

Automatic Tracking, Feature Extraction and Classification of *C. elegans* Phenotypes

Wei Geng*, Pamela Cosman, *Senior Member, IEEE*, Charles C. Berry, Zhaoyang Feng, and William R. Schafer

Abstract—This paper presents a method for automatic tracking of the head, tail, and entire body movement of the nematode *Caenorhabditis elegans* (*C. elegans*) using computer vision and digital image analysis techniques. The characteristics of the worm's movement, posture and texture information were extracted from a 5-min image sequence. A Random Forests classifier was then used to identify the worm type, and the features that best describe the data. A total of 1597 individual worm video sequences, representing wild type and 15 different mutant types, were analyzed. The average correct classification ratio, measured by out-of-bag (OOB) error rate, was 90.9%. The features that have most discrimination ability were also studied. The algorithm developed will be an essential part of a completely automated *C. elegans* tracking and identification system.

Index Terms—*C. elegans*, classification, computer vision, data mining, feature extraction, image processing, Random Forests, tracking.

I. INTRODUCTION

THE nematode *Caenorhabditis elegans* (*C. elegans*) is widely used for studies of nervous system function and development. *C. elegans* is a free-living worm, approximately 1 μm in length, that lives in the soil and feeds on bacteria. It has a simple nervous system containing 302 neurons, and the precise position, cell lineage, and synaptic connectivity of each of these neurons is known. However, despite its anatomical simplicity, the *C. elegans* nervous system mediates surprisingly diverse and intricate patterns of behavior. The sense organs of *C. elegans* are capable of perceiving and responding to a wide range of environmental conditions, including heavy and light touch, temperature, volatile odorants, osmotic and ionic strength, food, and other nematodes. Each of these sensory modalities in turn regulates many aspects of the animal's behavior, including the rate and direction of movement, the rates of feeding, egg-laying, defecation, and the process of mating.

Manuscript received July 17, 2003; revised January 19, 2004. This work was supported in part by the National Institutes of Health (NIH) and in part by a grant from the National Institute on Drug Abuse. Portions of this work were presented at the 2003 Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA. *Asterisk indicates corresponding author.*

*W. Geng was with the Department of Electrical and Computer Engineering, University of California at San Diego, La Jolla, CA 92093-0407 USA. He is now with I.D. Analytics Inc., San Diego, CA 92193 USA (e-mail: wei_geng@yahoo.com).

P. Cosman is with the Department of Electrical and Computer Engineering, University of California at San Diego, La Jolla, CA 92093-0407 USA.

C. C. Berry is with the Division of Biostatistics, Department of Family and Preventive Medicine, University of California at San Diego, La Jolla, CA 92093 USA.

Z. Feng and W. R. Schafer are with the Division of Biology, University of California at San Diego, La Jolla, CA 92093-0349 USA.

Digital Object Identifier 10.1109/TBME.2004.831532

Because a particular neuron can be positively identified based on its position, it is possible to eliminate the function of an individual neuron or group of neurons through single cell laser ablation. Moreover, because of their short generation time, completely sequenced genome, and accessibility to germline transformation, these animals are highly amenable to molecular and classical genetics. Thus, in *C. elegans* it is relatively straightforward to evaluate the functions of particular neurons or gene products by characterizing the effects of mutations or neuronal ablations on the animal's behavior.

This approach to understanding nervous system function relies on the ability to obtain precise quantitative measurements of behavioral abnormalities seen in mutant and cell-ablated animals. Although gross behavioral defects can often be discriminated qualitatively by simple observation, precisely defining these differences can be challenging without quantitative measurements of parameters such as the curvature and amplitude of body bends. Furthermore, many behavioral abnormalities are reliably detected only using quantitative methods. For example, some behavioral events, such as oviposition, occur on a time scale that precludes evaluation by long-time human observation [25], [28]. In other cases, qualitative differences between mutant and normal strains are not apparent to the eye, but can readily be discriminated through quantitative analysis. A number of mutants exhibiting altered egg-laying patterns or hyperactive locomotion fall into this category.

A variety of computer-driven systems for automated recording and/or analysis of *C. elegans* behavior have been described. Some systems have been designed to observe multiple animals at low magnification and track the position of each animal over time [8], [9] (it is necessary to record at low magnification to keep all tracked animals in the field of view of the microscope). Such systems make it possible to measure large-scale behavioral features such as the rate and direction of movement and the frequency of reversals in direction. However, because the animals are observed at low magnification, it is not possible to obtain more detailed information about their body posture and morphology. To gain such information, we recently developed a system designed to follow an individual animal at high magnification [1]. To keep the animal from leaving the field of view, a tracking program directs the movement of a motorized stage to maintain the worm in the center of the field. In this way, it is possible to follow the position of the animal over long time periods and comprehensively measure multiple features that define behavioral and morphological abnormalities of nematode mutants. By using 94 such features, we were able to classify representative mutant types using a binary decision tree algorithm (CART). However, although

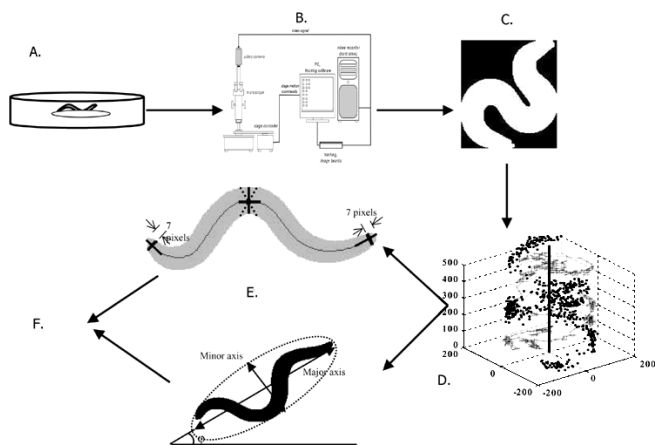


Fig. 1. (A) Fourth-stage larvae were picked the evening before the experiment and tracked the following morning on a fresh plate. (B) A data acquisition system containing a high power microscope and a stage controller was used to track and record the worm locomotive information. (C) Image processing steps removed noise and separated worm bodies from the background. (D) Head and tail tracking. (E) Feature extraction step extracted a total of 253 features from the binary and gray image sequence. (F) The data were then fed through the Random Forests algorithm for classification.

this system performed well at distinguishing visibly different mutant phenotypes, it was less effective at distinguishing types with more subtle differences.

In this paper, we present new automatic methods for feature extraction and phenotype classification based on morphological operations and on head and tail recognition. The methods presented in this paper extend our previously described system [1], [12] in several ways. First, we identify and track the head and tail separately, and extract features that characterize them. Secondly, a new ensemble learning scheme (Random Forests) is applied to this large-scale data set. Finally, using our new features and improved classifier, we have demonstrated the successful classification of a significantly larger data set consisting of wild type and 15 mutant strains.

Our approach can be divided into several well-defined stages, presented in Fig. 1. After video images acquisition, the images are first segmented to isolate the worm body from the background and remove noise and undesired components. Next, the head and tail are recognized for entire video sequences. Feature extraction is applied, to extract the useful information from the segmented objects and the head and tail locations. Finally, a classifier based on Random Forests operates on the characteristics extracted by the previous stages. Each of these steps will be presented in the next sections.

II. MATERIALS AND METHODS

A. Acquisition of the Video Images

Routine culturing of *C. elegans* was performed as described [6]. Animals were grown on standard nematode growth medium (NGM) seeded with *E. coli* strain OP50; all experiments were conducted in the presence of freshly-seeded OP50 lawns. All worms analyzed in these experiments were young adults; fourth-stage larvae were picked the evening before the experiment and tracked the following morning after cultivation at 22°. Since locomotion behavior shows reproducible and

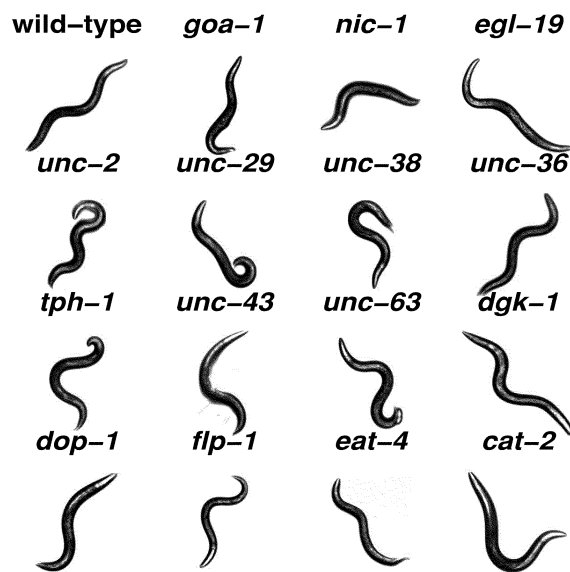


Fig. 2. Representatives of wild type and 15 mutants. Descriptions of these mutants are summarized in Table I.

stereotyped changes following the transfer of an animal to a new culture dish [15], experimental animals were transferred to new plates and tracked 5 min later to ensure a valid comparison between experiments. We used wild type worms and 15 mutants: N2-wild-type; *goa-1*(n1134) – *Goα* subunit; *nic-1*(lj22)-type1 glycosyltransferase; *unc-36*(e251)-VGCC $\alpha2/\delta$ subunit; *unc-38*(x20)-nAChR α subunit; *unc-29*(x29)-nAChR β subunit; *egl-19*(n582)-L-type VGCC $\alpha1$ subunit; *unc-2*(mu74)-non-L-type VGCC $\alpha1$ subunit; *tph-1*(mg280)-tryptophan hydroxylase; *unc-63*(x13)-nAChR α subunit; *dgk-1*(nu62)-diacylglycerol kinase; *unc-43*(e755)-CaMKII; *dop-1*(ev748)-D1 dopamine receptor; *flp-1*(yn2)-Fa-related neuropeptide; *eat-4*(ky5)-vesicular glutamate transporter; *cat-2*(e1112)-tyrosine hydroxylase.

C. elegans locomotion was tracked with a stereomicroscope mounted with a charge-coupled device video camera [1]. A computer-controlled tracker was used to maintain the worms in the center of the optical field of the stereomicroscope during observation. To record the locomotion of an animal, an image frame of the animal was snapped every 0.5 s for at least 5 min. Among those image pixels with values less than or equal to the average value minus three times the standard deviation, the largest connected component was found. The image was then trimmed to the smallest axis-aligned rectangle that contained this component, and saved as eight-bit grayscale data. The dimensions of each image, and the coordinates of the upper left corner of the bounding box surrounding the image were also saved simultaneously as the references for the location of an animal in the tracker field at the corresponding time point when the images are snapped. The stereomicroscope was fixed to its largest magnification (50 X) during operation. Depending on the type and the posture of a worm, the number of pixels per trimmed image frame varied. The number of pixels per millimeter was fixed at 312.5 pixel/mm for all worms. A total of 1596 video sequences of at least 5 min each were analyzed

TABLE I
DESCRIPTIONS OF MUTANT TYPES

Strain	Defective molecule	Description	Source
Wild-type	N/A	normal	Mendel et al., <i>Science</i> 267, 1652 (1995); Segalat, et al., <i>Science</i> 267, 1648 (1995)
<i>goa-1</i>	α subunit of G-protein G_o	hyperactive, defective male mating	Mendel et al., <i>Science</i> 267, 1652 (1995)
<i>nic-1</i>	type1 glycosyltransferase	dumpy, moves poorly	J. Kim and W. Schafer, unpublished
<i>egl-19</i>	L-type VGCC $\alpha 1$ subunit	moderate bloating, slow and floppy	Lee et al., <i>EMBO J.</i> 16, 6066 (1997)
<i>unc-2</i>	non-L-type VGCC $\alpha 1$ subunit	weak kinker, sluggish, thin	Schafer et al., <i>Nature</i> 375, 73 (1995)
<i>unc-29</i>	nicotinic receptor β subunit	weak kinker, head region stiff, moves better in reverse	Fleming et al., <i>J. Neurosci.</i> 17, 5843 (1997)
<i>unc-38</i>	nicotinic receptor α subunit	weak kinker, sluggish, slightly dumpyish	Fleming et al., <i>J. Neurosci.</i> 17, 5843 (1997)
<i>unc-36</i>	voltage-gated calcium channel (VGCC) $\alpha 2/\delta$ subunit	very slow, thin loopy at rest	Brenner, <i>Genetics</i> 77, 77 (1974)
<i>tph-1</i>	tryptophan hydroxylase	bloated, slow moving	Sze et al., <i>Nature</i> 403, 506 (2000)
<i>unc-43</i>	CaMKII	slow, lazy, slightly rippling movement	Reiner et al., <i>Nature</i> 402, 199 (1999)
<i>unc-63</i>	nicotinic receptor subunit	weak kinker, slow, inactive	Lewis et al., <i>Neurosci.</i> 5, 967 (1980)
<i>dgk-1</i>	diacylglycerol kinase	hyperactive for locomotion and foraging	Nurrish et al., <i>J. Neurosci.</i> 24, 231 (1999)
<i>dop-1</i>	D1 dopamine receptor	locomotion normal	Sanyal et al., <i>EMBO Journal</i> 23, 473 (2004)
<i>flp-1</i>	Fa-related neuropeptide	hyperactive, loopy uncoordinated movement	Nelson, et al., <i>Science</i> 281, 1686 (1998)
<i>eat-4</i>	vesicular glutamate transporter	foraging abnormal	Lee, et al., <i>J. Neurosci.</i> 19, 159 (1999)
<i>cat-2</i>	tyrosine hydroxylase	defective in food-swallowing behavior	Lee et al., <i>J. Neurosci.</i> 16, 6066 (1997)

containing 100 videos of *goa-1*, *nic-1*, *egl-19*, *unc-36*, *unc-38*, *tph-1*, *unc-63*, *unc-43*, *dop-1*, *flp-1*, *cat-2*, *eat-4*, 99 of *unc-2*, 98 of *unc-29*, and 99 of *dgk-1*. Representatives of each worm type are shown in Fig. 2 and descriptions are summarized in Table I to help the reader understand the classification problem visually.

B. Segmentation of the Worm Body

The segmentation process is presented in Fig. 3. The first operation is a local thresholding using a 5×5 moving window. The center pixel inside the moving window was assigned to 1 when the mean value of the window was less than 70% of the background pixel value or the standard deviation was larger than 30% of the mean value. Otherwise, the center pixel was assigned to 0 as background. Next, the sequential algorithm for component labeling was used to remove unwanted small objects [18].

A morphological closing operator (binary dilations followed by erosions) [14] cleaned up the spots inside the worm body. In order to avoid occasional false contours and exterior holes [formed by severe worm body bending as shown in Fig. 5(A)] being filled by excessive closing operations, we also generate a reference binary image in parallel by filling the holes that have compactness (defined as $perimeter^2/area$) greater than 25 after local thresholding. Since exterior holes tend to be round, the compactness was used to avoid filling large exterior holes (>100 pixels). Thus, two binary images were generated after local thresholding. The one with the closing operation sometimes contains excessive pixels, whereas the binary image after the hole filling operation tends to have fewer pixels than desired on other occasions. The difference between these two binary images provides a good indication of whether or not the segmentation is successful and of which binary image is better to use. Figs. 4 and 5 illustrate the segmentation process. Com-

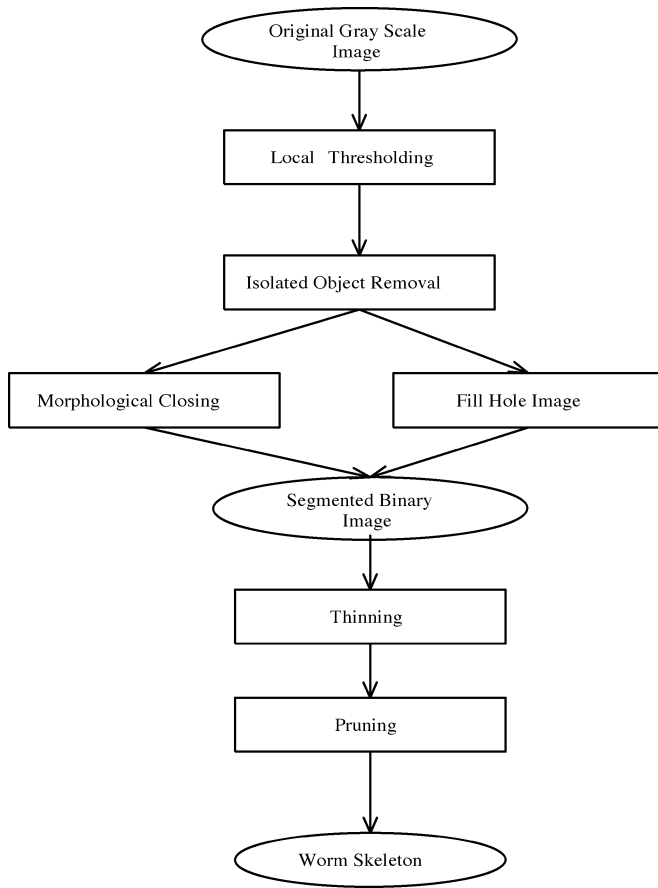


Fig. 3. General description of the segmentation process.




Fig. 4. (A) Gray level image. (B) Original binary image after local thresholding operation. (C) Binary image after closing operation. (D) Binary image after hole filling operation. (E) Difference between C and D. (F) Final binary image with the excess pixels removed.

paring Figs. 4 and 5, we see that the hole-filling result [Fig. 4(D)] is more correct than the closing result [Fig. 4(C)] for one case, but the closing result [Fig. 5(C)] is superior to hole-filling [Fig. 5(D)] in another case. By generating both hole-filling and closed versions of each image, and analyzing the difference between them, the algorithm is able to determine which pixels to include in the final binarization. For example, in Fig. 4(E), the excessive pixels caused by the closing operation are represented as the largest connected component in the difference image. The worm bodies along both sides of the arrow object's




Fig. 5. (A) Gray level image. (B) Binary image after local thresholding operation. (C) Binary image after closing operation. (D) Binary image after hole filling operation. (E) Difference between c and d. (F) Final binary image with inside crack filled.

second eigen-direction were wider than 20 pixels (the typical worm body width), indicating the existence of excessive pixels. The second eigen-direction is the direction that is perpendicular to the principal component direction, and is calculated as $\theta = \pi/2 + \tan^{-1}(2 * \sum_{k=1}^n (x_k * y_k) / (\sum_{k=1}^n x_k^2 * \sum_{k=1}^n y_k^2))$, where (x_k, y_k) are the coordinates of the pixels in the object after centering. In Fig. 5, the missing pixels inside the worm body caused by thresholding are represented as the curved object in the difference image. The worm body portions along both sides of the curved object's second eigen-direction were narrower than 20 pixels, indicating missing pixels inside the worm body.

Following binarization, a morphological skeleton was obtained by applying a skeletonizing algorithm [27]. Redundant pixels on the skeleton were eliminated by thinning. To avoid branches on the ends of skeletons, the skeleton was first shrunk from all its end points simultaneously until only two end points were left. These two end points represent the longest end-to-end path on the skeleton. A clean skeleton can then be obtained by growing out these two remaining end points along the unpruned skeleton by repeating a dilation operation [Fig. 6(A)–(D)].

C. Tracking and Head and Tail Recognition

Even though a simple tracking system was able to follow worm centroid movement, the head and tail information were not extracted in our earlier work. Because of the highly deformable nature of the worm's body, many conventional image matching and tracking algorithms do not apply to this problem. To address these problems, we have applied three spatial and temporal clues that human observers use to recognize the head and tail sections. Even though the entire worm body could travel a large distance (in camera coordinates) between two consecutive recording frames which were taken 0.5 s apart, the head and tail locations relative to the body centroid (worm body coordinates) tend to change little, much as a rigid body would behave (Fig. 7). The other two clues are: the worm's tail area is darker than the head (having to do with fat distribution), and the head moves more frequently than the tail (having to do with foraging behaviors). The detailed procedure, illustrated in Fig. 8, is as follows.

- 1) From recorded grayscale images, the above segmentation procedure was applied. For each video frame, the




Fig. 6. (A)–(D) Skeleton generating process. (A) Gray level image acquired from a video sequence containing the worm body and part of track. (B) Corresponding binary image after thresholding. (C) Skeleton after applying skeletonizing algorithm and redundant pixel removal. (D) Clean skeleton after pruning. (E)–(F) show typical skeletons from two different worm types. (E) Typical *unc-2* skeleton. (F) Typical *egl-19* skeleton.

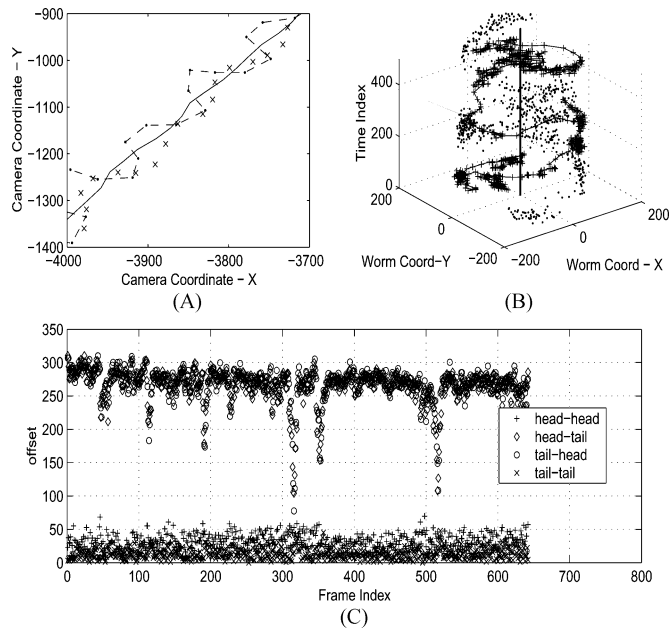


Fig. 7. Worm movement characteristics and their usage for tracking. (A) A portion of track in camera view. Solid line represents the worm body centroid movement. (x) and (·) represent worm’s tail and head location respectively, as they wiggle around the travel direction. (B) Three-dimensional plot of head and tail movement in worm coordinates. The centroid movement is represented as the vertical line in the (0, 0, t) location. The tail locations (+) are connected, showing the circular movement around the centroid. The head locations, marked by dots, tend to locate opposite the corresponding tail locations. (C) Location offset of heads and tails in worm coordinates for two consecutive frames. The head-head and tail-tail correspondences have smallest offsets of the four.

grayscale image and its corresponding binary image and skeleton were stored.

- 2) The two end points of the skeleton are potential head and tail locations. We assign the end points to two groups for each uninterrupted video segment according to the following rules: Let $endpt1x(t)$ denote the x coordinate of the end point in frame t that was assigned to group 1. Similar definitions hold for




Fig. 8. Image processing and head/tail extraction procedure. (A) Original grayscale image. Notice there is an egg object nearby that needs to be removed by cleaning. (B) Binary image after segmentation, cleaning. The worm skeleton generated from thinning, pruning process is superimposed on the binary image. Two end points of the skeleton are candidates of head and tail locations. (C) Two perpendicular lines (to skeleton line fitting) at 1/6 of skeleton location. Deleting the pixels along these separation lines divides the worm body into head, tail and middle sections. (D) Head and tail sections of the grayscale image can be easily obtained by indexing cutoff portions from the binary image from (C).

$endpt1y(t)$, $endpt2x(t)$ and $endpt2y(t)$. Now we use $endpt1(t) = [endpt1x(t), endpt1y(t)]$ to denote the vector of spatial coordinates for end point 1 in frame t . Let $endptA(t + 1)$ and $endptB(t + 1)$ denote the vectors of spatial coordinates for the two end points in frame $t + 1$ that have not yet been assigned to group 1 or group 2. Let $(M, N) = \arg \min_{(m,n)} dist(endptm(t + 1), endptn(t))$, ($m \in \{A, B\}, n \in \{1, 2\}$). Then $endptM(t + 1)$ will be assigned to group N provided that $(M, N) \neq \arg \max_{(m,n)} dist(endptm(t + 1), endptn(t))$, ($m \in \{A, B\}, n \in \{1, 2\}$). The condition statement is to avoid head and tail locations being accidentally flipped. If the condition is not met, the current frame is marked as “undecided” and the grouping process restarts from the next frame to avoid potentially spreading errors.

- 3) To isolate the head and tail sections from the rest of the body, we identify two points on the skeleton that are at 1/6 skeleton-length away from each end point. We compute the best fit line to 9-pixel-long segments from the skeleton list surrounding the two identified pixels. The lines are then rotated by 90° to get perpendicular lines. Lines that are $+5^\circ$ and -5° off from the perpendicular are also generated. The line with the shortest distance traversing the binary image is chosen as the separation line between the head/tail and the rest of the body. The end sections are

separated from the rest by deleting binary pixels along the separation lines.

- 4) Using the binary image and the end point locations as an index to the grayscale image, we calculate the median brightness of the two end sections for each frame. The means of these values for group 1 and 2 are calculated for the segment. If the difference between these two mean values is at least 20% of the larger mean value, the group with the higher average brightness value is labeled as the head.
- 5) Mutant types with digestive abnormalities have smaller brightness differences between head and tail. For these (brightness difference $\leq 20\%$), a secondary decision rule is introduced to compare the local movement distance for the two end points. The group with higher total movement distance is labeled as the head. This procedure was applied independently for each video segment. Segments are separated by missing frames, failed segmentation, or undecided frames.

D. Feature Extraction

All of the software for binarization, skeletonization, and feature extraction was coded either in C or MATLAB and implemented on UNIX machines. Some features (e.g., the area of the worm, that is, the number of pixels which make up the single binary object in the frame) could be computed on a single frame; these were computed for all 600 frames in the sequence. The average value, the maximum value and the minimum value were then computed for these 600 measurements. Other features could not be extracted from a single frame, for example, the movement between two frames, or the movement within 10 s (20 frames). Since there are 600 frames total in a sequence, the movement between two frames could be computed 300 times if we take pairs of frames in a nonoverlapping fashion, or it could be calculated 599 times taking pairs of frames in a sliding window or overlapping fashion. Likewise, for the movement within 20 frames, we could compute 581 values for overlapping 20-frame intervals. Quantities of this type were calculated in a sliding window fashion. The average, max, and min were computed from this set of numbers.

Some of the maximum and minimum values are outliers introduced by noise or errors during image capture and processing. To avoid using these extreme values, it was more useful to summarize the group statistics with such quantities as the 90th and 10th percentile values out of the population of 600 numbers. For the remainder of this paper, the terms max and min are used to denote the 90th and 10th percentile values. The measured features included the minimum, maximum, and average values of the following: distance moved in 0.5, 5, 10, 15, 20, 25, and 30 s; and 5 min, number of reversals in 10, 20, 40, 60, 80, 100, and 120 s; and 5 min, worm area, worm length, width at center and head/tail, ratio of thickness to length, fatness, eccentricity and lengths of major/minor axes of best-fit ellipse, height and width of minimum enclosing rectangle (MER), ratio of MER width and height, ratio of worm area to MER area, angle change rate, head/tail/center brightness, local head/tail/center movement relative to centroid, and head-centroid-tail angle. The area, angle change rate, and movement features were calculated for the

head, tail, center, and entire worm body respectively. We now describe in detail how several of these features were extracted from the image data.

- 1) *Body length and area*: The worm length can be readily calculated by counting the number of pixels on the skeleton. The area that the head, center and tail occupy were also measured by counting the binary pixels in these sections.
- 2) *Body thickness/width/fatness*: The worm thickness (width/length) was measured at the center, head, and tail positions of the worm skeleton (the center position was the value at the center of the skeleton pixel list; the head and tail positions were defined as the position which is 7 pixels away from head and tail end points identified by the tracking algorithm). The thickness measurement methods were as described in [1]. The center width was defined as the average width of the center section. The same widths were measured for the head and tail sections. We also defined the worm's fatness as the ratio of worm area to length.
- 3) *Local movement*: Many features characterize the global movement using the absolute distance traveled by the worm body centroid over various fixed time intervals. We also measured the relative offset of the head with regard to the centroid across the frames as an indication of the worm's head movement. This offset was defined as the movement of the head when the worm centroids were aligned on top of each other from one frame to the next. The tail movement was also measured. These measurements calculated how much the individual body parts move relative to the rest of the body.
- 4) *Angle change rate*: The angle change, an important feature for distinguishing different worm types, is defined as $R = 1/(n-1) \sum_{i=1}^{n-1} \theta_i/L$, where $\theta_i = \tan^{-1}(y_{i+2} - y_{i+1}/x_{i+2} - x_{i+1}) - \tan^{-1}(y_{i+1} - y_i/x_{i+1} - x_i)$, L is the worm length, and $(x_i, y_i), (x_{i+1}, y_{i+1}), \dots$ are the locations of consecutive points that are 5 pixels apart along the worm skeleton, and n is the number of such points along the skeleton. A larger angle change rate means that a worm has sharper body bends. Fig. 6(E), (F) shows typical skeletons from two different mutant types. The angle change rate is 15.51 for the skeleton in Fig. 6(E) compared with 8.45 in Fig. 6(F). The angle change rates were also calculated separately for head, tail, and center regions.
- 5) *Brightness*: Variations in fat distribution and absorption of nutrients cause some mutant types to become more transparent than others. The transparency can be measured by the median pixel value of the head, center, tail, and whole body regions (Fig. 8).
- 6) *Symmetry*: To measure the unbalanced muscle behavior of uncoordinated mutants, we characterized the way a worm body deviates from a perfect sine wave. These features include the amplitude (defined as the absolute distance from points on the skeleton to the line connecting the head and tail), the sum of signed distances to the line connecting the head and tail, the angle between the line connecting head to centroid, and the line connecting tail to centroid, and

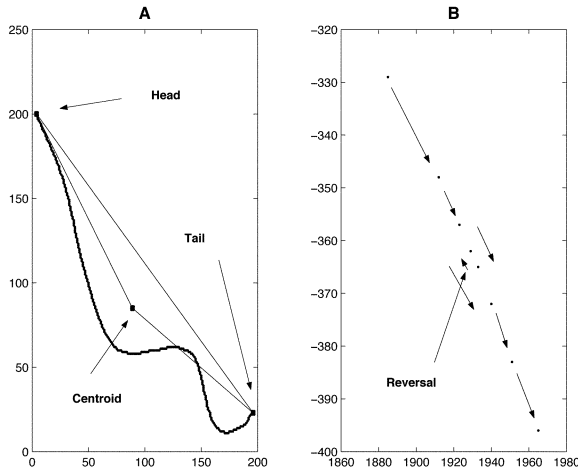


Fig. 9. Feature examples. (A) Worm skeleton, head, tail, and centroid locations. The length, angle of head to tail, head to centroid, tail to centroid lines provide symmetry information. Worm amplitude can also be measured. (B) A portion of track left by centroid. The reversal location is marked.

the distances between head and centroid, and between tail and centroid [Fig. 9(A)].

- 7) *Reversals*: Reversals were interesting characteristics during movement. They are characterized by the distance and frequency of the worm moving back into the recent previous path. We kept a moving window to record the previous 20 centroid locations. A reversal was detected when the new centroid was closer to any of the 19 previous centroid locations than to the most recent past [Fig. 9(B)].

E. Random Forests Classification and Feature Selection

Based on qualitative descriptions of uncoordinated mutant phenotypes, we expected the features measured by our system would provide useful quantitative definitions for specific mutant types [10], [16]. To assess the ability of these image features to provide effective characterization of *C. elegans* mutant phenotypes, we used the Random Forests algorithm [3], [4], [20], a classifier that uses bagging of classification trees to improve performance in distinguishing different mutant types. Random Forests utilizes an ensemble learning scheme. Instead of generating a single classification tree, many trees (to make up the forests) are generated independently by bootstrapping from the original data. A simple majority vote is taken for prediction. In addition to constructing each tree with a different bootstrap sample of the data, Random Forests adds an additional layer of randomness by splitting at each node using a random subset of predictors instead of using the best split among all features as is done in CART [5]. These two layers of randomness turn out to perform very well compared to many other classifiers including discriminant analysis, support vector machines and neural networks, and the method is robust against overfitting [3]. An estimate of the error rate can be obtained by predicting using “Out-of-bag” (OOB) data, which are the data (around 36% of the data) that were not used in each bootstrap sample. The classification error rate is, thus, defined as the aggregated OOB prediction error rate. Given enough trees being grown, the OOB error rate is quite accurate [7]. There are two free parameters

TABLE II
HEAD AND TAIL IDENTIFICATION RESULTS. DATA WERE COLLECTED FROM 161 5-MIN VIDEO SEQUENCES (2 Hz) FROM 16 DISTINCT MUTANT TYPES

Worm type	Total Frames	Recognition Wrong	Grouping Wrong	Error Percent
<i>goa-1</i>	6193	0	1	0
<i>unc-29</i>	4679	0	0	0
w.t.	6057	5	0	0
<i>egl-19</i>	5503	1	0	0
<i>cat-2</i>	4908	0	5	0
<i>dop-1</i>	4954	0	0	0
<i>dgk-1</i>	4892	0	0	0
<i>eat-4</i>	5014	0	0	0
<i>ffb-1</i>	4942	0	30	0
<i>nic-1</i>	11817	1	0	0
<i>unc-38</i>	4853	3	0	0
<i>unc-63</i>	4926	0	0	0
<i>unc-43</i>	9908	251	0	0.03
<i>tph-1</i>	10552	51	0	0
<i>unc-2</i>	10361	69	0	0.01
<i>unc-36</i>	10559	2613	3	0.24
Total	111233	3030	33	0.02

First column shows the mutant type. Second column shows the total number of frames in the videos. The number of frames that had head recognized as tail due to the tail section being lighter is listed in column 3. The number of frames that had head recognized as tail due to grouping errors is listed in column 4. The average error rate is around 2% for 111 233 frames tested.

(the number of trees in the forest and the number of random features considered at each split). Random Forests also provides four measures of feature importance that can be used for model reduction. One of these measures of feature importance, defined as the average lowering of the margin across all samples when this feature is randomly permuted, was used because it was more robust against noise (Liaw, Personal Communications). For each sample, the margin is defined as the proportion of votes for its true class minus the maximum of the proportion of votes for each of the other classes.

III. RESULTS

A. Tracking and Head and Tail Recognition

The tracking and head and tail recognition procedure was tested on 161 5-min video sequences (sampled at 2 Hz) from 16 mutant types including more than 111 000 image frames. The videos were played back with the worm’s tail marked by the algorithm for a human observer to verify. Experimental results are shown in Table II. The method produces excellent results as the average correct identification rate is around 98%.

B. Classification

For the classification, the forest was made up by 5000 trees ($n_{\text{trees}} = 5000$). At each split, 15 features were randomly selected to be considered for splitting ($m_{\text{try}} = 15$), which is approximately the square root of the total 253 features used. The confusion matrix, represented by OOB errors, is shown in Table III. The classification success rates are listed along the shaded main diagonal while the off-diagonal entries represent the misclassification error rates. The average success rate is 90.9%, showing a high degree of success at identifying the correct mutant type even if presented with a single example recording.

TABLE III
CLASSIFICATION RESULT FROM RANDOM FORESTS WITH 253 FEATURES, 5,000 TREES AND 15 RANDOM FEATURES TO SPLIT ON AT EACH NODE. THE OOB ESTIMATE OF ERROR RATE IS 90.09%

worm type	w.t.	goa-1	nic-1	unc-36	unc-38	unc-29	egl-19	unc-2	tph-1	unc-63	dgk-1	unc-43	dop-1	flp-1	eat-4	cat-2
w.t.	93								1		2		1			3
goa-1	1	93								2	1	1			2	
nic-1			98	1						1						
unc-36				97		1		2								
unc-38			3	2	83	11				1						
unc-29	1				21	74		1		1						
egl-19				2			97									
unc-2			1	1				97								
tph-1			1					3	89	1			2	1	3	
unc-63			6	1	1	3				87			1		1	
dgk-1	1	1								3	91					3
unc-43							1					98		1		
dop-1				1			1		2				88	1	7	
flp-1							3		1	1	5	1	1	85	2	1
eat-4			4	1					4	1					90	
cat-2	4									1	1		1	2		91

TABLE IV
IMPORTANT FEATURES IDENTIFIED BY RANDOM FORESTS. THE OOB ESTIMATE OF ERROR RATE BY USING THESE 25 FEATURES IS 85.9%

Feature	Description
MVHLFMAX	Maximal centroid movement in 0.5 sec
HTMVRAVG	Local head and tail movement ratio
PRP10MAX	Maximal centroid movement in 5 seconds
TOTMOVE	Total centroid movement in 5 minutes
TAILBRMAX	Maximal tail brightness
TAILBRAVG	Average tail brightness
PRP20MAX	Maximal centroid movement in 10 seconds
HTBRRMIN	Head and tail brightness ratio
HDHFTOTMV	Sum of head movements in 0.5 sec
LNECRMIN	Length/eccentricity max
ANCHRMAX	Max angle change rate
TLHFTOTMV	Sum of tail movements in 0.5 sec
LNGTHAVG	Average Length
TLMVHFMAX	Maximal tail movements in 0.5 sec
MVHLFAVG	Average centroid movement in 0.5 sec
TAILBRMIN	Minimal tail brightness
HTBRRAVG	Average head to tail brightness ratio
LNGTHMIN	Minimal length
HDMVHFAVG	Average head movements in 0.5 sec
HDMVHFMAX	Maximal head movements in 0.5 sec
HTBRRMAX	Maximal head to tail brightness ratio
LNGTHMAX	Maximal length
AVGBRMAX	Maximal body brightness
ANCHSMAX	Maximal standard deviation of angle change rate
TLTHKMIN	Minimal tail thickness

The important features identified by Random Forests are shown in Table IV. If we run the classification procedure using only the top 25 features (10% of the total features) identified by Random Forests, we achieved 86% classification accuracy.

IV. DISCUSSION

A. Effect of Number of Trees and Number of Features Used at Each Split

Fig. 10(A) shows the effect of the number of features selected at each split on the error rate with 5000 trees constructed.




Fig. 10. (A) Effect of number of features selected at each split on OOB error > rate. Five-thousand trees are used. The horizontal line represents the > error rate using 15 features. (B) Effect of number of trees used to construct Random Forests on OOB error rate. Fifteen features are used at each split.

The errors are stable between 10 and 100 features ($error = [0.088, 0.095]$) and trend upward slightly afterwards. Fig. 10(B) shows the effect of the number of trees used when 15 features are selected at each split. The error converges quickly after 800 trees ($error = [0.090, 0.096]$) are constructed. Both plots indicate that the results are not sensitive to the selection of these two parameters.

B. Comparison to Related Work

A number of automated systems have been developed for quantifying specific behavioral parameters in *C. elegans* as well as other genetically-tractable organisms. For example, there are systems in [25] that reveal the alternative behavioral states controlled by serotonin in egg-laying behavior; systems [9] that

study solitary (move slow) and social feeding (move fast) behaviors, systems [22] that investigate the behavioral mechanism of chemotaxis by studying the speed and turning rate during chemotaxis in gradients of the attractants ammonium chloride or biotin; systems [21] that report progressive increases in both locomotor activity and stereotyped behavior known as “reverse tolerance” or “behavioral sensitization” caused by repeated intermittent doses of cocaine; systems [2] that show that acute responses to cocaine and nicotine are blunted by pharmacologically induced reductions in dopamine levels by measuring the effect of psychostimulants on fly behavior. Each system is capable of measuring some specific behavioral parameters, but there is no automated system that is designed to classify a large number of the mutant types by evaluating large numbers of behavioral parameters simultaneously.

C. Comparison to Related Work

However, it is interesting to compare our system with human experts. A preliminary comparison with human observers reveals that our automatic system outperforms the human dramatically. For example, in one experiment we conducted, an experienced observer was presented with 100 1-min videos of *unc-36* and *unc-2* (50 of each), and identified only 50% of *unc-36* and 90% of *unc-2* frames correctly. Running the same experiment on *dgk-1* and *goa-1*, the observer identified 84% of *goa-1* and 52% of *dgk-1* correctly. Using 1-min videos for the human observer reduces the experiment time by 80% compared to 5-min videos, and in any case, the human observer typically makes his decision within the first 70 frames (35 s). The human observer in our experiment is a *C. elegans* expert with more than 20 years of experience working in this field. The system produces over 93% correctness for each of these types against 15 other types combined. Furthermore, from the features extracted by the system, the head and tail brightness difference and total movement and reversals are the top features distinguishing *unc-2*/*unc-36* and *goa-1*/*dgk-1* pairs, respectively. These features are hard to quantify by eye.

D. Applications for Computer Vision-Based Quantification of Mutant Phenotypes

Hundreds of genes have been identified in *C. elegans* that affect behavior and morphology in specific ways. Our long-term aim is to collect data on large numbers of mutant types and effectively classify them according to their phenotypic similarity. With an increasing data set, it becomes progressively more challenging to identify features that effectively classify and distinguish the large variety of worm types. The image processing methods developed here, including new features that require accurately identifying the head and tail regions of the animal, allowed us to achieve high classification accuracy even for a data set involving 16 different mutant types with subtle distinctions that are hard to classify by eye. These methods also hold promise for investigating the clustering of behavioral and morphological patterns seen in different mutant worm types. In a recent study [12], we investigated the natural clustering of *C. elegans* behavioral and morphological phenotypes using data collected by our automated tracking system. From a complex

data set consisting of 253 features measured from recordings of 797 individuals representing 8 distinct genotypes, we used principal component analysis to represent each mutant type as a cloud of data points in low-dimensional feature space. We also used k-means clustering and Euclidean distance measurements to explore the natural structure of the behavioral data and to compare the similarities of mutant phenotypic patterns. Encouragingly, the phenotypic classes inferred from the cluster analysis matched the known molecular similarities of the mutants that were grouped together; for example, two nicotinic receptor mutants formed a single cluster, as did two calcium channel mutants. Together, these results provided a precise and comprehensive definition of several important *C. elegans* phenotypes, and demonstrate that mutant phenotypes can be clustered using a complex behavioral and morphological signature based on quantitative image features. The image processing and computer vision methods developed in this report are basic building blocks for studying specific *C. elegans* behaviors. For example, we have developed methods [13] to automatically detect egg-laying events which allow the study of egg-laying events on a large scale. Future studies with multi-animal behaviors such as mating, social feeding, etc, can also take advantage of the methods developed. The algorithms developed here for tracking, head/tail recognition and feature extraction will be an essential part of a completely automated *C. elegans* tracking and identification system.

ACKNOWLEDGMENT

The authors would like to thank the Caenorhabditis Genetics Center for strains, and M. Palm, M. Orlov, and D. Poole for data collection. They would also like to thank an anonymous reviewer for the suggestions on potential usefulness of and other applications for the automated system.

REFERENCES

- [1] J. Baek, P. Cosman, Z. Feng, J. Silver, and W. R. Schafer, “Using machine vision to analyze and classify *C. elegans* behavioral phenotypes quantitatively,” *J. Neurosci. Meth.*, vol. 118, p. 921, 2002.
- [2] R. Bainton, L. Tsai, C. Singh, M. Moore, W. Neckameyer, and U. Heberlein, “Dopamine modulates acute responses to cocaine, nicotine and ethanol in *Drosophila*,” *Curr. Biol.*, vol. 10, pp. 187–194, 2000.
- [3] L. Breiman, “Random forests,” *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [4] —, (2002) Manual on Setting Up, Using and Understanding Random Forests v3.1. [Online]. Available: <http://citeseer.nj.nec.com/breiman01random.html>
- [5] L. Breiman, J. Friedman, R. Olshen, and C. Stone, *Classification and Regression Trees*. Belmont, CA: Wadsworth, 1984.
- [6] S. Brenner, “The genetics of *Caenorhabditis elegans*,” *Genetics*, vol. 77, pp. 77–94, 1974.
- [7] T. Bylander, “Estimating generalization error on two-class datasets using out-of-bag estimates,” *Machine Learning*, vol. 48, pp. 287–297, 2002.
- [8] R. Dhawan, D. B. Dusenbery, and P. L. Williams, “Comparison of lethality, reproduction, and behavior as toxicological endpoints in the nematode *Caenorhabditis elegans*,” *J. Toxicol. Environ. Health*, pt. A, vol. 58, pp. 451–462, 1999.
- [9] M. de Bono and C. I. Bargmann, “Natural variation in a neuropeptide y receptor homolog modifies social behavior and food response in *C. elegans*,” *Cell*, vol. 94, pp. 679–689, 1998.
- [10] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, 2nd ed. New York: Wiley, 2002.
- [11] A. Fraser *et al.*, “2000 functional genomic analysis of *C. elegans* chromosome I by systematic RNA interference,” *Nature*, vol. 408, pp. 325–330, 2001.

- [12] W. Geng, P. Cosman, J.-H. Baek, C. Berry, and W. R. Schafer, "Quantitative classification and natural clustering of *C. elegans* behavioral phenotypes," *Genetics*, vol. 165, pp. 1117–1136, 2003.
- [13] W. Geng, P. Cosman, M. Palm, and W. R. Schafer, "C. elegans egg-laying detection and behavior study using image analysis," *EUROSIP J. Appl. Signal Processing*, submitted for publication.
- [14] R. Gonzalez and R. Woods, *Digital Image Processing*, 2nd ed. Englewood Cliffs, NJ: Prentice Hall, 2002.
- [15] L. A. Hardaker, L. A. Singer, E. Kerr, G. T. Zhou, and W. R. Schafer, "Serotonin modulates locomotory behavior and coordinates egg laying and movement in *Caenorhabditis elegans*," *J. Neurobiol.*, vol. 49, pp. 303–313, 2001.
- [16] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*. New York: Springer, 2002.
- [17] J. Hodgkin, "Male phenotypes and mating efficiency in *Caenorhabditis elegans*," *J. Genetics*, vol. 103, pp. 43–64, 1983.
- [18] R. Jain, K. Rangachar, and B. Schunck, *Machine Vision*. New York: McGraw-Hill, 1995.
- [19] S. Kim, D. Poole, L. Waggoner, A. Kempf, D. Ramirez, A. Treschow, and W. R. Schafer, "Genes affecting the activity of nicotinic receptors involved in *C. elegans* egg-laying behavior," *Genetics*, vol. 157, pp. 1599–1610, 2001.
- [20] A. Liaw and A. Wiener. (2002) Classification and regression by random forest. *R Newsletter* [Online]. Available: <http://www.r-project.org>
- [21] C. McClung and J. Hirsh, "Stereotypic behavioral responses to free-base cocaine and the development of behavioral sensitization in *Drosophila*," *Curr. Biol.*, vol. 8, pp. 109–112, 1998.
- [22] J. Pierce-Shimomura, T. Morse, and S. Lockery, "The fundamental role of pirouettes in *C. elegans* chemotaxis," *J. Neurosci.*, vol. 19, no. 21, pp. 9557–9569, 1999.
- [23] J. Sulston and H. Horvitz, "Post-embryonic cell lineages of the nematode *Caenorhabditis elegans*," *Dev. Biol.*, vol. 56, pp. 110–156, 1977.
- [24] J. Sulston, E. Schierenberg, J. White, and J. Thomson, "The embryonic cell lineage of the nematode *Caenorhabditis elegans*," *Dev. Biol.*, vol. 100, pp. 64–119, 1983.
- [25] L. E. Waggoner, G. T. Zhou, R. W. Schafer, and W. R. Schafer, "Control of alternative behavioral states by serotonin in *Caenorhabditis elegans*," *Neuron*, vol. 21, pp. 203–214, 1998.
- [26] J. E. White, E. Southgate, N. Thomson, and S. Brenner, "The structure of the *Caenorhabditis elegans* nervous system," *Philos. Trans. Roy. Soc. Lond. (Biol.)*, vol. 314, pp. 1–340, 1986.
- [27] T. Y. Zhang and C. Y. Suen, "A fast parallel algorithm for thinning digital patterns," *Comm. ACM*, vol. 27, no. 3, pp. 236–239, 1984.
- [28] G. T. Zhou, W. R. Schafer, and R. W. Schafer, "A three-state biological point process model and its parameter estimation," *IEEE Trans. Signal Processing*, vol. 46, pp. 2698–2707, Oct 1998.
- [29] P. Zipperlen, A. Fraser, R. Kamath, M. Martinez-Campos, and J. Ahringer, "Roles for 147 embryonic lethal genes on *C. elegans* chromosome I identified by RNA interference and video microscopy," *EMBO J.*, vol. 20, pp. 3984–3992, 2001.



Wei Geng received the B.S. degree in computer science from Beihang University, Beijing, China, in 1993, and the M.S. degree from Northeastern University, Boston, MA, in 1997, and the Ph.D. degree from the University of California, San Diego, CA, in 2004, both in electrical engineering.

From 1993–1995, he was a Research Assistant at the National Laboratory on Machine Perception of Peking University, Beijing, China. From 1997–2001, he was an Embedded Software Engineer of Hewlett-Packard Company, San Diego, CA. He is currently a

Senior Analytical Engineer of I.D. Analytics Inc., San Diego, CA. His research interests are in the areas of data mining and image analysis.



Pamela Cosman (S'88–M'93–SM'00) received the B.S. degree with Honor in electrical engineering from the California Institute of Technology, Pasadena, in 1987, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1989 and 1993, respectively.

She was an NSF postdoctoral fellow at Stanford University, Stanford, CA, and a Visiting Professor at the University of Minnesota, Twin Cities, during 1993–1995. In 1995, she joined the faculty of the department of Electrical and Computer Engineering at the University of California, San Diego, where she is currently a Professor. Her research interests are in the areas of image and video compression and processing.

Dr. Cosman is the recipient of the ECE Departmental Graduate Teaching Award (1996), a Career Award from the National Science Foundation (1996–1999), and a Powell Faculty Fellowship (1997–1998). She was an associate editor of the *IEEE COMMUNICATIONS LETTERS* (1998–2001), a guest editor of the June 2000 special issue of the *IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS* on "Error-resilient image and video coding," and was the Technical Program Chair of the 1998 Information Theory Workshop in San Diego. She is currently an associate editor of the *IEEE SIGNAL PROCESSING LETTERS*, and a senior editor of the *IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS*. She is a member of Tau Beta Pi and Sigma Xi. Her web page address is <http://www.code.ucsd.edu/cosman/>.



Charles C. Berry was born in Ithaca, New York, on July 30, 1946. He received the B.A. degree in sociology from Dartmouth College, Hanover, NH, in 1969 and the Ph.D. degree in social relations from The Johns Hopkins University, Baltimore, MD, in 1973. Since 1973, he has worked at the University of California, San Diego. His area of interest is biostatistics.



Zhaoyang Feng received the B.S. degree in neurobiology (excellent thesis) from the University of Science and Technology, Huafei, China, in 1992, the M.S. degree in computer engineering from Syracuse University, Syracuse, NY, in 2001, and the Ph.D. degree in biochemistry and molecular biology from the State University of New York (SUNY) Upstate Medical University, Syracuse in 2001.

From 1992–1995, he was a Research Assistant at the Institute of Biophysics of the Chinese Academy of Sciences, Beijing, and from 1995–1996 at the Institute for Molecular Biology, University of Oregon, Eugene. Since 2001, he has been a Postdoctoral Fellow in the Division of Biological Sciences at the University of California, San Diego. His research concerns the quantitative analysis of behavior in nematodes.



William R. Schafer received the A.B. degree in biology (*summa cum laude*) from Harvard University, Cambridge, MA, in 1986, and the Ph.D. degree in biochemistry from the University of California, Berkeley, in 1991.

From 1992–1995, he was a Life Sciences Research Foundation Postdoctoral Fellow with the Department of Biochemistry and Biophysics, University of California, San Francisco. Since 1995, he has been on the faculty of the Division of Biological Sciences at the University of California, San Diego, where he is currently

Associate Professor and Vice-Chair of Neurobiology. His research concerns the molecular and cellular basis for nervous system function and the relationship between genes and behavior.

Dr. Schafer is a member of the Genetics Society of America, the Society for Neuroscience, and the International Society for Neuroethology. He received the John Belling Prize for the outstanding Ph.D. thesis in Genetics at UC Berkeley between 1988 and 1993. He also has been the recipient of a Young Investigator Awards from the National Alliance for Schizophrenia and Depression and the Beckman Foundation, an Alfred P. Sloan Fellowship, and a Klingenstein Fellowship in Neuroscience. In 2001, he received a Presidential Early Career Award from the National Institutes of Health.