Robust tracking and posture description for laboratory rodents using active shape models

C. J. TWINING, C. J. TAYLOR, and P. COURTNEY
University of Manchester, Manchester, England

We are in the process of developing an automated image analysis system, which uses deformable models of shape, learned from image examples, to interpret video images of rodents. Active shape models provide a compact description of the shape of the animal in a way that enables the postures that differentiate various behaviors to be distinguished. They also model the image profile across the shape boundary. We show how these features allow automatic, robust segmentation of the explicit object of interest. Rather than just detecting movement or changes from background in the image, the system can focus on objects that are of the correct shape and appearance. The modeling of the image profiles also allows the system to distinguish between the actual animal and image artifacts. We show how these techniques are being extended to extract postural information, which can then be integrated with positional data to produce a model of behavior.

Rats and mice are widely used laboratory animals. Their patterns of behavior have been categorized, and the typical postures that constitute these categories have been described, for both the single rat (Renner & Seltzer, 1991; Timmermans, 1978) and the mouse (Eisenberg, 1968; Van Abeelen, 1963). The postures involved in social interaction have also been described for the rat (Grant & Mackintosh, 1963; Timmermans, 1978), as has the ontogenical development of posture (Geisler, Westerga, & Gramsbergen, 1993). Researchers are interested in how these patterns of behavior are modified—through exposure to pharmacological agents, in the context of drug testing and development, and through the generation of novel genotypes and, hence, novel phenotypes in mutagenesis programs. Both of these activities are typically large-scale, with the use of multiple dose regimes and multiple animals to generate reliable statistics in pharmacological research and through the large number of mutants produced in random mutagenesis, which are then screened for possibly useful new animal models (e.g., Brown, Stanford, & Schellinck, 2000; Rogers et al., 1997).

Behavioral modification can be measured for both conditioned and unconditioned responses. Unconditioned response to a standardized environment, which requires no pretest training period, is potentially less time consuming and is the case we will consider here.

A number of test paradigms have been developed, such as the elevated plus-maze (Hogg, 1996; Lister,

1987; Pellow & File, 1986), the elevated zero-maze (Shepherd, Grewal, Fletcher, Bill, & Dourish, 1994), the open arena, the dark/light box, social interaction analysis, and so forth. As the methodology has been developed, further novel tests have been suggested in an attempt to increase the discriminatory ability of such tests, such as the T-maze (Zangrossi & Graeff, 1997) and the canopy stretched-attend test (Grewal, Shepherd, Bill, Fletcher, & Dourish, 1997). However, it has become apparent that simple measures, such as general locomotor activity or area entry frequencies in elevated maze tests, do not have an unambiguous correlation with such quantities as level of *anxiety*, and that a more complete ethological analysis, which includes measures of *risk-assessment* indicators, such as stretched-attend postures and head dipping, is needed (Rodgers & Dalvi, 1997; Weiss, Wadsworth, Fletcher, & Dourish, 1998).

In the measurement of murine behavior, there are various indirect automated methods available. Infrared light beam interruptions can be used to monitor locomotor activity, area entries, and the occurrence of such activities as rearing. Balance devices have been developed that translate the movement patterns of a test cage into defined behavioral patterns of the animal (Schlingmann, Van de Weerd, Baumans, Remie, & Van Zutphen, 1998). However, these and similar methods tend to be limited in both their range of applicability and the range of behaviors that can be detected. A more general approach involves the use of video cameras; data can then be analyzed either in real time or from videotape. There are automated video-tracking systems available (e.g., Sams-Dodd, 1995), which can monitor either the position of the center of mass of each animal or the movement of colored markers attached to animals (e.g., Spinka, Sustr, & Newberry, 1999). However, data collection on specific behaviors usually relies on a human observer keying in annotation.

This work was supported by BBSRC Grant BIO10484 under the Bioinformatics Initiative. The authors thank C. Dourish and his colleagues from Vernalis Ltd. for sharing with us their experience in this area and for providing the video data used in this paper. Correspondence concerning this article should be addressed to C. J. Twining, Imaging Science and Biomedical Engineering, Stopford Building, University of Manchester, Oxford Road, Manchester M13 9PL, England (e-mail: carole.twining@man.ac.uk).

Figure 1. Rat in the zero maze.

This is labor intensive, subjective, and prone to operator bias. There is also a natural limit to the amount of annotated information that can be added. Typically, observers record either the occurrence of a behavior at a certain point in time (e.g., stretched-attend, line crossing) or the beginning and end of an activity, such as grooming. But there is no simple way to attach quantitative measures to such things as exactly how a given stretched-attend was performed. Image analysis techniques offer the prospect of an automated system that can provide a continuous set of quantitative measures of behavior.

IMAGE DATA

We will consider videotape of adult male Sprague-Dawley rats in a zero-maze, with dimensions as in Shepherd et al. (1994). The maze consists of an elevated annular platform, marked with eight equally spaced reference lines and divided into open and closed quadrants, the closed quadrants having semitransparent red Perspex walls. The source of illumination is either normal white or dim red lighting. Individual images are digitized from monochrome VHS videotape at a resolution of 768×576 pixels. As can be seen in Figure 1, there are several points to note about these images from the point of view of image analysis. The background to the animal and the appearance of the animal itself seem to vary, owing to the colored walls. There are image artifacts, such as reflections and shadows. Finally, the animal occupies a relatively small proportion of the frame. Ideally, a system should be able to provide robust automated tracking of the animal, in that it should not be confused by these

image artifacts and should be able to automatically locate the animal, despite its small apparent size.

Previous work on extracting image information relating to posture (Heeren & Cools, 2000; Noldus, 2000; Rousseau, Van Lochem, Gispen, & Spruijt, 2000) has relied on first extracting a binary image, containing either a connected region or an outline. This is done by calculating the difference between the image and a reference background image. Filtering and thresholding then reduce this difference image to a binary image. There are, however, drawbacks to this type of approach. It can be difficult to distinguish between the animal and the image artifacts by using a simple threshold, since both represent differences from the background image. Also, such approaches extract only a relatively small proportion of the available image information relating to posture.

Image Preprocessing

Interlacing of images digitized from the videotape display is caused by the animal's moving in the interval between the scanning of the two half frames. By taking one of the half frames, consisting of alternate lines of the image, a bilinear interpolation is performed to reconstruct the full image. The difference between this reconstructed image and a reference background image is then calculated, and a binary image is extracted. An image mask is then created by blurring the edges of the binary image in the region of interest. Data in the unmasked region are then replaced by data from the original image. This process is illustrated in Figure 2.

Note that we do not rely on this image-differencing procedure to extract the shape of the animal; it is used to re-

Figure 2. Left: the original image with the region of interest shown. Center: the blurred mask constructed from the difference image. Right: the final output image.

move those parts of the image that are truly a constant part of the background. We hence obtain a signal not only from the animal, but also from its reflection and shadows. Replacing the unmasked region with data from the original image allows us to retain the information that will then enable us to distinguish between the animal and these image artifacts.

ACTIVE SHAPE MODELS AND MODEL-BASED IMAGE INTERPRETATION

Active shape models (ASMs; Cootes, Hill, & Taylor, 1994; Cootes & Taylor, 2000) are combined deformable models of object shape and local gray-level appearance. These and related model-based techniques have been applied to a range of difficult image analysis problems, such as locating, interpreting, and measuring structures in medical images (Cootes et al., 1994; Redhead, Kotcheff, Taylor, Porter, & Hukins, 1997; Solloway, Hutchinson, Waterton, & Taylor, 1997), tracking and performing mea-

surements on the nematode worm *Caenorhabditis elegans* (Butcher, Cootes, Courtney, Gill, & Lithgow, 1999), locating and tracking farm animals (Bulpitt, Boyle, & Forbes, 2000; Magee & Boyle, 1999; Marchant & Onyango, 1995; Sumpter & Bulpitt, 1998), and tracking pedestrians (Baumberg & Hogg, 1994). The important defining feature of such model-based approaches is that the model is first trained on a set of example images containing instances of the object of interest. It is then able to interpolate between these examples to produce a continuous range of valid instances of the object. Thus, when presented with a novel image, a model-based system can restrict its interpretations to only those that contain a valid instance of the object, in terms of both object shape and object appearance.

ASM Training: Shape

To train a model, each of a set of M training images is labeled with a set of n landmark points (e.g., Figure 3). Some of these points are placed on specific features; for

Figure 3. Example of a marked-up image.

Figure 4. Progression of a search, showing (from the left) the starting position and the results after 1, 5, and 10 iterations.

example, in the case of the rat, the nose (Point 1), ears (Points 4–8 and 27–31), or the base of the tail (Points 17 and 18), whereas others are spaced around the outline.

Given the x - and y -coordinates of the landmarks for image α , we then concatenate these to produce a vector:

$$\mathbf{r}^\alpha = (x_1^\alpha, \dots, x_n^\alpha, y_1^\alpha, \dots, y_n^\alpha).$$

Aligning the marked-up examples, using Procrustes analysis (Goodall, 1991), eliminates the four degrees of freedom associated with the pose of the point set (i.e., scale,

Figure 5. Variation about the mean shape for the first three modes, for -3 , -1.5 , 0 , 1.5 , and 3 standard deviations.

Figure 6. The values of the first two modes for the training set. Gray triangles, turning counterclockwise; black triangles, walking counterclockwise; gray circles, grooming; black circles, walking clockwise; squares, turning clockwise.

translation, and rotation). Each example shape is then represented by an aligned coordinate vector $\tilde{\mathbf{r}}^\alpha$. Since image points tend to move in a correlated manner, the set of data points representing the training set will have some structure, clustered about the mean shape $\mathbf{r}^0$. To analyze this structure, we apply principal component analysis (Anderson, 1984). First, we construct the covariance matrix C :

$$C_{ij} = \frac{1}{M} \sum_{\alpha=1}^M (\tilde{r}_i^\alpha - r_i^0)(\tilde{r}_j^\alpha - r_j^0).$$

We then solve for the set of eigenvectors and eigenvalues of C :

$$C_{ij} v_j^a = \lambda^a v_i^a, \text{ with } \mathbf{v}^a \cdot \mathbf{v}^a = 1 \forall a.$$

The set of eigenvectors $\{\mathbf{v}^a\}$, ordered in terms of non-increasing eigenvalue, represents the set of principal directions associated with the data points, with the associated eigenvalues being the variance of the data set along that direction. Only the first t , say, of these directions will have a significant variance; these directions represent the first t modes of variation of the data set about the mean shape. Each point in the data set can then be decomposed into these modes thus:

$$b_a^\alpha \equiv b_a(\tilde{\mathbf{r}}^\alpha) = \frac{1}{\sqrt{\lambda^a}} P_{ai}^T (\tilde{r}_i^\alpha - r_i^0) \text{ for } a = 1 \text{ to } t,$$

where P is the $(2n \times t)$ matrix of the first t eigenvectors. The vectors $\mathbf{b}^\alpha$ lie in a t -dimensional space, which we call *shape space*. Each of the components b_a^α , taken across

the data set, has unit variance and zero mean. Given an arbitrary point in shape space, we can then construct the linear mapping back to the space of aligned point sets thus:

$$\tilde{\mathbf{r}}_i(\mathbf{b}) = r_i^0 + \sum_{a=1}^t \sqrt{\lambda^a} P_{ia} b_a.$$

The difference between an aligned shape vector $\tilde{\mathbf{r}}^\alpha$ and its reconstruction $\tilde{\mathbf{r}}(\mathbf{b}^\alpha)$ is minimal, provided the number of modes, t , is large enough.

Given a novel point set $\mathbf{r}$, it can then be represented by its four pose parameters, calculated by aligning it to the mean shape, and its shape vector, calculated as above.

For a linear ASM, the space of valid shapes (the valid shape region) is defined to be the set of all shapes whose corresponding shape vector $\mathbf{b}$ has a modulus less than some given value. This is equivalent to the assumption that the distribution of training shapes can be modeled by a multivariate Gaussian, where surfaces of equal probability density become spheres in shape space.

The first advantage of this method of modeling shape is that it is object specific; we model only those modes of shape variation that are actually seen, as opposed to, say, modeling an arbitrary closed contour, as in Heeren and Cools (2000). Also, the set of landmark points is not confined to modeling just a simple contour; we could have included landmarks for features, such as the eyes, lying within the outline of the animal, had such features been visible in our training set. This approach to the modeling of shape does not make any assumptions about which measurements of the shape are important in terms of distinguishing the postures of interest, in contrast with

Figure 7. Examples of the active shape model searching on images from the training set. From the left: (1) correct area 78%, incorrect area 3%; (2) correct 91%, incorrect 8%; (3) correct 92%, incorrect 12%; (4) correct 94%, incorrect 7%.

such approaches as that of Rousseau et al. (2000). Rather, we let the mathematical analysis of the variation actually seen determine which sets of combined measurements (i.e., the modes) are important.

ASM Training: Local Gray Level

The final step in constructing an ASM is to include information about the image appearance in the neighborhood of each landmark point. In each example image, we construct the normals to the boundary at each landmark point. We then sample the gray levels along these normals. This set of gray-level profiles at each landmark can then be modeled in an analogous way to shape, leading to a separate gray-level model for each landmark. Further details can be found in Cootes and Taylor (1993) and the references therein.

ASM Search

We now explain briefly how an ASM is used to search an unseen image. Suppose we have some initial estimate of the location of the animal, either from motion detection, say, or from a successful search on a previous frame. Then, for each landmark point, we sample the gray levels in the neighborhood of that point to see whether we can locate a new position for that landmark that will give a better fit to the gray-level model. This gives us a new, suggested set of positions for the landmarks. We use the shape model to represent this set of suggested points as closely as possible, while constraining our shape to lie within the valid shape region for our model. This gives a new set of landmark positions. This process is iterated until it converges onto the animal.

The progression of an example search is shown in Figure 4. Starting from a set of points that roughly corre-

sponds to the position and orientation of the animal, it can be seen that the model points rapidly converge to the correct outline. Note that the weaker image signals from the extended forepaw and the image artifact are correctly ignored, since they are not included in the model.

This method of interpreting the image has several desirable properties. First, the information from the gray-level models tells us how good a local fit we have achieved. So, for example, if some landmarks have converged on a reflection edge, rather than on the edge of the actual animal, this will be indicated by their relatively poor fit to their gray level models. Second, if the gray levels indicate a good fit, we know that it is then highly likely that we have achieved a fit to the actual animal, since we have constrained ourselves to a valid shape throughout the search. Finally, a description of the shape of the animal has been extracted automatically, since this is encoded in the shape vector corresponding to the resultant set of landmarks.

EXPERIMENTS

Model Building

A set of 106 images of a rat in the zero-maze, similar to that shown in Figure 1, was chosen as the training set. It consisted of images digitized from sequences of the rat walking around the maze, both clockwise and counterclockwise, turning on the maze, and grooming. Care was taken to ensure that the set contained images with the animal in all positions around the maze, in both the open and the closed quadrants, and also images that contained significant image artifacts (e.g., as in Figure 2, which contains reflections, shadows, and also a stripe along the back of the animal, caused by the top edge of

Figure 8. Histogram of the distribution of the size of the overlap (black) and nonoverlap (gray) areas from search, as a percentage of the area of the marked-up region for the 45 examples from the training set.

the Perspex wall). These images were then preprocessed as described above.

Each image was marked up by hand with a set of 33 landmark points, as in Figure 3. It was decided not to include the tail in landmarking for several reasons. First, points along the length of the tail are not particularly distinctive in terms of the local gray-level appearance. Second, it was observed that the tail often disappeared from several video images, owing to its rapid motion.

An ASM was built from this training set. It was found that 20 modes of variation were sufficient to encompass 98% of the variation in the training set.

It is interesting to compare this figure with the 17 degrees of freedom (five pairs of complex harmonics less 3 degrees of freedom relating to orientation and size) used by Heeren and Cools (2000) when describing the rat outline in terms of Fourier descriptors. However, they explicitly decided not to represent such features as the ears, so that they were modeling a smoother contour. Also, their method used only one point of correspondence (the base of the tail). An ASM reconstruction of shape explicitly contains more information than just the contour, since points either directly correspond to or are spatially related to specific features, so it is reasonable that our representation requires slightly more degrees of freedom.

In Figure 5, we illustrate examples of the shapes that can be generated by the model. Starting from the mean shape, we vary each of the first three modes in turn. From this, we can see that the main action of the first mode

consists of the rat's bending, the second corresponds to stretching and compression. The third mode is not so clear, although it can be seen that, in one direction at least, it is related to the posture adopted during grooming. The other higher modes correspond to more complicated changes of shape.

We now consider the distribution of the training set in shape space. Figure 6 shows the set of shapes from the training images plotted in the first two dimensions of shape space.

The first point to note is that these two modes alone are enough to distinguish between actions such as walking, turning, and grooming and to establish the direction of walking or turning. The points corresponding to turning were actually taken from sequences including a turn. Hence, those corresponding to the beginning or the end of a turn sequence lie close to the shape observed in walking, whereas others, corresponding to the point of maximum bending, lie at the extrema of the distribution.

However, an important point to note is that the distribution of the data is nonlinear; a valid shape region that is a sphere either will not include all the training data or, if it does, will also include invalid shapes. The effects of this will become apparent in the next section.

Searching with the ASM

The model was first tested on a set of 45 images, taken from the training set, of walking and grooming. To estimate the location and orientation of the animal, the

Figure 9. Examples of the active shape model searching on unseen images of walking, showing automatic initialization and the output of the search process.

image was first thresholded at a level corresponding to 85% of the maximum intensity. Note that we are *not* relying on the thresholding to give us the exact outline, but only the brightest pixels in the image, which we can reliably assume to be from the animal, and not from the image artifacts.

The mean position, the principal axis, and the second and third centered moments along this axis were then calculated, both for the binary image and for the mean shape from the model. Matching these parameters then gave us the initial pose for the mean shape. It was found that, in all 45 cases, we were able to correctly establish the orientation of the animal. Examples of search results are shown in Figure 7, together with the initial position of the mean shape.

To evaluate the efficacy of the search, we compared the region defined by the mark-up points with that found in search. The areas of overlap and nonoverlap with the marked-up region were calculated as a fraction of the area of the marked-up region. These results are given in Figure 8. The means and standard deviations of overlap and nonoverlap areas are $89.0\% \pm 4.4\%$ and $8.9\% \pm 5.2\%$, respectively.

It was found that, in general, the results for the images of grooming tended to be slightly worse than those for walking. This is understandable, given that the initial estimate for the position of the mean shape tends to be less accurate in these cases, as can be seen from the examples given in Figure 7.

The system was also tested on a set of 63 unseen images of walking, using the method for the initial estimate, as above. In all 63 cases, the orientation was correctly established. Examples of the search results are given in Figure 9. It can be seen that the search gives a reasonable

estimate of the outline, even in cases in which there are significant image artifacts present.

There are, however, some problems with constraining the shape during search. This is not unexpected, given the observation above that the distribution of the training set was seen to be nonlinear. In Figure 10, we give an example of a failed linear search, showing self-intersections, plus the result of searching on the same image, from the same initial position, using a trial nonlinear ASM. This was constructed by performing a density estimate on the distribution of training shapes (Cootes & Taylor, 1999). The valid shape region is then defined as all points lying within a particular contour. An example of a valid shape region so defined, in the case of the first 2 modes only, is given in Figure 11. In the actual nonlinear ASM, the valid shape region is defined over all 20 modes simultaneously.

DISCUSSION

We have shown that an ASM can perform well in search, giving a robust, automated method of image segmentation and, hence, tracking. The model was automatically initialized on single images (i.e., without the use of motion cues), by estimating the orientation and location of the animal. We obtained correct estimates of the orientation on all of the 108 images considered. Given this initialization, we have shown that the ASM can reliably extract the outline of the animal. Furthermore, we have demonstrated that the ASM provides a useful compact description of shape, as can be seen by the differentiation of the various postures/behaviors included in the training set. However, the point-to-feature correspondence from the search results is generally not precise enough to allow reliable extraction and classification of posture. It

Figure 10. Example of a failed linear search, with self-intersections (top right of upper image), and the same search performed using a trial nonlinear active shape model.

was observed that the distribution of the training set in shape space was nonlinear; this is often the case in objects that display modes related to bending. This means that the linear ASM does not always constrain the shape in an appropriate manner during search, as is shown in Figure 10, where it allows self-intersections to occur. We have also presented preliminary results from a nonlinear ASM, designed to address this problem.

It should be noted that all the experimental results presented above relate to searching on single images. In practice, a system would be searching on a sequence of images, taking the search result from a previous frame as additional data for use in estimating the initial position and the posture for the next frame. This should improve the point-to-feature correspondence in search results,

since it will effectively allow us to follow a feature between frames and will allow us to build explicit models of motion. As regards the expected computational efficiency of such a system, we note that related techniques applied to the task of tracking pedestrians (Baumberg & Hogg, 1994) were able to achieve frame rates of up to 30 frames per second. However the greater range of shape variation seen for rats (i.e., 20 linear modes used in our model, as opposed to 14 for pedestrians [Baumberg & Hogg, 1994] or 5 for a flock of ducks [Sumpter & Bulpitt, 1998]), plus the associated nonlinearity of the shape space, mean that such a high frame rate may not be achievable. However, we note that Heeren and Cools (2000) were able to obtain a useful temporal analysis of behavior (grooming), using a sampling rate of 10 frames per second.

Figure 11. An example valid shape region (bold line) calculated across the first 2 modes of the training set. Shading indicates the 20-dimensional density estimate integrated across the other 18 modes.

Of previous work on extracting postural data for rats, based on the difference-from-background procedure, some (e.g., Noldus, 2000) did not utilize the whole outline, but only a small set of discrete points, such as the nose, tail base, and center of gravity. Others, which used the whole outline (e.g., Heeren & Cools, 2000; Rousseau et al., 2000), considered only the case of animals in an open arena. Image data from the zero-maze, with the image artifacts and the variations in the appearance of the animal, present a more difficult image analysis problem. It can be seen that a single, fixed threshold value would not, in general, be sufficient to segment these images. To overcome this increase in image complexity, ASMs effectively utilize more of the available image data. The basic difference-from-background procedure relies on a single pass through the image to construct the initial estimate of the outline. Refinement of this outline then typically proceeds via a purely morphological analysis (i.e., erosion and dilation, as in Heeren & Cools, 2000), relying on the low-level concepts of connectedness and smoothness. In contrast, refinement of the outline during ASM search proceeds with continued reference both to the local image appearance and to the known, previously learned shape variation of the target object. Furthermore, ASMs also possess the advantage that, for each proposed outline, we have an estimate of the reliability of the segmentation. This can be calculated from the quality of fit of the image profile models at each landmark point. Such a reliability measure can then be used to automatically detect and reject the results of erroneous

searches, allowing reinitialization of the search procedure on the next image frame.

REFERENCES

- ANDERSON, T. W. (1984). *An introduction to multivariate statistical analysis* (2nd ed.). New York: Wiley.
- BAUMBERG, A., & HOGG, D. (1994). An efficient method for contour tracking using active shape models. In *Proceedings of the 1994 IEEE Workshop on Motion of Non-Rigid and Articulated Objects* (pp. 194-199). Los Alamitos, CA: IEEE Computer Society Press.
- BROWN, R. E., STANFORD, L., & SCHELLINCK, H. M. (2000). Developing standardised behavioral tests for knockout and mutant mice. *Institute for Laboratory Animal Research Journal*, **41**, 163-174.
- BULPITT, A. J., BOYLE, R. D., & FORBES, J. M. (2000, August). *Monitoring behavior of individuals in crowded scenes*. Paper presented at Measuring Behavior 2000, 3rd International Conference on Methods and Techniques in Behavioral Research, Nijmegen, The Netherlands.
- BUTCHER, D., COOTES, T. F., COURTNEY, P., GILL, M., & LITHGOW, G. J. (1999). Model-based image analysis of a model organism for life science research. In *Proceedings of the IEE 7th International Conference on Image Processing and Its Applications* (pp. 392-396). London: Institution of Electrical Engineers.
- COOTES, T. F., HILL, A., & TAYLOR, C. J. (1994). The use of active shape models for locating structure in medical images. *Image & Vision Computing*, **12**, 355-366.
- COOTES, T. F., & TAYLOR, C. J. (1993). Active shape model search using local grey-level models: A quantitative evaluation. In J. Illingworth (Ed.), *Proceedings of the 4th British Machine Vision Conference, 1993* (pp. 639-648). Guildford, U.K.: BMVA Press.
- COOTES, T. F., & TAYLOR, C. J. (1999). A mixture model for representing shape variation. *Image & Vision Computing*, **17**, 567-573.
- COOTES, T. F., & TAYLOR, C. J. (2000). *Statistical models of appearance for computer vision* [On-line document]. Available: www.isbe.man.ac.uk/~bim/Models/app_model.ps.gz.
- EISENBERG, J. F. (1968). Behavior patterns. In J. A. King (Ed.), *Biology*

- of *Peromyscus (rodentia)* (pp. 451-495). Washington, DC: American Society of Mammologists.
- GEISLER, H. C., WESTERGA, J., & GRAMSBERGEN, A. (1993). Development of posture in the rat. *Acta Neurobiologiae Experimentalis*, **53**, 517-523.
- GOODALL, C. (1991). Procrustes methods in the statistical analysis of shape. *Journal of the Royal Statistical Society: Series B*, **53**, 285-339.
- GRANT, E. C., & MACKINTOSH, J. H. (1963). A description of the social postures of some laboratory rodents. *Behaviour*, **21**, 246-259.
- GREWAL, S. S., SHEPHERD, J. K., BILL, D. J., FLETCHER, A., & DOURISH, C. T. (1997). Behavioural and pharmacological characterisation of the canopy stretched attend posture test as a model of anxiety in mice and rats. *Psychopharmacology*, **133**, 29-38.
- HEEREN, D. J., & COOLS, A. R. (2000). Classifying postures of freely moving rodents with the help of Fourier descriptors and a neural network. *Behavior Research Methods, Instruments, & Computers*, **32**, 56-62.
- HOGG, S. (1996). A review of the validity and variability of the elevated plus-maze as an animal model of anxiety. *Pharmacology Biochemistry & Behavior*, **54**, 21-30.
- LISTER, R. G. (1987). The use of a plus-maze to measure anxiety in the mouse. *Psychopharmacology*, **92**, 180-185.
- MAGEE, D., & BOYLE, R. (1999). Feature tracking in real world scenes (or how to track a cow). In *Proceedings of the IEE Colloquium on Motion Analysis and Tracking* (pp. 21-27). London: Institution of Electrical Engineers.
- MARCHANT, J. A., & ONYANGO, C. M. (1995). Fitting grey level point distribution models to animals in scenes. *Image & Vision Computing*, **13**, 3-12.
- NOLDUS, L. (2000). Digital imaging, ultra-sound measurement and infrared thermography: New techniques to refine behavioral observation. *Noldus News*, **7**, 1-6.
- PELLOW, S., & FILE, S. E. (1986). Anxiolytic and anxiogenic drug effects on exploratory activity in an elevated plus-maze: A novel test of anxiety in the rat. *Pharmacology, Biochemistry & Behavior*, **24**, 525-529.
- REDHEAD, A. L., KOTCHEFF, A. C. W., TAYLOR, C. J., PORTER, M. L., & HUKINS, D. W. L. (1997). An automated method for assessing routine radiographs of patients with total hip replacements. *Journal of Engineering in Medicine*, **211**, 145-154.
- RENNER, M. J., & SELTZER, C. P. (1991). Molar characteristics of exploratory and investigatory behavior in the rat (*Rattus norvegicus*). *Journal of Comparative Psychology*, **105**, 326-339.
- RODGERS, R. J., & DALVI, A. (1997). Anxiety, defence and the elevated plus-maze. *Neuroscience & Biobehavioral Reviews*, **21**, 801-810.
- ROGERS, D. C., FISHER, E. M., BROWN, S. D., PETERS, J., HUNTER, A. J., & MARTIN, J. E. (1997). Behavioral and functional analysis of mouse phenotype: SHIRPA, a proposed protocol for comprehensive phenotype assessment. *Mammalian Genome*, **8**, 711-713.
- ROUSSEAU, J. B. I., VAN LOCHEM, P. B. A., GISPEN, W. H., & SPRUIJT, B. M. (2000). Classification of rat behavior with an image-processing method and a neural network. *Behavior Research Methods, Instruments, & Computers*, **32**, 63-71.
- SAMS-DODD, F. (1995). Automation of the social interaction test by a video-tracking system: Behavioural effects of repeated phencyclidine treatment. *Journal of Neuroscience Methods*, **59**, 157-167.
- SCHLINGMANN, F., VAN DE WEERD, H. A., BAUMANS, V., REMIE, R., & VAN ZUTPHEN, L. F. M. (1998). A balance device for the analysis of behavioural patterns of the mouse. *Animal Welfare*, **7**, 177-188.
- SHEPHERD, J. K., GREWAL, S. S., FLETCHER, A., BILL, D. J., & DOURISH, C. T. (1994). Behavioural and pharmacological characterisation of the elevated zero-maze as an animal model of anxiety. *Psychopharmacology*, **116**, 56-64.
- SOLLOWAY, S., HUTCHINSON, C. E., WATERTON, J. C., & TAYLOR, C. J. (1997). The use of active shape models for making thickness measurements of articular cartilage from MR images. *Magnetic Resonance in Medicine*, **37**, 943-952.
- SPINKA, M., SUSTR, P., & NEWBERRY, R. (1999). A colorful Ethovision masquerade, or how to trace pig maneuvering automatically. *Noldus News*, **6**, 6-7.
- SUMPTER, N., & BULPITT, A. J. (1998). Learning spatio-temporal patterns for predicting object behaviour. In P. H. Lewis & M. S. Nixon (Eds.), *Proceedings of the 9th British Machine Vision Conference, 1998* (Vol. 2, pp. 649-658). Guildford, U.K.: BMVA Press.
- TIMMERMANS, P. J. A. (1978). *Social behaviour in the rat*. Unpublished doctoral dissertation, Katholieke Universiteit Nijmegen, The Netherlands.
- VAN ABELEN, J. H. F. (1963). Mouse mutants studied by means of ethological methods: 1. Ethogram. *Genetica*, **34**, 79-94.
- WEISS, S. M., WADSWORTH, G., FLETCHER, A., & DOURISH, C. T. (1998). Utility of ethological analysis to overcome locomotor confounds in elevated maze models of anxiety. *Neuroscience & Biobehavioral Reviews*, **23**, 265-271.
- ZANGROSSI, H., & GRAEFF, F. G. (1997). Behavioral validation of the elevated T-maze: A new animal model of anxiety. *Brain Research Bulletin*, **44**, 1-5.

(Manuscript received September 27, 2000;
accepted for publication May 10, 2001.)