

Low-Level Visual Processing of Biological Motion

George Mather, Kirstyn Radford and Sophie West

Proc. R. Soc. Lond. B 1992 **249**, 149-155

doi: 10.1098/rspb.1992.0097

References

Article cited in:

<http://rsob.royalsocietypublishing.org/content/249/1325/149#related-urls>

Email alerting service

Receive free email alerts when new articles cite this article - sign up in the box at the top right-hand corner of the article or click [here](#)

To subscribe to *Proc. R. Soc. Lond. B* go to: <http://rsob.royalsocietypublishing.org/subscriptions>

Low-level visual processing of biological motion

GEORGE MATHER, KIRSTYN RADFORD AND SOPHIE WEST

Laboratory of Experimental Psychology, University of Sussex, Brighton BN1 9QG, U.K.

SUMMARY

Biological motion displays depict a moving human figure by means of just a few isolated points of light attached to the major joints of the body. Naive observers readily interpret the moving pattern of dots as representing a human figure, despite the complete absence of form cues. This paper reports a series of experiments which investigated the visual processes underlying the phenomenon. Results suggest that (i) the effect relies upon responses in low-level motion-detecting processes, which operate over short temporal and spatial intervals and respond to local modulations in image intensity; and (ii) the effect does not involve hierarchical visual analysis of motion components, nor does it require the presence of dots which move in rigid relation to each other. Instead, movements of the extremities are crucial. Data are inconsistent with current theoretical treatments.

1. INTRODUCTION

The ability to recognize fellow members of one's species by sight is clearly vital for survival and evolution. Visual images of articulated biological forms are inherently complex: when a person walks across an observer's field of view he or she projects a complex pattern of translatory, elliptical and pendular movements on the observer's retina, because of the articulated structure of rigid body parts (e.g. elliptical motion of each shoulder, pendular motion of each arm, superimposed on the overall translation of the body across the field of view). In a series of striking demonstrations, Johansson (1973) found that even if the walking person is completely invisible except for about a dozen strategically placed points of light at the major joints, naive observers immediately perceive the moving cluster of lights as a walking person. Any single static view of the person is reported as an unstructured collection of dots. How is this remarkable perceptual feat achieved? It is widely held that the human visual system has evolved efficient but specialized processes for the analysis of such 'biological motion' stimuli. Theoretical models have so far set the phenomenon apart from the simple rigid movements typically studied in motion research, classifying it as an example of 'event perception'. Accordingly, the complexity of the moving object is emphasized. Sophisticated high-level visual processes are said to extract the motion of rigid body components individually, and then structure them to form the final percept (Cutting & Proffitt 1981; Webb & Aggarwal 1982). However, few psychophysical data are available on the perception of biological motion which can be used to infer the nature of the visual processes involved. The studies reported here address two questions. First, how are the movements of individual dots in a biological motion display encoded by the visual system? Second, how are

these separate signals integrated to form a representation of a walking figure?

2. GENERAL METHODS

Observers viewed a simulation of biological motion on a computer screen, developed using the algorithm described in Cutting (1978). The shifting pattern of dots generated by a walking figure was sampled to create 40 static views during a single step cycle (the period between the figure striking the same posture, such as left foot extended). When this series of static frames was presented in rapid succession, observers reported a compelling impression of a walking figure, as expected. In each frame, 12 points were plotted to define the figure, two each (left and right) for the shoulders, elbows, wrists, hips, knees, and ankles. They simulated the pattern generated by a sideways view of a person walking on a treadmill. In other words, the dot displacements contained no translatory component, only elliptical and oscillatory components. A square outline border surrounded the figure. The border subtended 6.4° on a side, and the distance between the figure's shoulder and hip dots subtended 1.6° . Figure 1*a* shows one frame from the animation sequence used in the experiments.

Two tasks were done in different experiments to assess biological motion analysis in differing conditions. In the 'direction' task, the direction the simulated walker faced (left against right) varied randomly from presentation to presentation, and the observer had to identify direction after each presentation. In the 'coherence' task, half of the presentations depicted a coherent display in which all dots were consistent with a person facing left (or right; see figure 1*a*). The other half of the presentations were incoherent in that the dots belonging to the upper body (shoulders, elbows, wrists) were consistent with a left-facing (or right-facing) walker whereas dots belonging to the lower body (hips, knees, ankles) were consistent with a right-facing (or left-facing) walker (see figure 1*b*). In this task observers had to discriminate between coherent and incoherent presentations. The reason for using two different tasks was to check that subjects were not basing their discrimination on some cue in the display unique to the




Figure 1. Views taken from apparent motion sequences depicting a sideways view of a walking human figure. Points represent the locations of the major joints. (a) A coherent view in which all points represent a right-facing figure. (b) An incoherent view in which upper body points are consistent with a left-facing figure and lower body points are consistent with a right-facing figure. In experimental stimuli, 60 randomly plotted background dots were also present, removing positional cues. Figure is drawn to scale. The background square subtended 6.4 arc deg on a side.

particular task, rather than on extracting a representation of a walking figure. Further, two measures were adopted to prevent observers making judgements on the basis of predictable spatial cues rather than motion cues (i.e. static locations of individual dots in particular frames). First, the horizontal location of the walker within the square border varied randomly from presentation to presentation. Second, 60 randomly located dots were plotted in each frame, in addition to the 12 belonging to the figure, and the positions of these irrelevant dots changed randomly from frame to frame. Thus, in a static view of any individual frame, it was not possible to discriminate the dots belonging to the figure from the rest. However, when the animated sequence was viewed, it gave the appearance of a figure striding through a light snowstorm (these irrelevant dots were omitted from figure 1 to make the walker dots visible). Note that all 12 dots defining the figure were always visible, even when they should have been occluded by other parts of the body, to avoid unwanted local cues (e.g. 'if leftward moving dots disappear, then the walker must be facing left').

A trial consisted of a single brief presentation of the walking figure, usually containing only a subset of the 40 frames defining a full step cycle. The particular set chosen varied randomly from trial to trial. After the presentation the observer was given a forced choice between two responses. For experiments using the direction task, the observer was required to respond 'left' or 'right' using two buttons; for experiments using the coherence task, the observer was required to respond 'coherent' or 'incoherent'. Direction or coherence varied randomly from trial to trial, as did the parameters defining different conditions in each experiment. When statistical significance levels are reported during presentation of results, they were derived from analyses of variance on arcsine-transformed data.

A pool of 15 observers participated in the experiments, either five or six serving in each. All except the authors were naive. All experiments after the first involved a mixture of new subjects and those who had served in at least one previous experiment, allowing us to check for long-term practice effects. Each new subject was first shown a demonstration of a biological motion display similar to that used in the experiments, and instructed on the nature of the

response they were to make (judge direction or coherence). The first experimental session which followed was treated as a practice session, and its data discarded. No more practice sessions were given for that observer in any subsequent experiments.

3. SPATIOTEMPORAL PROPERTIES OF BIOLOGICAL MOTION

Perceptual and physiological evidence shows that the human visual system possesses a specialized low-level motion-detecting system which operates over very short time intervals, processing changes in the image which occur over a time span of approximately 50 ms (see, for example, Baker & Braddick 1985). However, various motion displays have been described in the literature which cannot stimulate this low-level process, yet still give an impression of movement (see, for example, Mather 1988; Chubb & Sperling 1988). A higher-level motion process must be invoked to explain these effects. According to available data, the high-level process can operate over much longer intervals than the low-level process (up to about 500 ms).

It is not clear which of the two putative processes is involved in the perception of biological motion. Because of their inherent complexity, it is generally assumed that biological motion stimuli are analysed by the high-level process. An experiment was done to determine whether the time dependency of the percept was consistent with the involvement of low-level processes or high-level processes. Subjects viewed animation sequences containing 20 frames in a single presentation, and were required to make a coherence discrimination after each presentation. Each frame in the sequence was visible on the display screen for 40 ms. The time interval between one frame disappearing and the next appearing on the display (inter-frame interval) varied in different presentations between 0 ms, 60 ms and 120 ms. During the interval (if any)




Figure 2. Results of an experiment to investigate the effect of inter-frame interval on perception of biological motion. Each data point is the mean percentage correct across five subjects, based on 40 trials per condition per subject. Vertical bars represent standard errors. Subjects discriminated coherent presentations (figure 1*a*) from incoherent presentations (figure 1*b*). Filled symbols represent results when frames were presented in the correct order, and open symbols represent results when frame order was randomized.

between frames, all dots disappeared, leaving only the outline border visible. If perception of biological motion depends on higher-level processes then there should be little effect of varying inter-frame interval in the range used. But if low-level processes are important then subjects' ability to do the task should deteriorate at the longer inter-frame intervals. To ascertain whether the measures described in the previous paragraph were successful in eliminating cues based on spatial structure, a control condition was included. In control presentations the motion information was corrupted simply by randomizing the order in which frames appeared in the animation sequence.

Results for five subjects are shown in figure 2, which plots mean percentage correct discrimination of coherence as a function of inter-frame interval. The solid squares represent results using correctly ordered animation sequences, and open squares represent results using disordered sequences. Subjects could make the discrimination reliably only when the animation sequence was correctly ordered and when the inter-frame interval was zero, suggesting an important role for low-level motion processes. The effect of frame ordering was highly significant ($p < 0.01$), as was the interaction between frame ordering and inter-frame interval ($p < 0.01$).

As inter-frame interval increased, simulated walking speed slowed down below naturalistic levels, and this may have contributed to poorer performance at longer intervals. A second experiment was done to test for this possibility. Ten frames were presented in each trial, with frame duration fixed at 24 ms. Inter-frame interval was set to 12 ms, 48 ms or 120 ms in different trials. The actual frames presented were sampled from the full 40-frame step cycle at intervals of either one frame (i.e. consecutive, as in the first experiment), two

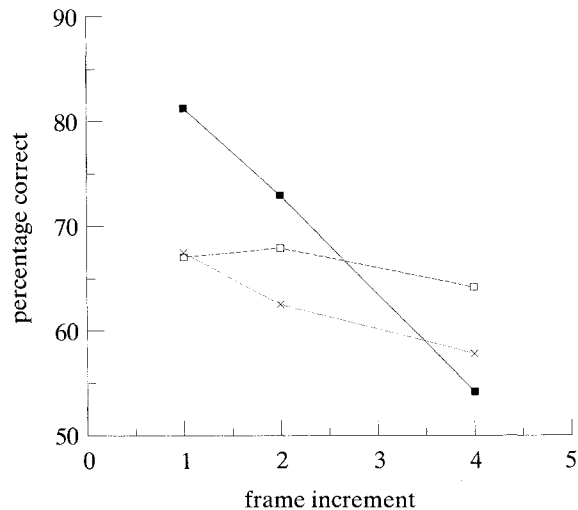


Figure 3. Results of an experiment in which subjects discriminated the direction the walking figure faced (left against right, all presentations were coherent), as a function of frame increment (abscissa) and inter-frame interval. Filled squares, inter-frame interval = 12 ms; open squares, inter-frame interval = 48 ms; crosses, inter-frame interval = 120 ms. Standard errors are omitted for clarity, but were similar to those plotted in figure 2.

frames, or four frames. In the 3×3 factorial combination of three inter-frame intervals and three frame increments, three conditions involved natural walking speed (Carlsoo 1972): consecutive frames at 36 ms asynchrony, every second frame at 72 ms asynchrony, and every fourth frame at 144 ms asynchrony. Subjects were required to make direction discriminations rather than coherence discriminations, to determine whether data in the first experiment were specific to the task involved. Each subject accumulated 40 observations in each of the nine conditions over five experimental sessions.

Results based on five subjects are shown in figure 3. The effects of inter-frame interval and frame increment interacted significantly ($p < 0.02$): when inter-frame interval was short, performance was optimal at shorter frame increments; at longer inter-frame intervals, performance was lower for all frame increments. Conditions which were matched for walking speed showed wide differences in performance (81%, 68% and 58%), so we conclude that walking speed was not responsible for the pattern of data in the first experiment. Instead, performance is determined jointly by inter-frame interval and frame increment. Why is frame increment important? As frame increment increased, the frame-to-frame displacement of each dot in the figure increased. Low-level motion processes are known to operate only over shorter displacements (typically estimated at below 0.25 arc deg), and would therefore respond poorly at larger frame increments. To assess the plausibility of this account, we computed the mean frame-to-frame displacement of each dot in the figure at different frame increments. Computed values are shown in figure 4. At larger frame increments, more dots approach the upper displacement limits of low-level processes.




Figure 4. Computed mean frame-to-frame displacement (over the full step cycle) of each dot in the walking figure, as a function of frame increment. The three bars for each dot represent frame increments of 1 (pale shading), 2 (unshaded) and 4 (dark shading) frames.

4. EFFECT OF INTENSITY MODULATION ON BIOLOGICAL MOTION

Results of the first experiments suggest that low-level motion processes are important in the analysis of biological motion. As a second test of this hypothesis, an experiment was done in which only ordered sequences, short inter-frame intervals, and consecutive frames were used, but dot intensity was manipulated. In one condition, the contrast of all dots was fixed during the animation sequence, and in a second condition the intensity of each dot in each animation frame alternated randomly between darker and brighter than the background. Responses in low-level motion processes should be severely disrupted by such random contrast reversals during animation (Chubb & Sperling 1988). Subjects performed direction discrimination and coherence discrimination tasks in different experimental sessions, and trials involving fixed and random contrast were presented in random order within a session. Session order was randomized across subjects. In each session, half of the presentations were from the fixed contrast condition (all dots either darker or lighter than the background, selected at random for each trial). In the remaining presentations (random contrast), the intensity of each dot in each frame of the animation sequence (both figure and background dots) was chosen at random from two values, one darker than the background and one lighter than the background. Intensities were: background, 80 cd m^{-2} ; bright dots, 140 cd m^{-2} ; dark dots, 20 cd m^{-2} ; outline border, 140 cd m^{-2} . For all presentations, the duration of each frame was fixed at 36 ms, with no inter-frame interval, to simulate a natural walking speed.

Results are shown in figure 5. In both tasks, performance was good when dot contrast remained fixed but poor when dot contrast reversed randomly during the animation sequence. The effect of contrast reversal was statistically significant for both tasks ($p < 0.01$). These data confirm the results of the first

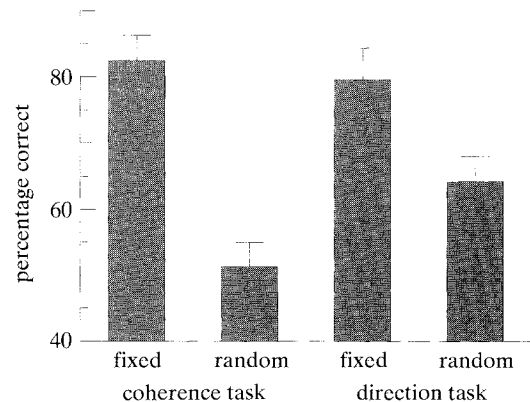


Figure 5. Results of an experiment to investigate the effect of random reversals in dot contrast on perception of biological motion. Each bar represents the mean percentage correct for five subjects, with 40 trials per subject per condition. Vertical bars represent standard errors. The two left-hand bars depict results from the coherence task, and right-hand bars depict results from the direction task. Results show that performance in both tasks was good when dot contrast against the background remained fixed, but poor when contrast reversed randomly, as expected if perception is mediated by low-level intensity-based motion mechanisms.

experiments: low-level motion-detecting processes have a major role in the analysis of biological motion. Previous work has shown that they have a major role in the perception of depth through motion (Mather 1989; Doshier *et al.* 1989), so these experiments establish a link between analysis of biological motion and analysis of kinetic depth stimuli.

The dots representing the walking figure were embedded in random noise, so subjects had to segment the figure from the noise to perform the tasks. The ability to segment an image on the basis of motion has been used as a defining feature of low-level processes. One could argue that our displays forced the system to depend on low-level responses. If this is the case, then the most one can assert is that low-level responses are sufficient, but perhaps not necessary, for biological motion perception. In fact, data obtained from early pilot experiments which did not use random background dots were also consistent with the importance of low-level processes. In one experiment, four subjects made coherence judgements from viewing short sequences of one, two, four or eight frames, at a frame duration of 20 ms. Mean percentages correct were: one frame, 61.5%; two frames, 73%; four frames, 80%; eight frames, 86.5%. Consistently correct discriminations from such brief, rapidly presented stimuli are most plausibly explained by the involvement of low-level processes. An experiment on the perception of motion in subjective figures, associated with high-level processes, found that no motion was seen at all unless the duration of each frame exceeded 80 ms (Mather 1988). We decided to include random background dots in subsequent experiments because pilot data obtained without noise showed that, at longer frame durations, subjects were scoring over 70% correct, even when given just one static view of the walker. As mentioned in the General Methods section, we wished to avoid any opportunity for them to exploit such positional information.

5. EFFECTS OF DOT OMISSION ON BIOLOGICAL MOTION

The first set of experiments established the perceptual subsystem involved, but how does the visual system extract meaningful structure from the encoded movements of individual dots? There are two general theoretical approaches, one in the perceptual literature, and one in the computational literature. The perceptual approach (Cutting & Proffitt 1981) proposes that analysis of these stimuli proceeds in a hierarchical fashion. The motion of the torso is analysed first (based on the motion of shoulder and hip dots), which then acts as a frame of reference for subsequent analysis of elbow and knee movements. Finally, wrist and ankle motion is analysed, using the elbow and knee motions as frames of reference. For example, Cutting and Proffitt (1981, p. 269) state that: 'One perceives in the knee only its motion relative to the hip, rather than the more complex motion resulting from compounding this pendular motion with the elliptical motion imparted to the knee by the hip'. A different approach has been adopted in the computational literature (Webb & Aggarwal 1982). The problem has been presented as one of grouping dots together into representatives of body parts, on the basis of constraints in the image. The only successful algorithm to date utilizes the 'fixed axis assumption', which asserts that the motion of an object, or the rigid parts of an object, consists of translation and rotation about an axis that is fixed in direction for short periods of time. This assumption is correct for human body articulation, and Webb & Aggarwal have successfully applied an algorithm based on it to biological motion sequences. The main prerequisite for the application of the algorithm is that 'each rigid part in each jointed object has at least two visible points' (Webb & Aggarwal 1982, p. 117). Rigidity is a common assumption made in computational approaches to visual analysis.

An experiment was designed to test these two theories. A set of four stimulus conditions was used in which selected dots were omitted from the animation sequence that subjects viewed in a particular presentation. In the control condition, all 12 dots were visible. In the other three conditions, only eight dots were visible: either the shoulder and hip dots were omitted, or the elbow and knee dots, or the wrist and ankle dots. The perceptual theory based on hierarchical analysis would predict that performance will be worse when dots higher up in the hierarchy are omitted, as the analysis of other dots depends on their frame of reference. Thus omission of shoulder and hip dots should have the worst effect, followed by omission of elbow and knee dots. Omission of wrist and ankle dots should have least effect, as no other dots depend on them for their analysis. In contrast, the computational theory based on rigid rotation would predict that omission of elbow and knee dots will be most damaging because, without these dots, only the shoulder and hip dots move in rigid relation to each other. When shoulder and hip or wrist and ankle dots are omitted, then all remaining dots move in rigid




Figure 6. Results of an experiment to investigate the effect of selective omission of dots from the walking figure. Four different presentation conditions were used. In baseline presentations, all 12 dots were plotted, as in figure 1. In other presentations either the shoulder and hip dots, the elbow and knee dots, or the wrist and ankle dots were omitted. Bars represent mean percentage correct discrimination in each presentation condition (± 1 s.e.). In both (a) the coherence task and (b) the direction task, performance is good in all conditions except that in which wrist and ankle dots were omitted, contradicting predictions of both current models.

relation to each other. Omission of wrist and ankle dots should be least damaging because fewest rigid connections are removed. Six subjects each performed in four sessions, two requiring coherence discriminations and two requiring direction discriminations (session order was randomized). Within a session there were ten presentations of each condition, order randomized. Frame duration was 36 ms, with no inter-frame interval. Dot intensity was 100 cd m^{-2} , background intensity 20 cd m^{-2} .

Results are shown in figure 6. For both tasks, performance was virtually unaffected by omission of either shoulder and hip or elbow and knee dots, but dropped to chance levels when wrist and ankle dots were omitted. The effect of presentation condition was statistically significant ($p < 0.01$). In defence of the fixed-axis model, one could argue that when some dots are omitted from the display the system attempts to make rigid connections between the remaining dots by using the fixed-axis constraint, and this new set of connections determines the pattern of performance. However, when wrist and ankle dots are omitted, 'correct' rigid connections can still be made between nearest pairs of remaining dots (shoulder to elbow, and hip to knee), but performance is near to chance. The




Figure 7. Results of a control experiment to assess direction discrimination in conditions which give little or no impression of biological motion. Left-hand bar, mean percentage correct in baseline condition containing 12 dots, appropriately positioned; middle bar, results when only a single dot is presented (the ankle dot); right-hand bar, results when all 12 dots are presented, but the vertical position of each dot varied randomly within the viewing frame from trial to trial. Performance is good only in the condition containing intact biological motion displays.

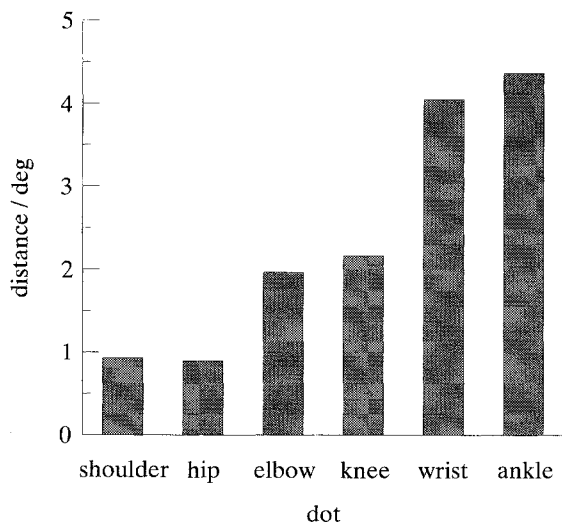


Figure 8. Total distance travelled by each dot during a single step cycle of the walker display, computed by summing the frame-to-frame displacements of each dot during the 40-frame sequence.

distance between shoulder and wrist dots, and between hip and ankle dots, varies during the step cycle. So, when elbow and knee dots are omitted, rigid connections between the shoulder and wrist dots, and the hip and ankle dots, cannot be sustained unless it is assumed that the person is moving in a way which violates the fixed axis constraint (arms and legs remaining rigid and swinging out to the side of the body as well as back and forth). Yet performance in this condition is very good.

The importance of wrist and ankle dots revealed in this experiment is consistent with results of the earlier experiment which manipulated frame increment. At larger frame increments, all dots except wrist and ankle dots still underwent short frame-to-frame displacements (figure 4), yet performance deteriorated. Data contradict both current theories of biological motion,



Figure 9. A composite view of dot positions over the 40-frame step cycle. Only dots on one side of the body are shown, in every second frame, for clarity.

and therefore suggest that a major revision is required in current conceptions of how the system analyses these stimuli. Analysis does not proceed hierarchically, nor does it require the presence of rigid motion in all dots.

6. CONTROL OBSERVATIONS

Perhaps there is sufficient information in the displays used in these experiments for subjects to make consistent discriminations without gaining an impression of biological motion. Several measures were adopted to minimize the possibility that other cues could be used (random background dots, variation in the horizontal position of the walker). As a further check, a control experiment was done to assess discrimination performance in displays which gave little or no impression of biological motion. Performance in a baseline condition containing all 12 dots was compared with performance in two other conditions. In one condition only one dot was visible, an ankle dot. In the other condition all 12 dots were visible, but in each trial a random vertical offset was added to the position of each dot for the duration of the animation sequence (constrained so that the dots always fell inside the viewing frame). This manipulation preserved the motion trajectory of each dot, and any cues it might offer, but destroyed the relations between the dots which defined the structure of the moving figure.

Six subjects participated, viewing ten-frame animation sequences in a direction discrimination task. Frame duration was 33 ms, with 0 ms inter-frame interval. All other parameters matched those in previous experiments. Results are shown in figure 7. Performance in the baseline condition was significantly

higher than in the other conditions ($p < 0.0001$), suggesting that reliable discrimination is closely associated with intact biological motion perception. Note that performance in the single dot and unstructured conditions reached 60%. Performance in other experiments rarely dropped below this level, suggesting that sufficient cues are available to support performance slightly above chance levels even in the absence of biological motion percepts.

7. GENERAL DISCUSSION

Why are dots at the extremities so important? An obvious difference between sets of dots in the figure is in the amount of movement. Perhaps extremities are most important because they move furthest. We computed the total distance travelled by each dot during a single step cycle. Data are plotted in figure 8. Shoulder and hip dots move relatively little, elbow and knee dots move more, and wrist and ankle dots move most. However, distance travelled seems inadequate as an explanation for the pattern of discrimination: omission of elbow and knee dots had no effect on discrimination (figure 6), despite the fact that they move twice as much as the shoulder and hip dots. A more profitable approach may be to look at how the dots move. Figure 9 plots the positions of all dots on one side of the body in every second frame of the animation sequence (other dots and frames are omitted for clarity). The movements of shoulder and hip dots are relatively simple and symmetrical, and not unique to a walking human. Their trajectories could be created by the ends of oscillating rods. Wrist and ankle dot motions are more informative because they reflect the resultant effect of several moving body parts, and are therefore very characteristic of a walking human. It is difficult to imagine what else could give rise to their complex, asymmetrical trajectories. The visual system

may rely heavily on detecting such characteristic movement patterns during recognition of moving images, rather than on constructing a full structured representation of the body. A similar theoretical distinction between characteristic patterns and structural descriptions can be found in the literature on recognition of static patterns (see, for example, Perrett & Harries 1988).

This research was supported by the SERC, under the Image Interpretation Initiative.

REFERENCES

- Baker, C. & Braddick, O. 1985 Temporal properties of the short-range process in apparent motion. *Perception* **14**, 181–192.
- Carlsoo, S. 1972 *How man moves*, London: Heinemann.
- Chubb, C. & Sperling, G. 1988 Drift-balanced stimuli: a general basis for studying non-Fourier motion perception. *J. opt. Soc. Am.* **A5**, 1986–2006.
- Cutting, J. 1978 A program to generate synthetic walkers as dynamic point-light displays. *Behavior Res. Meth. Instrum.* **10**, 91–94.
- Cutting, J. & Proffitt, D. 1981 In *Intersensory perception and sensory integration* (ed. R. Walk & M. Pick), pp. 249–273. New York: Plenum Press.
- Doshier, B., Landy, M. & Sperling, G. 1989 Kinetic depth from optic flow. I. 3D shape from Fourier motion. *Vision Res.* **29**, 1789–1813.
- Johansson, G. 1973 Visual perception of biological motion and a model for its analysis. *Percept. Psychophys.* **14**, 202–211.
- Mather, G. 1988 Temporal properties of apparent motion in subjective figures. *Perception* **17**, 729–736.
- Mather, G. 1989 Early motion processes and the kinetic depth effect. *Q. Jl exp. Psychol.* **41A**, 183–198.
- Perrett, D. & Harries, M. 1988 Characteristic views and the visual inspection of simple faceted and smooth objects: ‘tetrahedra and potatoes’. *Perception* **17**, 702–720.
- Webb, J. & Aggarwal, J. 1982 Structure from motion of rigid and jointed objects. *Artif. Intell.* **19**, 107–131.

Received 9 April 1992; accepted 27 May 1992




Figure 4. Computed mean frame-to-frame displacement (over the full step cycle) of each dot in the walking figure, as a function of frame increment. The three bars for each dot represent frame increments of 1 (pale shading), 2 (unshaded) and 4 (dark shading) frames.

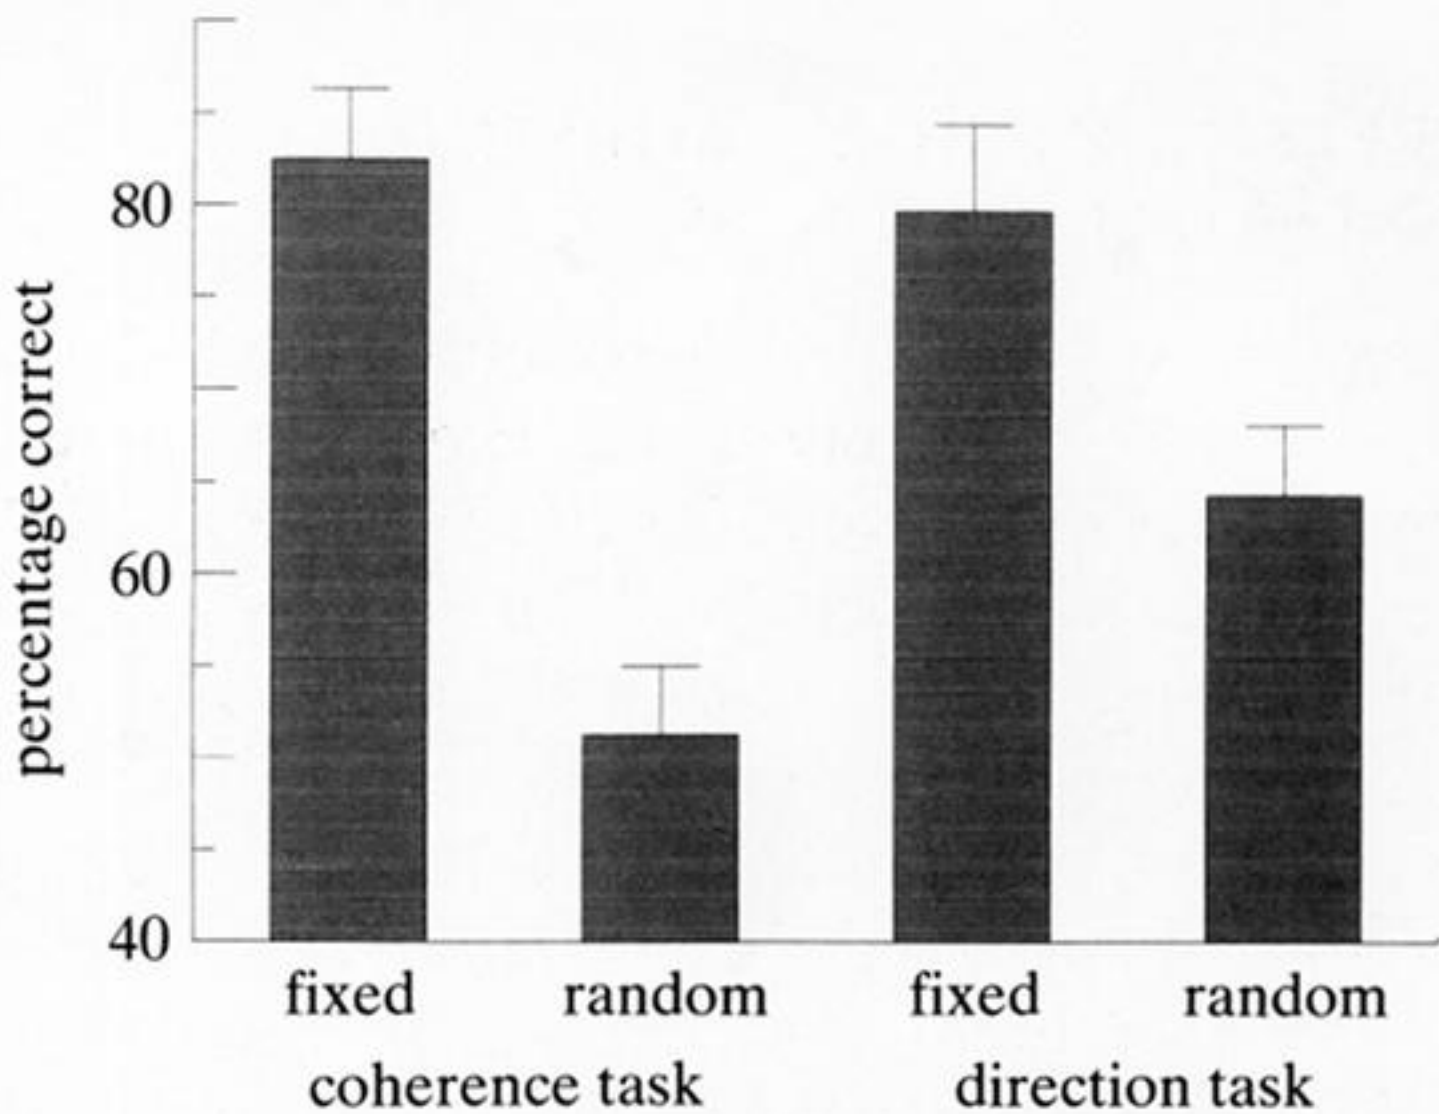


Figure 5. Results of an experiment to investigate the effect of random reversals in dot contrast on perception of biological motion. Each bar represents the mean percentage correct for five subjects, with 40 trials per subject per condition. Vertical bars represent standard errors. The two left-hand bars depict results from the coherence task, and right-hand bars depict results from the direction task. Results show that performance in both tasks was good when dot contrast against the background remained fixed, but poor when contrast reversed randomly, as expected if perception is mediated by low-level intensity-based motion mechanisms.




Figure 6. Results of an experiment to investigate the effect of selective omission of dots from the walking figure. Four different presentation conditions were used. In baseline presentations, all 12 dots were plotted, as in figure 1. In other presentations either the shoulder and hip dots, the elbow and knee dots, or the wrist and ankle dots were omitted. Bars represent mean percentage correct discrimination in each presentation condition (± 1 s.e.). In both (a) the coherence task and (b) the direction task, performance is good in all conditions except that in which wrist and ankle dots were omitted, contradicting predictions of both current models.




Figure 7. Results of a control experiment to assess direction discrimination in conditions which give little or no impression of biological motion. Left-hand bar, mean percentage correct in baseline condition containing 12 dots, appropriately positioned; middle bar, results when only a single dot is presented (the ankle dot); right-hand bar, results when all 12 dots are presented, but the vertical position of each dot varied randomly within the viewing frame from trial to trial. Performance is good only in the condition containing intact biological motion displays.




Figure 8. Total distance travelled by each dot during a single step cycle of the walker display, computed by summing the frame-to-frame displacements of each dot during the 40-frame sequence.