserved; only the "ring" has been rotated.) In the rotated space the projection of each agent's position onto the horizontal (i.e. the cosine of its angle relative to the horizontal) reflects its position along the friendly/hostile dimension. We regressed the subjects' mean friendliness rating against this variable and found a close fit ($r(10) = -.768, p < .01$, Figure 5). These data corroborate our hypothesis that the ring variable essentially reflects the degree of perceived friendliness or hostility each agent exhibited.

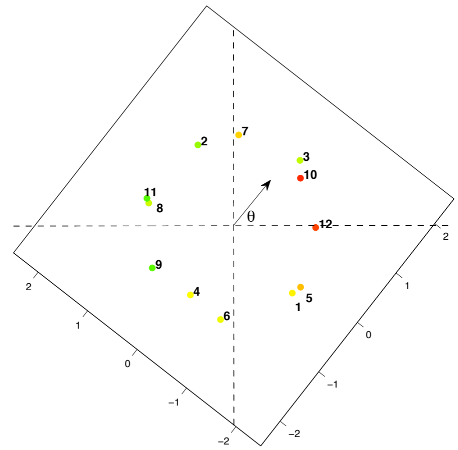

Figure 4: Pooled 2-dimensional MDS from Experiments 1 and 2. The 2D space is rotated about the origin so that the cosine of the agent angle (relative to the horizontal) best predicts how subjects rated the agents along the "hostility" versus "friendliness" dimension. "Hostility" versus "friendliness" is represented for each agent here with a color gradient from red (most hostile) to green (most friendly), with yellow being neutral.

General Discussion and Conclusions

These experiments were designed to probe the underlying structure of the agent space perceived by subjects as they watched autonomously programmed agents interacting in a dynamic scene. In Experiments 1 and 2, the MDS approach succeeded in revealing certain aspects of this perceptual space: a ring-like structure, which—in Experiment 3—we attempted to connect to a dimension of perceived hostility versus friendliness in the agents. One of the low-level parameters controlling the behaviors of the agents contributed to this more abstract percept: that which controlled inter-agent reactive behavior at one critical distance. We conclude that this reflected one perceptually critical inter-agent zone upon which subjects based their interpretations of the agents' intentional behavior.

From the results of Experiment 3, we further conclude that "hostility" versus "friendliness," or something akin to this di-

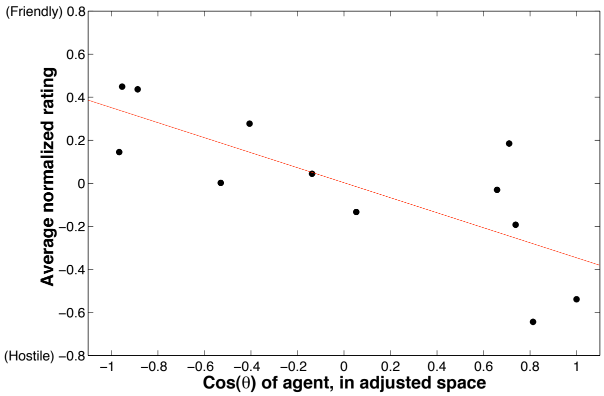

Figure 5: Cosine of the agent's angle in MDS space (see Fig. 4), plotted against how subjects, on average, rated them (from hostile to friendly). Best linear fit is drawn in red.

chotomy, appears to produce an especially salient partition in subjects' perceptual space. In other words, after first surmising that an object in the world has intentions (i.e., is animate), a next step for the cognitive machinery might be an attempt to guess whether these intentions are bad or good.

This work represents one step in what we hope is a fruitful new direction. Programming agents automonously, and asking how subjects' interpretations of these agents' behavior relates to the actual programs they are carrying out, allows one to pursue a true "psychophysics of intention," in which we explore the relationship between the perceived intention and the "actual" intention present in the agent's autonomous program. In future experiments, employing displays of potentially far more complex behavioral interactions, we hope to uncover correspondingly more complex structures in the intentionality percept.

Acknowledgments

This research was supported by the National Institutes of Health (NIH EY15888), the NSF IGERT program in Perceptual Science (NSF DGE 0549115), and a grant from the Hellenic University Club of New York.

References

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

Baker, C. L., Tenenbaum, J. B., & Saxe, R. R. (2006). Bayesian models of human action understanding. *Advances in Neural Information Processing 18*.

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

Braitenberg, V. (1984). *Vehicles*. Cambridge, MA: MIT Press.

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

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. 151–184). New York, NY: Oxford University Press.

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

Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *American Journal of Psychology*, *57*, 242–259.

Johnson, S. C. (2000). The recognition of mentalisic agents in infancy. *Trends in Cognitive Development*, *16*, 637–656.

Kahana, M. J., & Bennett, P. J. (1994). Classification and perceived similarity of compound gratings that differ in relative spatial phase. *Perception & Psychophysics*, *55*, 642-656.

Keil, F. C. (1994). The birth and nurturance of concepts by domains: The origins of concepts of living things. In L. Hirschfeld & S. A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture.* New York: Cambridge University Press.

Klein, J. (2002). breve: a 3d simulation environment for the simulation of decentralized systems and artificial life. *Proceedings of Artificial Life VIII, the 8th International Conference on the Simulation and Synthesis of Living Systems.*.

McAleer, P., & Pollick, F. E. (2008). Understanding intention from minimal displays of human activity. *Behavior Research Methods*, *40*(3), 830–839.

Pantelis, P. C., van Vugt, M. K., Sekuler, R., Wilson, H. R., & Kahana, M. J. (2008). Why are some people's names easier to learn than others? the effects of face similarity on memory for face-name associations. *Memory & Cognition*, *36*(6), 1182–1195.

Shepard, R. N. (1980). Multidimensional scaling, tree-fitting, and clustering. *Science*, *210*(4468), 390–398.

Sloman, S. A., Fernbach, P. M., & Ewing, S. (2009). Causal models: the representational infrastructure for moral judgment. In D. Bartels, C. W. Bauman, L. J. Skitka, & D. Medin (Eds.), *Moral judgment and decision making: The psychology of learning and motivation* (Vol. 50). San Diego, CA: Elsevier.

Takane, Y., Young, F. W., & Leeuw, J. de. (1977). Nonmetric individual differences multidimensional scaling: an alternating least squares method with optimal scaling features. *Psychometrika*, *42*(1), 7-67.

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.