Matching biological motion at extreme distances

Ian M. Thornton

Department of Cognitive Science, University of Malta, Msida, Malta

Zac Wootton

Psychology Department, Swansea University, Swansea, UK

Pille Pedmanson

Psychology Department, Swansea University, Swansea, UK

The goal of the current paper was to determine the maximum distance at which an actor could be placed so that an observer would still be able to interpret their behavior. Although we know a great deal about the limits of action perception, particularly through studies of biological motion processing, this question of distance has not been previously documented. We began by reviewing the sizes of point-light figures used in 100 previous studies of biological motion. We found that with an average figure height of 6.6° visual angle, actors were effectively 15 m from the observer, assuming average physical height of 1.75 m. No previous studies had explicitly examined extreme distances. Here, we introduce a new matching task in which we systematically varied the apparent distance of point-light figures relative to a fixed viewing position by manipulating size. Our results suggest that a variety of human actions could potentially be interpreted up to 1000 m away, a distance at which a human figure would subtend only 0.1° visual angle in height. Dynamic figures could be interpreted at further distances than static figures (Experiment 1), and upright figures were similarly processed more efficiently than inverted figures (Experiment 2). We discuss these findings in the context of the processing mechanisms thought to underlie action perception and suggest that the ability to match actions at extreme distance is another example of the robust nature of biological motion processing.

Introduction

Key to the survival of any species is an ability to react appropriately to environmental threat (Curio, Ernst, & Vieth, 1978; Darwin, 1872; Griffin, Blumstein, & Evans, 2000). For social animals like humans, a highly salient visual cue to impending threat often involves the action of others. Clearly, early recognition of whether an observed behavior warrants approach or avoidance may often be desirable. In the current paper, we use the point-light technique pioneered by Gunnar Johansson (1973) to explore the extreme limits of distance at which useful information might be extracted from human motion.

Point-light figures, in which a body is represented by light sources attached to the major joints, are known to give rise to very clear and compelling impressions of underlying action (Dittrich, 1993; Vanrie & Verfaillie, 2004). Such actions are thought to be processed by a range of mechanisms at multiple levels of the visual system (Bertenthal & Pinto, 1994; Bülthoff, Bülthoff, & Sinha, 1998; Mather, Radford, & West, 1992; Neri, Morrone, & Burr, 1998; Pavlova & Sokolov, 2000; Thornton, Pinto, & Shiffrar, 1998; Troje & Westhoff, 2006) via a distributed network of action-related brain areas (Downing & Peelen, 2011; Giese & Poggio, 2003; Grossman & Blake, 2002; Saygin, 2007; Thompson & Parasuraman, 2012).

Over the years, many properties of this stimulus class have been systematically manipulated in the laboratory in attempts to better understand the nature and the limits of biological motion processing (see Blake & Shiffrar, 2007; Thornton, 2006, for reviews). Interestingly, one very salient, ecologically valid parameter—the apparent distance of the actor from the observer—appears to have escaped systematic examination. Thus, we know very little about the range of real-world distances over which human observers might have access to the signals generated by the actions of others. To address this gap, we examined the speed and accuracy with which actions could be matched across a wide range of apparent distances, spanning 100 to 1000 m.

As we note in more detail below, only a handful of previous studies have made explicit reference to distance as a potentially interesting variable in the context of biological motion (Balk, Tyrrell, Brooks, & Carpenter, 2008; Jokisch & Troje, 2003; Legault, Troje, & Faubert, 2012; Luoma & Penttinen, 1998; Owens, Antonoff, & Francis, 1994; Tyrrell et al., 2009; Wood, Tyrrell, & Carberry, 2005). Similarly, although figure size—an important real world correlate of distance—has been varied, both within and across experiments, this has usually been done either to control for eccentricity effects or to accommodate particular equipment or display characteristics (e.g., Gurnsey, Roddy, Ouhnana, & Troje, 2008; Gurnsey, Roddy, & Troje, 2010; Ikeda, Blake, & Watanabe, 2005). To our knowledge, the current paper is the first attempt to explicitly measure how performance varies as an actor is moved away into the far distance.

It is unclear why distance has been effectively ignored in this area of research. Possibly it has just been assumed that the mechanisms responsible for biological motion processing are scale invariant. If this were the case, then probing for performance differences as figures move into the distance would be of little interest. Another, probably more likely, scenario is that the sparse abstract nature of the point-light displays have led researchers to miss the connection that maps the projected size of a figure on the screen to a real human actor located at some fixed distance in the outside world.

A new matching task to explore apparent distance

In order to explore the question of distance, we developed the new concurrent matching task, illustrated in Figure 1 and Movie 1. In this task, two flanking figures, to the left and right of the screen, always performed two different actions. In the current experiments, the actions were randomly selected on each trial from a subset of those available in Vanrie and Verfaillie's (2004) motion-capture database. The 12 actions we chose were: chop, jump, mow, paint, pump, saw, shoot, spade, sweep, tap, walk, and wave. The display and viewing conditions were calibrated so that the size of the flanking figures on the screen corresponded to the visual angle that would be subtended by a 1.75-m tall actor positioned at either 15 m (Experiment 1a) or 30 m (Experiments 1b & 2) away from the observer. The flanking figures were always positioned at the vertical midline of the screen and were separated to the right and left of the screen center by 10° (Experiment 1a) or 5° (Experiments 1b & 2).

On each trial, a central target figure performed the same action as either the left or the right flanker with equal probability. The observer's task was simply to match the action by pressing a left or right key on a standard keyboard. Crucially, the apparent distance of the central target was manipulated from trial to trial within the range 100–1000 m by appropriate size scaling. Specifically, the target figure was scaled in accordance with the curves shown in Figure 2, to mimic the reduction of visual angle that would occur when a real human figure receded in depth. As the target figure remained centered on the vertical midline, the feet also appeared elevated relative to the flanking figures.

The starting frame and the in-depth orientation (i.e., 360° rotation around the vertical axis) of all three figures was independently randomized on each trial to minimize low-level similarity. The individual points on each figure consisted of a single pixel. The figures were orthographically projected on a black featureless background, so that the most salient cue to apparent distance was the global change in size of the central figure relative to the flanking figures.

The use of a concurrent matching task is appealing for a number of reasons. Most importantly in the current context, the presence of standard flanking figures allowed us to convey a sense of relative distance between target and flankers. The interpretation of the display in which the central target was being moved in depth rather than simply changing in size was easy to explain and appeared to be intuitively grasped by all of our observers.

More generally, matching also makes it possible to include a wide range of actions and viewing conditions while keeping the task and dependent variables constant. For example, the left/right facing or forward/backward articulation decisions often made to walking actions do not generalize well to other forms of behavior. Matching makes it possible to separate and standardize the task-relevant decision from the nature of the action being performed. In the General discussion we return to further consider both the merits and limitations of using matching to measure biological motion performance.

Cues to distance

From the above description it should be clear that there are two main cues to apparent distance in the current paper. The primary cue, and the one that will be the focus of our analysis, are changes to the visual size of the target figure. Details of these changes are provided later in this section. Additionally, as the scaled target figure remained vertically centered on the screen, its apparent elevation also changed. We chose to do this for two reasons. First, this helps to remove low-level alignment cues that may have aided local matching with the feet of the flanking figures. Second,
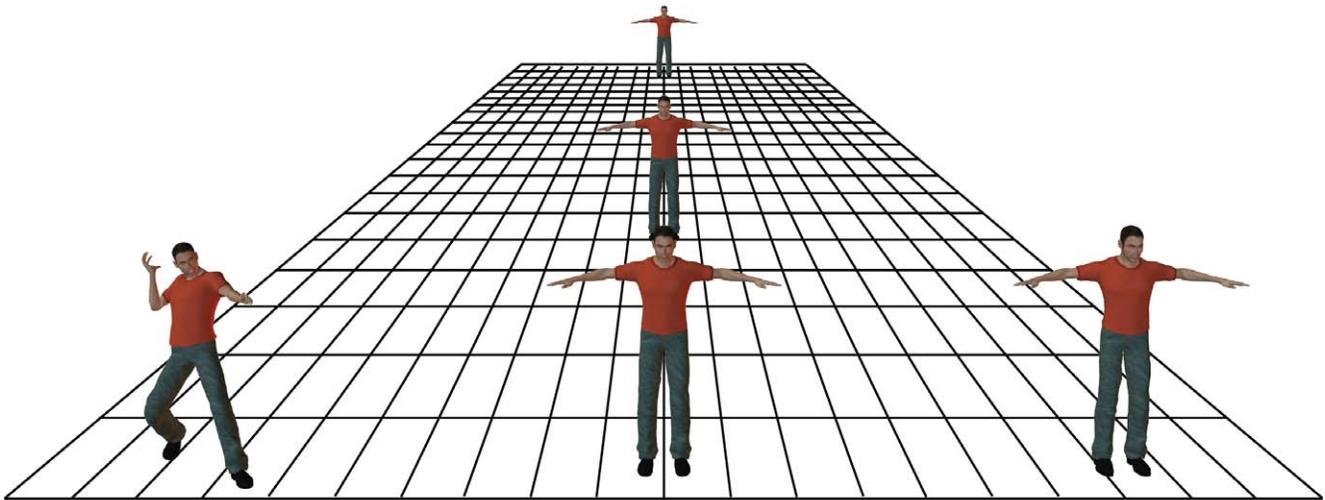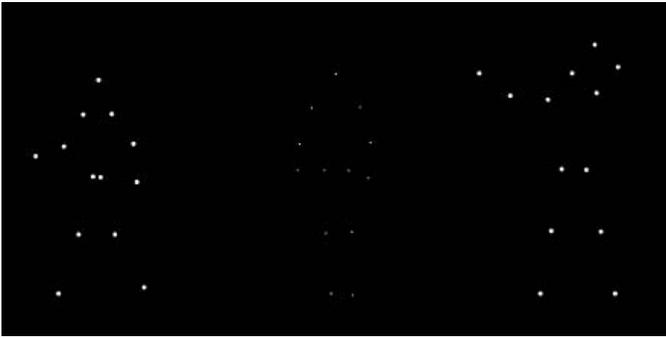
Figure 1. (A) Schematic illustration of the matching scenario used in the current work. Two flanking figures were held at a constant apparent distance and always performed two different actions. A central target figure that matched either the left or right flanker appeared at a series of simulated distances within the range 100–1000 m. On each trial the in-depth orientation of each figure, and the exact starting or static pose, were independently randomized to avoid low-level synchronization or image matching. The task for the observer was to indicate with a key press, whether the target matched the left or right flanker. (B) In the experimental display point-light figures were used. The actions were a subset of those from the Vanrie and Verfaillie (2004) database. To provide a range of matching difficulties, we selected actions such that some pairs were roughly matched in terms of posture and degree of movement, while others would be quite different. The actions used, in alphabetical order, were: chop, jump, mow, paint, pump, saw, shoot, spade, sweep, tap, walk, and wave. See text for further details.

the change in elevation provides a consistent perspective cue that reduces the likelihood of our displays being interpreted as containing a target that simply shrinks or grows in place.

While other cues to distance and depth could have been included, there are a number of reasons to specifically concentrate on size. For example, while binocular cues are relatively easy to introduce in the context of biological motion (e.g., Bülthoff et al., 1998;

Jackson & Blake, 2010), at the real-world distances we were interested in, (i.e., 100 m and beyond) this cue is unlikely to play a substantial role (although see Palmisano, Gillam, Govan, Allison, & Harris, 2010). Additionally, previous research suggests that disparity cues would almost certainly be overridden by the two-dimensional form/motion cues in the figures (Bülthoff et al., 1998).

Movie 1. An illustration of the display and trial structure of the matching task used throughout this paper. The goal is to indicate which flanker performs the same action as the central target. The size of the central figure is scaled down in each subsequent trial, following the curve shown in Figure 2A. At a viewing distance of 60 cm these central figures subtend a visual angle consistent with an actor standing approximately, 15, 25, 50, 100, 200, 300, 400, and 500 m away, respectively. Note that for ease of viewing, the dots of the flanking figures have been scaled up. During experimental trials, these would also have been drawn as single pixels. The flanking figures have also been brought closer to the target and were not constrained by the edge of the screen during the experiment. The correct sequence of target actions is as follows: walk, saw, spade, jump, spade, walk, chop, and wave.

The inclusion of ground plane textures or additional scene structure may well have improved the overall sense of distance in our displays. However, this would have made it impossible to compare our findings to previous biological motion studies in which such visual clutter is almost always absent. Also, our use of single pixels to represent the joint dots resulted in global figures that were of fairly low contrast. We chose single pixel dots to avoid local size scaling cues to distance, cues that would have been difficult to systematically manipulate over our range of interest. Our concern, then, was that the addition of further scene content might easily overwhelm or mask the biological motion figures.

Perhaps the most compelling reason to focus on size is that it is almost certainly the most reliable and familiar cue to distance in our everyday experience. Figure 2 shows the relationship between visual size and distance for a human adult figure assumed to be 1.75 m tall. In Panel A, we show the full range of size changes that would occur from looking at a person on opposite sides of the same room (e.g., 10 m) to seeing someone in the far distance (1000 m). In Panel B, we have zoomed in on the extreme distances of particular interest in the current work.

As a consequence of this relationship, we are constantly exposed to changes in the apparent visual size of those around us directly caused by variations in
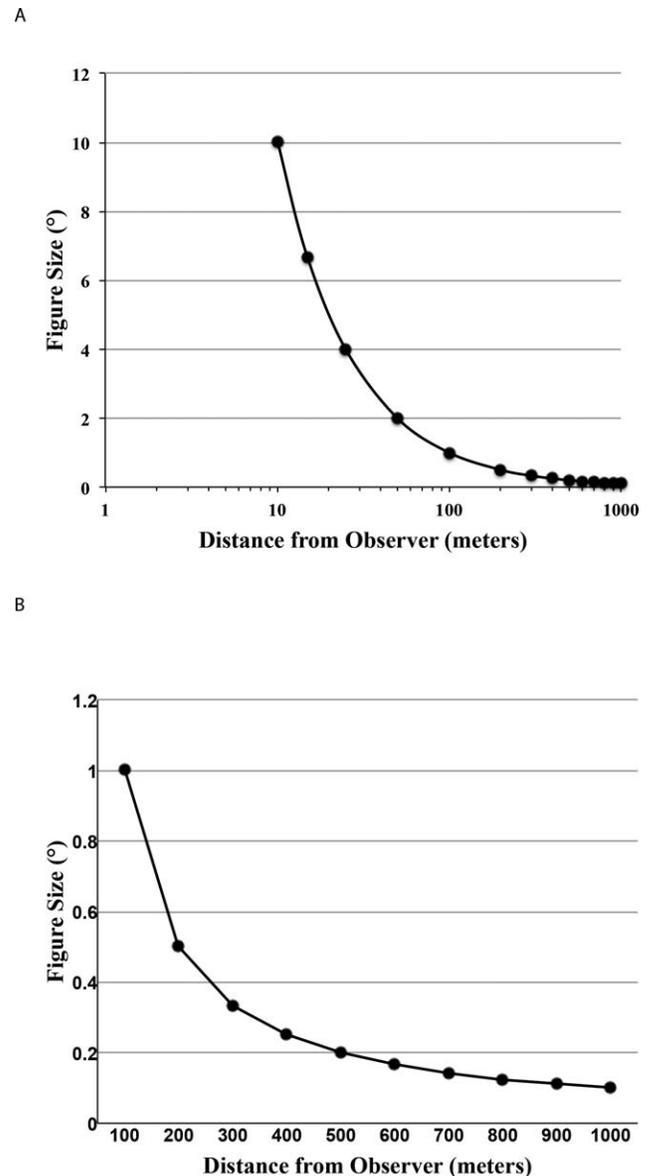

Figure 2. (A) A log linear plot of the visual angle subtended by a human figure with a physical height of 1.75 m seen over the range of distances 10–1000 m. (B) A linear plot of the size variation in the range of interest in the current study, 100–1000 m. For reference, the extreme values of 500 and 1000 m subtend 0.20° and 0.10° visual angle, respectively. See Table 1 for full details.

distance. We may seldom be aware of such variation, at least in part because of size constancy, the mechanism, that, when violated, makes the Ames Room so powerful (Ames, 1952). Nevertheless, as a person moves away from us, the size of the image they project to our eye gets systematically smaller. Glancing down from a window on to a busy street or watching a sporting event in a stadium, for example, immediately presents figures that vary in size quite dramatically.

Figure 3. Stadiometric range finding. This image shows an augmented reality view of Strait Street in Valletta, Malta, taken with the orienteering iPad app SpyGlass (Happymagenta, Ltd.). The scale on the left of the image replicates the size/distance curve shown in Figure 2. By placing a human figure within the scale, one can estimate the approximate distance. Here, the figure at the top of the stairs appears to be a little less than 100 m from the viewer. These scales are used in a variety of applications, perhaps most commonly in telescopic gun sights.

As an aside, we note that outside of academia, the reliability of human size as a cue to distance has long been exploited in Stadiometric range finding applications, particularly in telescopic gun sights. Here, the feet of a human figure are aligned with the base of a scale drawn within the viewfinder—a scale that would essentially replicate the curve shown in Figure 2—and the height of the figure is then used to establish distance. Figure 3 shows a view through a nonlethal rangefinder within the orienteering app SpyGlass (Happymagenta, Ltd.). The stadiometric scale can be seen at approximately 9:00 in this image. The human figure at the top of the stairs is roughly 100 m away from the viewer.

Of course, the use of this single cue does make it possible to interpret our matching displays as containing a figure that physically shrinks or grows in size, rather than an actor of standard height changing position. Similarly, the displays could result from the

presence of children or adults of various heights appearing at a constant distance. As already mentioned, changes in elevation were included to make these interpretations less likely, as such targets would also have to be levitating. Therefore, while both of these interpretations are possible, the explicit instructions that were presented to observers—that displays contained a figure of a standard height being moved closer and further away—is both consistent with physical size/distance relationships, and probably most in line with the majority of their past experience.

Previous studies of distance and size

Within the mainstream biological motion literature, we could find only two previous laboratory studies

that made explicit reference to distance, one in relation to speed of locomotion (Jokisch & Troje, 2003) and the other dealing with perceptual deficits in older adults when processing figures in near space (Legault, Troje, & Faubert, 2012). Neither of these considered the current question by looking at extreme distance. Thus, it appears that very few out of the several hundred laboratory studies of biological motion have considered that the figure displayed on the screen might directly translate to a real actor at some fixed distance.

Outside of the laboratory, there has been much applied work aimed at exploiting sensitivity to biological motion in order to improve road safety (e.g., Balk et al., 2008; Luoma & Penttinen, 1998; Owens et al., 1994; Tyrrell et al., 2009; Wood et al., 2005). In a typical study, pedestrians might be placed along the edge of a closed road circuit wearing different types of reflective clothing. During real or simulated driving maneuvers, participants are then asked to detect the presence of pedestrians. Clothing that highlights the motion, rather than just human form, has been shown to dramatically improve the distance at which the presence of a human target can be detected (Wood et al., 2005). Such studies have considered quite extreme distances (e.g., up to 250 m). However, the focus on simple detection rather than interpretation of action limits their applicability in the current context.

Several laboratory studies have explicitly manipulated the visual size of the human figure within the same experiment, although this has typically been done to control for or explore eccentricity effects (Gurnsey et al., 2008; Gurnsey et al., 2010; Ikeda et al., 2005; Thornton & Vuong, 2004). More generally, between studies, there has been quite some variation in the size of figures used.

In Figure 4, we've plotted the frequency with which various figure sizes have been used in 100 previous studies of biological motion. There are several points to note. First, there is clearly a good deal of variation, with the majority of studies having employed figures in the range 2°–12°. Assuming a standard figure height of 1.75 m, this equates to a real-world actor who is located at a distance of between 8 and 50 m, with the average actor (6.6°) approximately 15 m away (Figure 2A). Second, the distribution is clearly positively skewed. As the tail extends much further to the right than the left we are likely to know much more about the processing of closer figures than we do about more distant figures. Third, we want to emphasize again that with the exception of the papers already mentioned, the choice of figure size in these studies appears to have been dictated by specific equipment or display characteristics, rather than any theoretical or ecological considerations relating to distance.
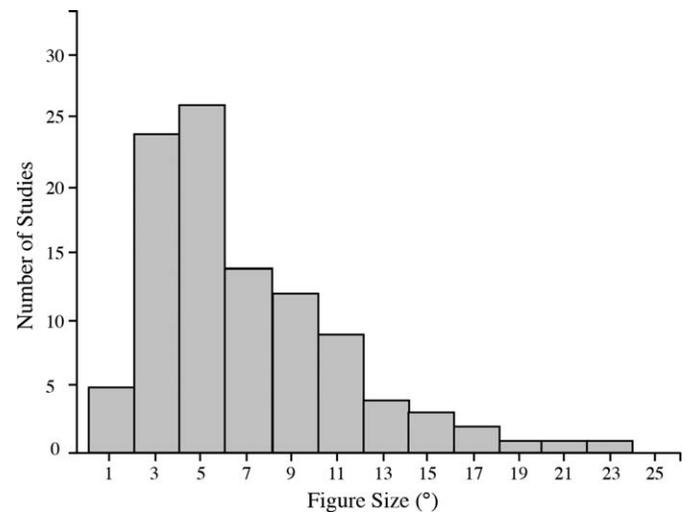

Figure 4. The distribution of figure heights sampled from 100 previous laboratory studies of biological motion. There is clearly quite a good deal of variation, with the majority of studies having employed figures in the range 2°–12°. Assuming a standard height of 1.75 m, this equates to a real-world actor who is somewhere between 8 and 50 m away. The average figure height was 6.6°, which is approximately 15 m away.

The current questions

In the current paper we used point-light stimuli to examine the speed and accuracy of action matching decisions across a wide range of apparent distances. Our goal was to establish the limits at which action interpretation would break down. We conducted two experiments.

In Experiment 1, we used upright actions to examine matching at a range of distances between 100–1000 m. In addition to standard dynamic point-light figures, we also interleaved trials in which the target and flankers were single static snapshots. The purpose of including static trials was to establish the contribution of the static posture to overall matching performance (Beintema & Lappe, 2002; Casile & Giese, 2005; Giese & Poggio, 2003; Lange & Lappe, 2006; Thirkettle, Benton, & Scott-Samuel, 2009; Thirkettle, Scott-Samuel, & Benton, 2010).

In Experiment 2, we examined approximately the same range of distances (200–1000 m), but now compared upright and inverted trials. In the inverted trials, both the target and the flanker were rotated in the picture plane by 180° (Barclay, Cutting, & Kozlowski, 1978; Pavlova & Sokolov, 2000; Sumi, 1984). The purpose of including this manipulation was to gauge the contribution of local and global processes to overall matching performance (Bertenthal & Pinto, 1994; Chang & Troje, 2009; Pavlova & Sokolov, 2000; Troje & Westhoff, 2006). In Experiment 2, all trials contained dynamic figures.

Apparent distance (meters)	Retinal size (degrees of visual angle)	Screen distance (cm)	On screen size (cm)	On screen size (pixels)
15	6.68	60	7.02	270
100	1.00	60	1.05	40
200	0.50	60	0.53	20
300	0.33	60	0.35	14
400	0.25	60	0.26	11
500	0.20	60	0.21	9

Table 1a. Simulated distance and screen height for Experiment 1a.

Apparent distance (meters)	Retinal size (degrees of visual angle)	Screen distance (cm)	On screen size (cm)	On screen size (pixels)
30	3.34	120	7.02	270
200	0.50	120	1.05	40
400	0.25	120	0.53	20
600	0.16	120	0.35	14
800	0.12	120	0.26	11
1000	0.10	120	0.21	9

Table 1b. Simulated distance and screen height for Experiments 1b and 2. *Notes:* Actor height assumed to be 1.75 m. Pixel size 0.026 cm.

Experiment 1

Here we report the results of two separate studies, that we label Experiments 1a and 1b. Our initial estimate of the maximum distance at which matching performance would fall to chance was 500 m, based on the size/distance relationships shown in Figure 2 and Table 1. Thus in Experiment 1a, we explored the range of distances between 100–500 m. However, it quickly became apparent that performance would remain well above chance beyond this estimated maximum distance. We therefore increased the range to include a maximum of 1000 m in Experiment 1b. We report data from both studies in order to provide a full picture of performance between 100–1000 m.

Participants

Eight observers (two female, six male) took part in Experiment 1a and 12 observers (six female, six male) in Experiment 1b. All participants had normal or corrected to normal vision. We increased the number of participants between Experiments 1a and 1b to allow us to balance and test for possible gender effects. All observers gave written informed consent and were told that the purpose of the experiment was to assess the perception of action at a distance. The study was reviewed by and received approval from the appropriate departmental Ethics Committee at Swansea University and therefore conformed to the ethical guidelines set out by the Declaration of Helsinki for testing human participants.

Stimuli

Stimuli consisted of either three static or three dynamic point-light figures, randomly interleaved across trials. Each dot was a single pixel drawn in white on a black background. These dots subtended 0.025° in Experiment 1a and 0.012° in Experiment 1b. The stimuli were orthographically projected and the only cues to distance were the overall height of the figure and its vertical elevation. All figures were aligned so that their centers were at the vertical midline of the screen. The flankers were offset 10° (Experiment 1a) or 5° (Experiment 1b) to the left and right of the horizontal center of the screen.

The actions used were a subset of those from the Vanrie and Verfaillie (2004) database. To provide a range of matching difficulties, we selected actions with both high and low levels of limb movement. We only chose actions with a vertical orientation to avoid simple matching, for example when one flanker was standing and the other sitting. Similarly, we avoided actions that were not periodic to avoid large motion transients that could provide additional matching cues when an animation was looped. The actions used, in alphabetical order, were: chop, jump, mow, paint, pump, saw, shoot, spade, sweep, tap, walk, and wave. See Vanrie and Verfaillie (2004) for detailed descriptions of the actions and the format of the motion capture files.
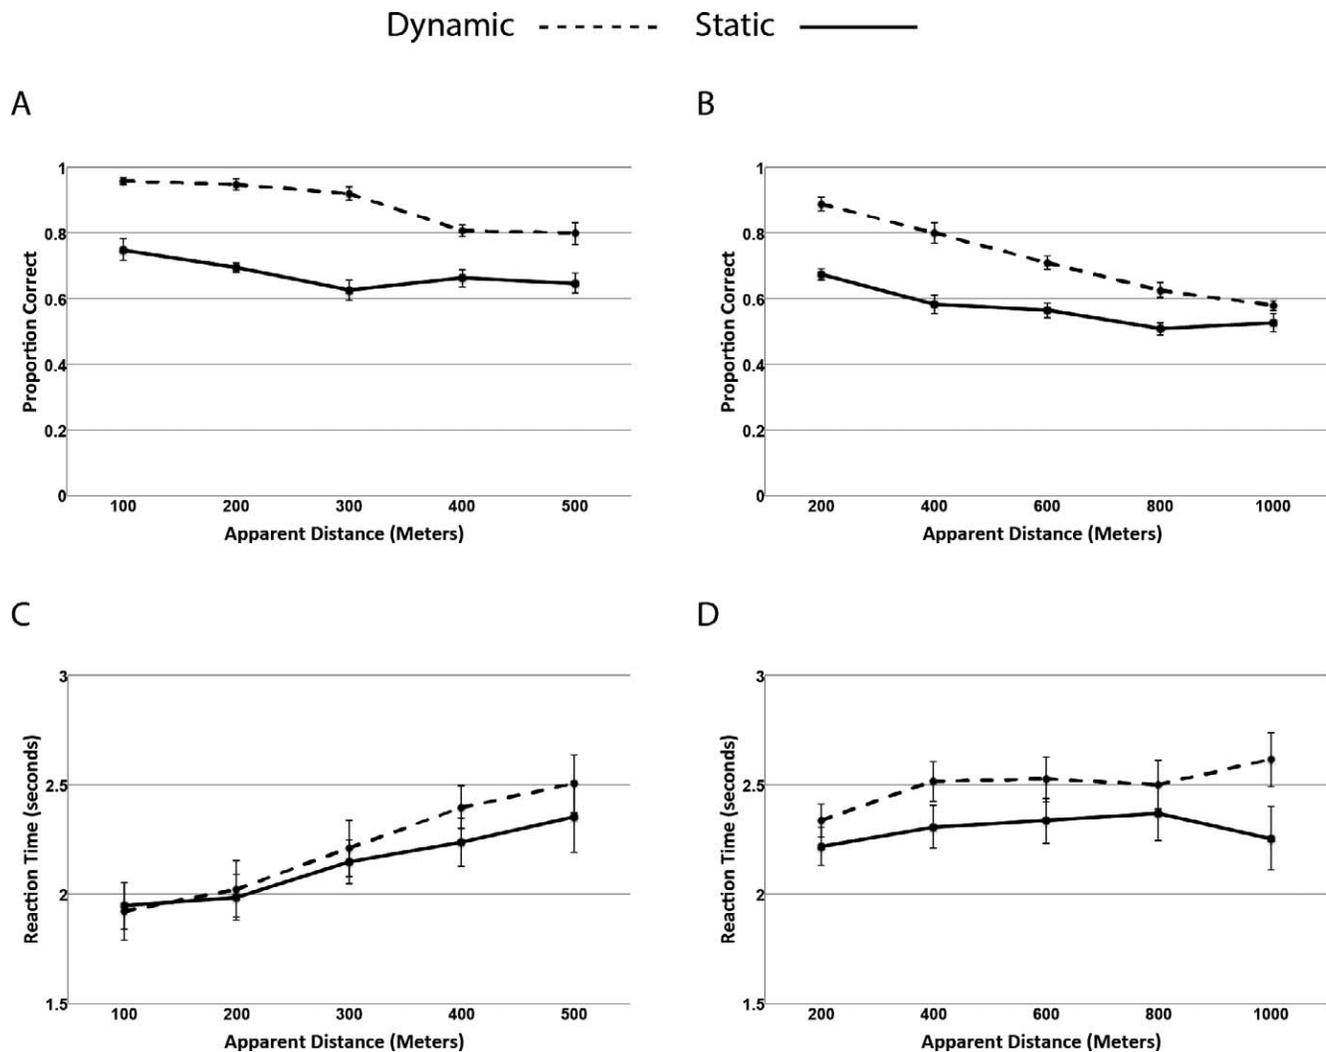
Figure 5. Performance in Experiments 1a and 1b in terms of both accuracy (Panels A–B) and speed (Panels C–D).

Actions were randomly selected on a trial-by-trial basis. The rotation around the vertical axis and the starting frame chosen for both static and dynamic stimuli was randomized independently for each figure in the display on a trial-by-trial basis. Thus each figure could begin to move from any point within the action cycle and face in any direction sampled from the full 360° orientation range. This was done to minimize feature-level overlap between targets and flankers. Custom written code was developed in Matlab, using the Psychophysics Toolbox extensions (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007; Pelli, 1997) and used to render the figures on a standard Macintosh Cinema Display (43 × 27 cm) with a resolution of 1680 × 1050 pixels. Animations were displayed at approximately 30 frames/s, and all three actions continued to play until a response was made. In Experiment 1a, viewing distance was fixed, via a chin rest, at 60 cm. In Experiment 1b,

the same setup was used to fix the viewing distance at 120 cm.

Figures were scaled in size following the curves shown in Figure 2. In Experiment 1a, the flankers were kept at a fixed distance of 15 m, based on the average figure size from our review of past literature (see Figure 4). Table 1a gives the details of visual angle and screen dimensions for the target distances between 100 and 500 m. In Experiment 1b, screen size was not altered relative to Experiment 1a, but screen-to-eye distance was used to double the apparent real-world range. Table 1b gives these parameters in detail.

Task and analysis

In Experiment 1a, each observer completed 10 blocks of 60 trials in a 2 (Static/Dynamic) × 2 (Target Left/Target Right) × 5 (Distance) × 3 (Repetition)

design, giving a total of 600 trials. All factors were randomized within a block of trials. In Experiment 1b, we used the same design but reduced the number of blocks to six, giving 360 trials. Pilot testing suggested stable estimates could be achieved with a smaller number of blocks and some observers in Experiment 1a had reported mild levels of fatigue in the last four blocks of trials.

To familiarize participants with the task and the actions, a practice block of 60 trials was performed in which the target always appeared at the same distance as the flankers. These collinear trials were also present as fillers during the experimental blocks but were not analyzed in the main experimental design.

Initial tests indicated that target side led to no significant main effects or interactions, so this factor was omitted from the main analysis. Data from Experiment 1a were thus analyzed using a 2 (Motion: Dynamic/Static) × 5 (Distance) repeated measures analysis of variance (ANOVA). Data from Experiment 1b used a 2 (Gender) × 2 (Motion: Dynamic/Static) × 5 (Distance) mixed ANOVA. Accuracy and reaction time (RT) data were analyzed separately.

Experiment 1a: Matching targets located between 100 and 500 m

The results of Experiment 1a are shown in Figure 5. For the accuracy data (Panel A), across the whole range of distances examined, there was a clear advantage for processing dynamic ($M = 0.89$, $SE = 0.02$) versus static ($M = 0.68$, $SE = 0.01$) actions, $F(1, 7) = 84.18$, $MSE = 0.01$, $p < 0.001$, eta_2 = 0.92.

It is clear from Figure 5, that in all conditions, performance remained well above chance (i.e., 50% correct). This is particularly obvious in the dynamic case, where performance is essentially at ceiling levels until 300 m, after which there is a gentle decline. Again, even at the most extreme distance examined here, performance with dynamic actions was extremely good ($M = 0.80$, $SE = 0.04$). Performance in the static condition also remained above chance, following a slightly different trajectory in which accuracy begins to decline immediately between 100 and 300 m, and then to level off at the extremes of the range. These overall patterns led to a significant main effect of Distance, $F(4, 28) = 14.78$, $MSE = 0.003$, $p < 0.001$, eta_2 = 0.68, and a Motion × Distance interaction, $F(4, 28) = 2.98$, $MSE = 0.005$, $p < 0.05$, eta_2 = 0.3.

For the reaction time data (Panel C) the most obvious pattern is the steady slowdown of responses for both dynamic and static action that begins at 300 m, reflected in a main effect of Distance, $F(4, 28) = 35.35$, $MSE = 20.53$, $p < 0.001$, eta_2 = 0.84. There was no main effect of Motion, and although there appears to be a trend for dynamic responses to slow more sharply than static responses, the Motion × Distance interaction was also not significant. Note that overall, reaction times in this matching task were fairly slow (i.e., > 2.0 s). This would seem to reflect the need to compare at least two actions, rather than simply responding to a single target.

Experiment 1b: Matching targets located between 200 and 1000 m

The results of Experiment 1b are also shown in Figure 5. Again, at all distances, dynamic trials ($M = 0.72$, $SE = 0.02$) were matched more accurately than static trials ($M = 0.57$, $SE = 0.01$), $F(1, 10) = 137.38$, $MSE = 0.005$, $p < 0.001$, eta_2 = 0.93. In contrast to Experiment 1a, there is now a clear linear trend to the dynamic data, with accuracy dropping consistently at approximately 4% per 100 m across the whole range of distances. However, it is important to note that even at the most extreme end of this range, performance was at 58%, a level that was still above chance according to a nonparametric binomial test ($p < 0.01$). Static performance appears to decrease slightly less steeply at the beginning of the distance range but falls to chance beyond 600 m. These patterns are reflected in a main effect of Distance, $F(4, 40) = 41.25$, $MSE = 0.005$, $p < 0.001$, eta_2 = 0.81, and a Motion × Distance interaction, $F(4, 40) = 5.32$, $MSE = 0.006$, $p < 0.01$, eta_2 = 0.35.

Again in contrast to Experiment 1a, over this range of distances participants took significantly longer to match dynamic ($M = 2.5$ s, $SE = 0.09$) as compared to static ($M = 2.3$ s, $SE = 0.1$) actions, $F(1, 10) = 16.60$, $MSE = 98.32$, $p < 0.01$, eta_2 = 0.62. The most likely explanation for this pattern is that participants could gain an advantage by continuing to sample the dynamic trials—and also possibly increase their confidence in the match—as additional cues would arise as the actions unfolded. In static trials, cues did not vary over time.

There was also a main effect of Distance, $F(4, 40) = 2.87$, $MSE = 47.74$, $p < 0.05$, eta_2 = 0.22, that appears to be driven by slightly faster responses at 200 m, the nearest distance tested. At 400 m and beyond, response times were fairly constant. Although there is slight divergence at 1000 m, with the dynamic condition increasing and the static condition decreasing, the Motion × Distance interaction did not approach significance.

The only influence of gender was in the speed of responses, where female participants ($M = 2.22$ s, $SE = 0.13$) were consistently faster than male participants ($M = 2.63$ s, $SE = 0.13$), giving rise to a main effect of Gender, $F(1, 10) = 5.87$, $MSE = 939.49$, $p < 0.05$, eta_2

= 0.37. As there were no accuracy effects and no interactions with either motion or distance, we will not discuss this finding any further.

Examining individual actions

In the current paper we included a variety of actions to ensure that our findings were not specific to a given set of movements. Previous research has suggested that the speed and accuracy of interpreting actions can vary quite considerably depending on the type of action (Dittrich, 1993; Vanrie & Verfaillie, 2004). Furthermore, use of a range of movements has helped shed light on how actions might be mentally represented (Giese, Thornton, & Edelman, 2008; Vangeneugden, Pollick, & Vogels, 2009) and processed in a top-down manner (Dittrich, 1993; Petrini et al., 2011) during perception. Here, we were interested in assessing whether any particular actions showed unusual distance scaling properties. In Figure 6, we've summarized dynamic and static responses to actions collapsing across experiment, observer, and distance. Error bars show the standard error of the mean computed across the 10 distance levels in Experiments 1a and 1b.

There are a number of points to note. First, it is clear that the dynamic advantage seen in both experiments is not being driven by a subset of actions but is present across the whole range of behaviors studied. Second, it appears that some actions (e.g., walk, wave, & jump) give rise to faster and more accurate responses than others (e.g., sweep, spade, pump). The rank ordering of the actions, from best to worse performance, was highly consistent across Experiments 1a and 1b (all rs > 0.65, ps < 0.05). It seems that the actions that support better performance are those that contain more whole-body or limb dynamics, an issue we return to in the General discussion. Finally, examination of the error bars suggests that distance is having a similar effect on all actions. That is, although there is some variability in the size of the error bars, there appears to be no extreme outliers. To confirm that the effect of distance was consistent across actions, we repeated the main experimental ANOVAs as a series of item analyses, collapsing across observers. For both experiments and both dependent variables, these analyses gave identical qualitative patterns as the original observer-based analyses.

Experiment 2

The results of Experiment 1 suggest that the ability to process action remains functional over a surprisingly large range of distances. It is also clear that this ability relies predominantly on dynamic aspects of the display rather than on static shape matching. What is not clear is whether motion at the local level (i.e., individual dots) or at the global level (i.e., the entire figure) is responsible for the observed performance (Bertenthal & Pinto, 1994; Mather et al., 1992; Thornton et al., 1998; Thurman & Lu, 2013; Troje & Westhoff, 2006; Wang, Zhang, He, & Jiang, 2010).

It is generally accepted that both sources of information contribute to the robust nature of biological motion processing (e.g., Thornton, Rensink, & Shiffrar, 2002; Thurman & Lu, 2013; Troje & Westhoff, 2006). Here though, given the nature of our matching task, it seems quite feasible that local motion would become favored over global motion at extreme distances. That is, as the distance increases and overall figure size decreases, participants might focus on matching the motion of one or two dots (e.g., wrist or ankle) as these could provide a quick and easy way to distinguishing two actions.

In Experiment 2, we addressed this question by including trials in which both target and flankers were turned upside down. This form of picture-plane inversion has become a standard tool for disrupting global processing of point-light figures (e.g., Chang & Troje, 2008, 2009; Pavlova & Sokolov, 2000; Sumi, 1984). Inversion is thought to disrupt the ability to extract configural information from body stimuli (Reed et al., 2003), in a similar way to that documented for faces (Rossion & Gauthier, 2002; Thompson, 1980; Yin, 1969). If our current task relies on matching familiar, global actions then this manipulation should significantly reduce performance (Dittrich, 1993). Conversely, as the local motion of the dots is unaffected by inversion, sustained performance in our task can be taken as an indication that matching individual dots was the key strategy adopted by participants in Experiment 1.

We should note that inversion is known to influence the interpretation of individual dots (Chang & Troje, 2008; Troje & Westhoff, 2006; Wang et al., 2010). When a terrestrial animal moves, it propels its limbs against gravity with a "characteristic ballistic-velocity profile" that could be a very reliable cue to animacy (Troje & Westhoff, 2006). When stimuli are turned upside down, the local patterns of acceleration no longer provide such a cue, thus giving rise to an independent local inversion effect (Chang & Troje, 2008, 2009). Here, however, we're not asking people to assess or detect animacy, but simply to match two patterns. If the predominant strategy in our task were to match dots then their animacy may be of little relevance. Also, to date, such "life detectors" have only been demonstrated for walking patterns. It is unclear whether such local inversion effects would be seen with the range of actions used in our study.
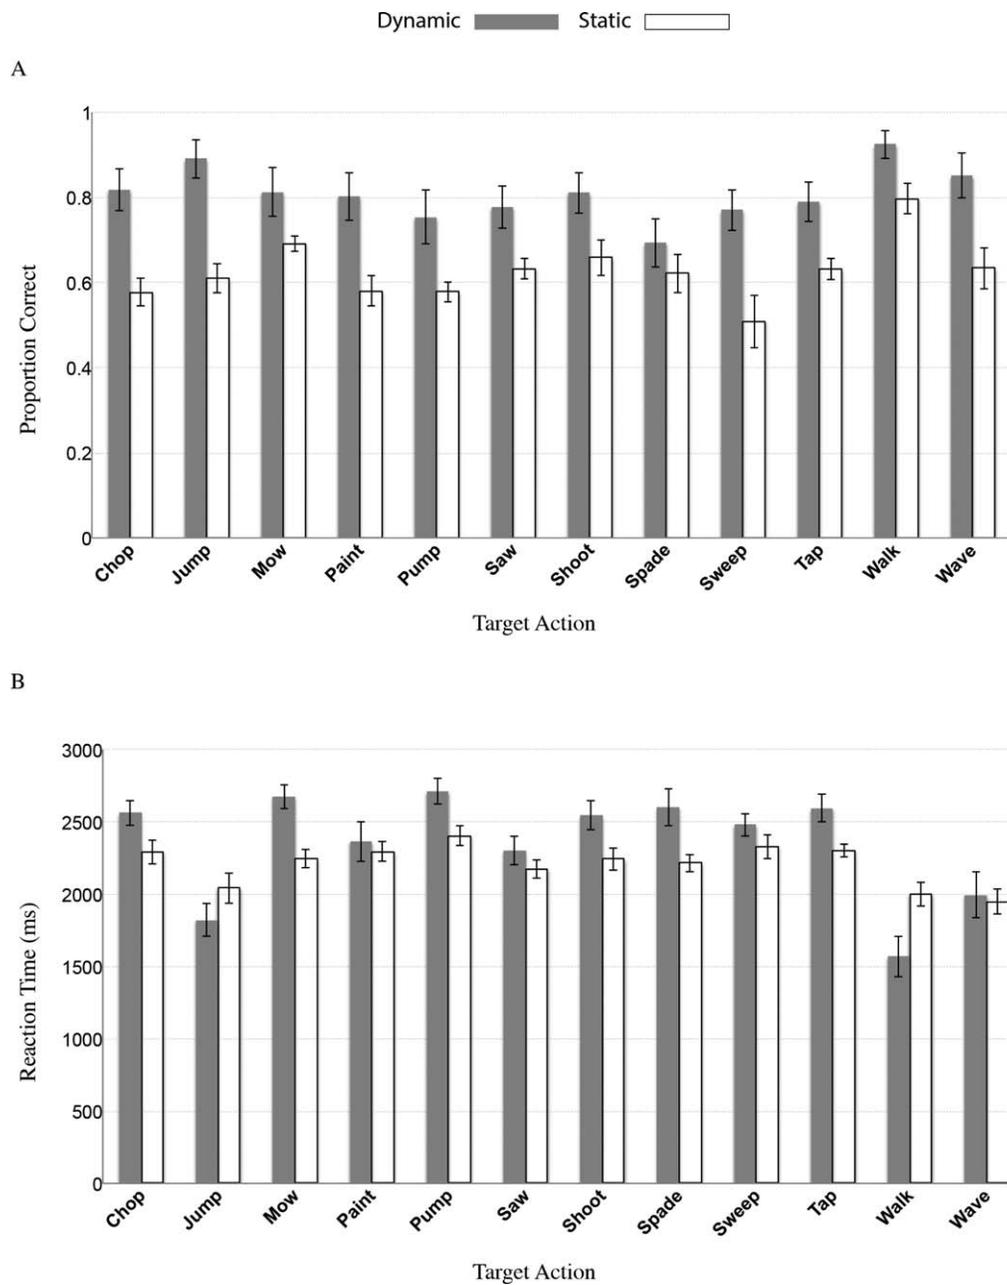
Figure 6. Summary of item analysis results, in terms of both accuracy (Panel A) and speed (Panel B). Data are collapsed across participant, distance, and experiment. Error bars show the standard error of the mean computed across the 10 distance intervals from both experiments.

As an aside, we should emphasize that our use of inverted figures is not an attempt to demonstrate the special status of biological versus nonbiological motion (Hiris, 2007; Neri et al., 1998; Poom & Olsson, 2002). Our goal in this paper is simply to assess the range of distances over which useful information could be extracted from moving human stimuli. We do not exclude the possibility that other mechanical or natural moving stimuli could also be processed at such distances. As noted above, the motivation for turning

our displays upside down was simply to assess the contribution of local and global processing, a manipulation that could equally be applied to other complex moving stimuli.

In Experiment 2a, we replicated the methods of Experiment 1b but removed the static trials and turned all of the dynamic trials upside down. We then compared performance in these inverted trials to the data obtained in the upright dynamic condition of Experiment 1b. We opted for this initial between-

subjects design as pilot testing indicated that when presented with inverted matching trials, participants had a tendency to slow down quite dramatically. Thus, this initial reduced design was aimed at verifying the presence of a speed/accuracy trade off.

Experiment 2b was again based on the design of Experiment 1b. Here, we used a within-subject design and interleaved upright and inverted dynamic trials. Apart from the introduction of inverted displays, the major modification to the design was to limit the overall viewing time to 3 s. The limit was imposed to equate exposure under upright and inverted conditions, based on our observations in Experiment 2a. The duration of 3 s was chosen as this was well within the time limit needed to produce above-chance dynamic performance in Experiment 1b. As our goal was essentially to determine whether performance in Experiment 1b relied on local or global performance, it was important to maintain a similar timeframe of responding. Thus, on each trial of Experiment 2b, the entire animation completed and disappeared from the screen. Only at this point could responses be made.

Participants

Twelve observers (eight female, four male) took part in Experiment 2a, and a further group of 12 observers (eight female, four male) in Experiment 2b. All participants had normal or corrected to normal vision. They all gave written informed consent and were told that the purpose of the experiment was to assess the perception of action at a distance. The study was reviewed by and received approval from the appropriate departmental Ethics Committee at Swansea University and therefore conformed to the ethical guidelines set out by the Declaration of Helsinki for testing human participants.

Stimuli, design, and analysis

The distance range, point-light actions, and task were identical to those used in Experiment 1b. The only significant changes were the removal of static trials, the inclusion of inverted trials, where all three figures were rotated in the picture plane by 180°, and the time limit imposed in Experiment 2b. In Experiment 2a, all 540 trials were shown in this inverted configuration. Participants completed three blocks of 180 trials. The within-block randomization was identical to Experiment 1b, except there were no static trials and participants saw 36 repetitions of each distance. Data from the inverted condition in this experiment was compared with the dynamic trials from Experiment 1b using a between-subjects design. Accuracy and reaction time data were analyzed using separate 2 (Orientation: Upright-Experiment 1b/Inverted-Experiment 2a) × 5 (Distance) ANOVAs.

In Experiment 2b, upright and inverted trials were interleaved within a block. Participants completed five blocks of 60 trials in a 2 (Orientation: Upright/Inverted) × 5 (Distance) × 6 (Repetition) design, giving a total of 300 trials. The within-block randomization was again identical to Experiment 1b. Accuracy and reaction time data were analyzed using separate 2 (Orientation: Upright/Inverted) × 5 (Distance) repeated measures ANOVAs.

Experiment 2a: Assessing speed–accuracy issues with inversion

Accuracy data from this experiment are shown in Figure 7A (solid line). Data from the upright dynamic condition of Experiment 1b are replotted for comparison (dashed line). It is immediately clear that the two conditions largely overlap. Accuracy with inverted displays ranged from approximately 82% ($SE = 2.2$) correct at 200 m to 60% ($SE = 1.8$) at 1000 m.

Importantly, even at this most extreme distance performance was significantly above chance, as confirmed by a nonparametric binomial test ($p < 0.05$). The ANOVA comparing data from the two experiments confirmed there was a significant main effect of Distance, $F(4, 88) = 59.89$, $MSE = 0.005$, $p < 0.001$, eta_2 = 0.73, but no main effect of orientation, $F(1, 22) = 0.65$, $MSE = 0.01$, n.s. Although upright performance at 200 m ($M = 0.89$, $SE = 0.02$) is consistently higher than inverted performance ($M = 0.82$, $SE = 0.02$), $t(22) = 2.33$, $p < 0.05$, there was no Orientation × Distance interaction, $F(4, 88) = 1.55$, $MSE = 0.005$, n.s.

The reaction time data, shown in Figure 7C, tell a very different story. Inverted trials ($M = 3.45$ s, $SE = 0.20$) were almost a second slower than upright trials ($M = 2.53$ s, $SE = 0.20$) leading to a significant main effect of Orientation, $F(1, 22) = 10.16$, $MSE = 2482.86.$, $p < 0.01$, eta_2 = 0.32. In both conditions, there was a tendency to slow down as distance increased, leading to a significant main effect of Distance, $F(4, 88) = 12.77$, $MSE = 99.10.$, $p < 0.001$, eta_2 = 0.37. The steeper increase in RT as a function of distance gave rise to a significant Orientation × Distance interaction, $F(4, 88) = 4.65$, $MSE = 99.10$, $p < 0.01$, eta_2 = 0.18.

Overall, these data confirm our initial pilot observations that inversion leads to a significant slowing of responses. Participants were still able to match inverted actions at extreme distance, but only by taking a considerably longer time to reach decisions, suggesting that matching in this condition is relying on strategies other than those that supported performance in the upright condition of Experiment 1b. To further explore
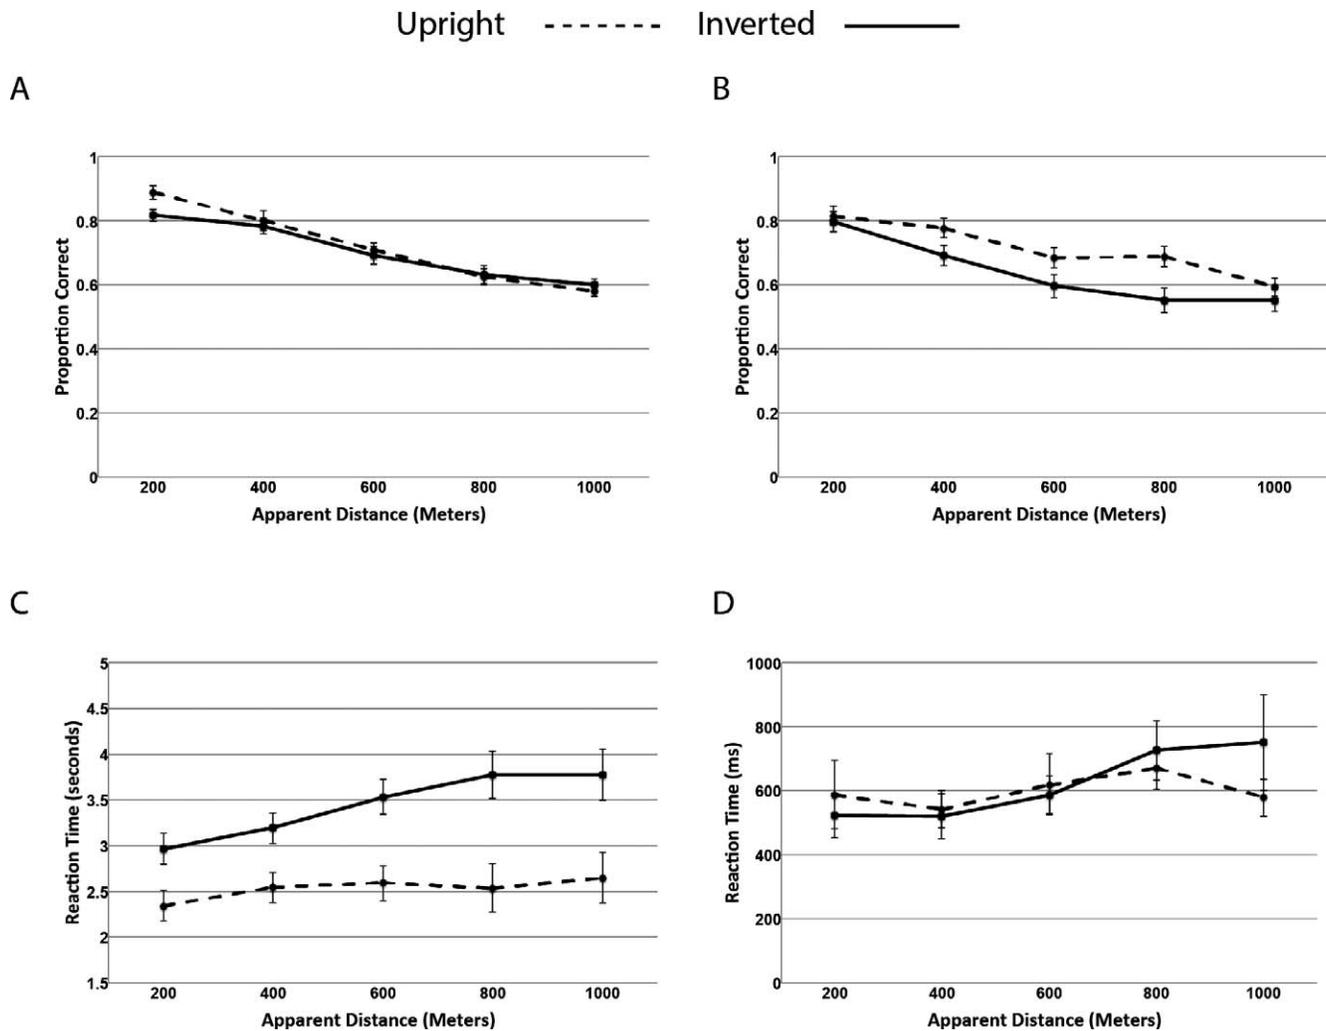
Figure 7. Performance in Experiments 2a and 2b in terms of both accuracy (Panels A–B) and speed (Panels C–D).

the relevance of the observed speed–accuracy trade-off, in Experiment 2b, we restricted the time available for making matching decisions.

Experiment 2b: Interleaving upright and inverted trials with a fixed duration

In Experiment 2b, participants were shown upright and inverted actions for a fixed duration of 3 s. Responses, could only be made after the stimuli had disappeared. Figure 7B, shows the accuracy data from this experiment. As in all previous datasets, there is a clear influence of distance on the accuracy of matching responses, leading to a significant main effect of Distance, $F(4, 44) = 23.25$, $MSE = 0.01.$, $p < 0.001$, eta_2 = 0.68. In contrast to Experiment 2a, across all levels of distance, upright figures ($M = 0.71$, $SE = 0.02$) were matched more accurately than inverted figures ($M = 0.66$, $SE = 0.02$), leading to a significant main effect of

Orientation, $F(1, 11) = 8.15$, $MSE = 0.01.$, $p < 0.05$, eta_2 = 0.43. There was no interaction, $F(4, 44) = 1.91$, $MSE = 0.01$, n.s. Importantly, as with Experiment 1b, matching performance in the dynamic upright condition remained above chance even at the most extreme distance of 1000 m (binomial test, $p < 0.05$). In the inverted condition, performance was at chance beyond 600 m.

Reaction time data are shown in Figure 7D. In contrast to the other studies in this paper, these times are measured from stimulus offset, and thus are substantially faster, remaining well below 1 s in all conditions (note the change of scale). There was a main effect of Distance, $F(4, 44) = 3.52$, $MSE = 33523.53$, $p < 0.05$, eta_2 = 0.24, which appears to reflect the overall slowing of responses as the range increases. There was no main effect of Orientation, $F(1, 11) = 0.19$, $MSE = 71853.79$, n.s., nor any interaction between Orientation × Distance, $F(4, 44) = 1.12$, $MSE = 57237.54$, n.s.

General discussion

The current paper set out to explore the range of distances over which it might be possible to accurately process human action. In Experiment 1, using calibrated point-light displays and a novel concurrent matching task, we found that dynamic actions could be matched at above chance levels, even at the extreme apparent distance of 1000 m. At this distance, where dynamic performance was 58% correct, the height of the entire figure was extremely small, subtending approximately 0.1° visual angle. Although dynamic performance never fell to chance levels, the clear linear trend in the data from Experiment 1b makes it possible to predict the distance at which matching would fail, a distance of approximately 1160 m. Across the whole range of distances we explored, there was a clear advantage for dynamic over static actions, and importantly, at the most extreme distances, only movement supported performance.

In Experiment 2, we investigated whether this movement advantage relied on local or global processing by introducing inverted trials. In Experiment 2a, when participants could free view the displays, the accuracy of inverted matching performance was equivalent to that seen for upright figures in Experiment 1b. However, in order to achieve this level, participants took almost a second longer to view the displays. This would suggest that matching based on local processing is possible, but is unlikely to account for the upright performance seen in Experiment 1b. To verify this suggestion, in Experiment 2b we limited the exposure time for both upright and inverted trials. Under these conditions, there was a clear accuracy advantage for upright trials.

Taken together, then, these results suggest that action perception remains functional over a surprisingly large range of distances and that this performance relies, at least in part, on the presence of familiar, global, dynamic patterns. Indeed, within the first 200 m, matching performance appears to be little affected and may even be considered invariant with respect to distance in this range. Beyond 200 m, however, there is a clear cost associated with accurate matching. This cost, which is initially reflected in response times, and beyond 400 m in accuracy, could signal a transition between different forms of processing, a point we return to shortly.

Although in our everyday experience, we would probably not be often called upon to make judgments at such extreme distances—indeed finding unobstructed views in a typical city landscape might be quite challenging—this sustained performance would appear to represent yet another demonstration of the robust and flexible nature of our ability to process human dynamics. From a functional perspective, we would thus expect human observers to be able to respond appropriately to suspicious or threatening behavior either at extreme physical distances or in reduced viewing conditions such as those that might obtain in modern surveillance settings.

We should note, however, that our current findings could be overestimating more general action processing ability. That is, point-light figures on an uncluttered background may, in a number of ways, serve as a near-optimal stimulus. Not only are the figures highly visible—a feature clearly exploited in the pedestrian safety work discussed earlier (Balk et al., 2008; Luoma & Penttinen, 1998; Owens et al., 1994; Tyrrell et al., 2009; Wood et al., 2005)—but our use of a single pixel to represent the joints means that self-occlusion would be minimized. When viewing a normal solid body figure at such extreme distances the visibility of the extremities, such as the wrist and ankles, might become obscured by other portions of the body much more quickly than in the current displays. Also, as the overall density of the points is much reduced at extreme distances, the average luminance increases proportionally. For the very smallest figures, the proximity of neighboring dots gives rise to an apparent "glow" that, at least subjectively, appears to aid in basic figure/ground segregation.

Our concurrent matching task may also be slightly less demanding than action identification (e.g., Dittrich, 1993; Vanrie & Verfaillie, 2004) or movement discrimination (e.g., Mather et al., 1992; Thornton et al., 1998). In our task we provide two candidate actions that are always visible for the observer to compare and contrast. These displays may thus support a range of inclusion and exclusion strategies that could boost performance. Of course, when trying to interpret the actions of a distant person, it seems likely that we would generate internal hypotheses as to what they were doing and use these to guide our interpretation. This sort of top-down, "active" processing (e.g., Bertenthal & Pinto, 1994; Bülthoff et al., 1998; Cavanagh, Labianca, & Thornton, 2001; Thornton, Rensink, & Shiffrar, 2002) has been suggested as a parallel route to understanding action, in addition to the bottom-up "passive" processing favored by early theories and models of biological motion perception (e.g., Giese & Poggio, 2003; Mather et al., 1992). It seems highly likely that in our task, observers are engaging in this sort of active processing. This is clear both from subjective reflection on what it feels like to do the task, and also by looking at the levels of baseline reaction times. In any event, our matching task would seem to both encourage and support such strategies by providing external models (i.e., the flankers) on which to base such active processing of the target and these factors may help to explain sustained levels of performance with distance.

Returning to the current data, it is useful to consider the nature of the mechanisms that might underlie sustained performance with distance and to reflect on whether our findings might be specific to biological motion or might generalize to other objects and dynamic events. In general, it seems clear that both form and motion contribute to our ability to understand the actions of others (Giese & Poggio, 2003; Lange & Lappe, 2006; Thompson & Baccus, 2011; Thurman, Giese, & Grossman, 2010). The precise contribution of these two sources of information remains controversial (Beintema & Lappe, 2002; Casile & Giese, 2005). Consistent with a role for form processing, matching in the static control condition was surprisingly good in the current experiment, remaining above chance out to distances of 600 m. However, it is also clear that the presence of motion was able to boost performance and support perception beyond the limits possible with static matching. It would seem then that performance in the current task, at least at the most extreme distances, is reliant on dynamic cues.

An important distinction that is often made in relation to dynamic cues is that between local (Chang & Troje, 2008; Mather et al., 1992; Troje & Westhoff, 2006) and global (Bertenthal & Pinto, 1994; Thornton et al., 1998) processing. It seems that both types of processing are possible, and the dominance of one versus the other is often dictated by specific task and display characteristics. Consistent with this notion, the results of Experiment 2 suggest that both types of cue can support matching at extreme distance but that there is a clear advantage for conditions that support global processing (Bertenthal & Pinto, 1994).

In one sense, we found the results of Experiment 2b rather surprising. The subjective impression of watching these displays changes quite dramatically at the more extreme distances. Within 600 m, a global impression of the figure is surprisingly easy to extract even though the entire stimulus is very small. Beyond this limit, however, the global impression appears to diminish and both ourselves and a number of observers indicated that they began to look for specific distinguishing features, such as joint or limb movement, rather than trying to match based on the global action. Further support for this idea comes from the pattern of between-action performance, shown in Figure 6, where actions that involve limb movements outside of the main bodyline, such as waving, jumping, and walking, seem to be processed most effectively. If we take these subjective impressions at face value, then the effects of inversion measured in Experiment 2b could relate more to disruption in the ability to locate or interpret limb or local dot motion rather than processing of the global figure (Troje & Westhoff, 2006).

Finally, although we have couched this discussion within the framework of biological motion research, it would be premature to suggest that the observed pattern of performance at extreme distances is unique to decisions regarding human or even, more generally, animate motion. As just mentioned, our task clearly lends itself to strategic feature matching that could be equally applied to other types of motion (e.g., mechanical sheering or rotation). Our hunch, based on converging evidence of specialization for processing human bodies and our everyday familiarity with human action, is that biological motion would be particularly robust over extreme distances. The results of Experiment 2b are certainly consistent with this notion. However, a definitive answer to this question must clearly wait for further research that implements further control stimuli and tasks.

Conclusions

The current paper makes a number of contributions to our understanding of action perception. First, our review of previous work identified a gap in the biological motion literature with respect to how distance might affect our ability to interpret human action. Second, we provided an overview of the range of visual sizes that have been used in a large sample of biological motion studies, finding that the average figure height was 6.6°. This size is consistent with an adult figure standing approximately 15 m away from the point of observation. Third, we conducted the first systematic test of performance over extreme distances. Our results suggest that dynamic figures can still be matched at 1000 m—a distance at which they subtend approximately 0.1° visual angle—but that performance is likely to fall to chance just beyond this limit. Finally, we have introduced a new form of concurrent matching task that can easily be adapted to provide quick and easy estimates of a range of action perception abilities.

Acknowledgments

Commercial relationships: none.
Corresponding author: Ian M. Thornton.
Email: ian.thornton@um.edu.mt.
Address: Department of Cognitive Science, University of Malta, Msida, Malta.

References

Ames, A. (1952). *The Ames demonstrations in perception*. New York: Hafner Publishing.

Balk, S. A., Tyrrell, R. A., Brooks, J. O., & Carpenter, T. L. (2008). Highlighting human form and motion information enhances the conspicuity of pedestrians at night. *Perception*, 37, 1276–1284.

Barclay, C. D., Cutting, J. E., & Kozlowski, L. T. (1978). Temporal and spatial factors in gait perception that influence gender recognition. *Perception & Psychophysics*, 23, 145–152.

Beintema, J. A., & Lappe, M. (2002). Perception of biological motion without local image motion. *Proceedings of the National Academy of Sciences, USA*, 99, 5661–5663.

Bertenthal, B. I., & Pinto, J. (1994). Global processing of biological motions. *Psychological Science, 5,* 221–225.

Blake, R., & Shiffrar, M. (2007). Perception of human motion. *Annual Review of Psychology*, 58, 47–73.

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

Bülthoff, I., Bülthoff, H. H., & Sinha, P. (1998). Top-down influences on stereoscopic depth-perception. *Nature Neuroscience*, 1(3), 254–257.

Casile, A., & Giese, M. A. (2005). Critical features for the recognition of biological motion. *Journal of Vision*, 5(4):6, 348–360, http://www.journalofvision.org/content/5/4/6, doi:10.1167/5.4.6. [PubMed] [Article]

Cavanagh, P., Labianca, A., & Thornton, I. M. (2001). Attention-based visual routines: Sprites. *Cognition, 80,* 47–60.

Chang, D. H. F., & Troje, N. F. (2008). Perception of animacy and direction from local biological motion signals. *Journal of Vision*, 8(5):3, 1–10, http://www.journalofvision.org/content/8/5/3, doi:10.1167/8.5.3. [PubMed] [Article]

Chang, D. H. F., & Troje, N. F. (2009). Acceleration carries the local inversion effect in biological motion perception. *Journal of Vision*, 9(1):19, 1–17, http://www.journalofvision.org/content/9/1/19, doi:10.1167/9.1.19. [PubMed] [Article]

Curio, E., Ernst, U., & Vieth, W. (1978). Cultural transmission of enemy recognition: One function of mobbing. *Science*, 202, 899–901.

Darwin, C. (1872). *The expression of emotions in man and animals*. London: John Murray.

Dittrich, W. H. (1993). Action categories and the perception of biological motion. *Perception*, 22, 15–22.

Downing, P. E., & Peelen, M. V. (2011). The role of occipitotemporal body-selective regions in person perception. *Cognitive Neuroscience*, 2(3-4), 186–203.

Giese, M. A., & Poggio, T. (2003). Neural mechanisms for the recognition of biological movements. *Nature Reviews: Neuroscience, 4*, 179–192.

Giese, M. A., Thornton, I. M., & Edelman, S. (2008). Metrics of the perception of body movement. *Journal of Vision*, 8(9):13, 1–18, http://www.journalofvision.org/content/8/9/13, doi:10.1167/8.9.13.

Griffin, A. S., Blumstein, D. T., & Evans, C. S. (2000). Training captive-bred or translocated animals to avoid predators. *Conservation Biology*, 14, 1317–1326.

Grossman, E., & Blake, R. (2002). Brain areas active during visual perception of biological motion. *Neuron*, 35(6), 1167–1175.

Gurnsey, R., Roddy, G., Ouhnana, M., & Troje, N. F. (2008). Stimulus magnification equates identification and discrimination of biological motion across the visual field. *Vision Research*, 48, 2827–2834.

Gurnsey, R., Roddy, G., & Troje, N. F. (2010). Limits of peripheral direction discrimination of point-light walkers. *Journal of Vision*, 10(2):15, 1–17, http://www.journalofvision.org/content/10/2/15, doi:10.1167/10.2.15. [PubMed] [Article]

Hiris, E. (2007). Detection of biological and nonbiological motion. *Journal of Vision, 7*(12):4, 1–16, http://www.journalofvision.org/content/7/12/4, doi:10.1167/7.12.4. [PubMed] [Article]

Ikeda, H., Blake, R., & Watanabe, K. (2005). Eccentric perception of biological motion is unscalably poor. *Vision Research*, 45, 1935–1943.

Jackson, S., & Blake, R. (2010). Neural integration of information specifying human structure from form, motion and depth. *Journal of Neuroscience*, 30, 838–848.

Johansson, G. (1973). Visual perception of biological motion and a model for its analysis. *Perception & Psychophysics*, 14, 201–211.

Jokisch, D., & Troje, N. F. (2003). Biological motion as a cue for the perception of size. *Journal of Vision*, 3(4):1, 252–264, http://www.journalofvision.org/content/3/4/1, doi:10.1167/3.4.1. [PubMed] [Article]

Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3?. *Perception*, 36(ECVP Abstract Supplement).

Lange, J., & Lappe, M. (2006). A model of biological motion perception from configural form cues. *Journal of Neuroscience*, 26(11), 2894–2906.

Legault, I., Troje, N. F., & Faubert, J. (2012). Healthy

older observers cannot use biological motion point light information efficiently within 4 meters of themselves. *iPerception*, *3*, 104–111.

Luoma, J., & Penttinen, M. (1998). Effects of experience with retroreflectors on recognition of nighttime pedestrians: Comparison of driver performance in Finland and Michigan. *Transportation Research Part F: Traffic Psychology and Behaviour*, *1*(1), 47–58.

Mather, G., Radford, K., & West, S. (1992). Low level visual processing of biological motion. *Proceedings of the Royal Society of London, Series B, 249,* 149–155.

Neri, P., Morrone, M. C., & Burr, D. C. (1998). Seeing biological motion. *Nature*, *394*, 894–896.

Owens, D. A., Antonoff, R. J., & Francis, E. L. 1994. Biological motion and nighttime pedestrian conspicuity. *Human Factors*, *36*(4), 718–732.

Palmisano, S., Gillam, B., Govan, D. G., Allison, R. S., & Harris, J. M. (2010). Stereoscopic perception of real depths at large distances. *Journal of Vision*, *10*(6):19, 1–16, http://www.journalofvision.org/content/10/6/19, doi:10.1167/10.6.19. [PubMed] [Article]

Pavlova, M., & Sokolov, A. (2000). Orientation specificity in biological motion perception. *Perception & Psychophysics*, *62*, 889–899.

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

Petrini, K., Pollick, F. E., Dahl, S., McAleer, P., McKay, L., Rocchesso, D., . . . Puce, A. (2011). Action expertise reduces brain activity for audio-visual matching actions: An fMRI study with expert drummers. *NeuroImage*, *56*(3), 1480–1492.

Poom, L., & Olsson, H. (2002). Are mechanisms for perception of biological motion different from mechanisms for perception of nonbiological motion? *Perceptual & Motor Skills*, *95*(3), 1301–1310.

Reed, C. L., Stone, V., Bozova, S., & Tanaka, J. (2003). The body inversion effect. *Psychological Science*, *14*, 302–308.

Rossion, B., & Gauthier, I. (2002). How does the brain process upright and inverted faces? *Behavioral and Cognitive Neuroscience Reviews*, *1*, 63–75.

Saygin, A. P. (2007). Superior temporal and premotor brain areas necessary for biological motion perception. *Brain*, *130,* 2452–2461.

Sumi, S. (1984). Upside-down presentation of the Johansson moving light-spot pattern. *Perception*, *13*, 283–286.

Thirkettle, M., Benton, C. P., & Scott-Samuel, N. E. (2009). Contributions of form, motion and task to

biological motion perception. *Journal of Vision*, *9*(3):28, 1–11, http://www.journalofvision.org/content/9/3/28, doi:10.1167/9.3.28. [PubMed] [Article]

Thirkettle, M., Scott-Samuel, N. E., & Benton, C. P. (2010). Form overshadows 'opponent motion' information in processing of biological motion from point light walker stimuli. *Vision Research*, *50*(1), 118–126.

Thompson, J. C., & Baccus, W. (2011). Form and motion make independent contributions to the response to biological motion in occipitotemporal cortex. *NeuroImage*, *59*(1), 625–634.

Thompson, J., & Parasuraman, R. (2012). Attention, biological motion, and action recognition. *NeuroImage*, *59*, 4–13.

Thompson, P. (1980). Margaret Thatcher - A new illusion. *Perception*, *9*, 483–484.

Thornton, I. M. (2006). Biological motion: Point-light walkers and beyond. In G. Knoblich, I. M. Thornton, M. Grosjean, & M. Shiffrar (Eds.), *Human body perception from the inside out* (pp. 271-305). New York:Oxford University Press.

Thornton, I. M., Pinto, J., & Shiffrar, M. (1998). The visual perception of human locomotion. *Cognitive Neuropsychology*, *15*, 535–552.

Thornton, I. M., Rensink, R. A., & Shiffrar, M. (2002). Active versus passive processing of biological motion. *Perception*, *31*, 837–853.

Thornton, I. M., & Vuong, Q. C. (2004). Incidental processing of biological motion. *Current Biology*, *14*, 1084–1089.

Thurman, S. M., Giese, M. A., & Grossman, E. D. (2010). Perceptual and computational analysis of critical features for biological motion. *Journal of Vision*, *10*(12):15, 1–14, http://www.journalofvision.org/content/10/12/15, doi:10.1167/10.12.15. [PubMed] [Article]

Thurman, S. M., & Lu, H. (2013). Complex interactions between spatial, orientation, and motion cues for biological motion perception across visual space. *Journal of Vision*, *13*(2):8, 1–18, http://www.journalofvision.org/content/13/2/8, doi:10.1167/13.2.8. [PubMed] [Article]

Troje, N. F., & Westhoff, C. (2006). The inversion effect in biological motion perception: Evidence for a "life detector?." *Current Biology*, *16*, 821–824.

Tyrrell, R. A., Wood, J. M., Chaparro, A., Carberry, T. P., Chu, B.-S., & Marszalek, R. P. (2009). Seeing pedestrians at night: Visual clutter does not mask biological motion. *Accidents Analysis and Prevention*, *41*(3), 506–512.

Vangeneugden, J., Pollick, F., & Vogels, R. (2009).

Functional differentiation of macaque visual temporal cortical neurons using a parametric action space. *Cerebral Cortex*, *19,* 593–611.

Vanrie, J., & Verfaillie, K. (2004). Perception of biological motion: A stimulus set of human point-light actions. *Behavior Research Methods, Instruments, & Computers*, *36*, 625–629.

Wang, L., Zhang, K., He, S., & Jiang, Y. (2010). Searching for life motion signals visual search asymmetry in local but not global biological-motion processing. *Psychological Science*, *21*(8), 1083–1089.

Wood, J. M., Tyrrell, R. A., & Carberry, T. P. (2005). Limitations in drivers' ability to recognize pedestrians at night. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *47*(3), 644–653.

Yin, R. (1969). Looking at upside down faces. *Journal of Experimental Psychology*, *81*, 141–145.