

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A. I. Memo 374

July 1976

A Computerized Look at Cat Locomotion

or

One Way to Scan a Cat

by

Glen Speckert

This paper describes a three phase project concerning the watching, analyzing, and describing of motions of a cat in various gaits. All data is based on two 16 mm films of an actual cat moving on a treadmill. In phase I, the low level issues of tracking key points on the cat from frame to frame are discussed. Phase II deals with building and using a graphics tool to analyze the data of phase I. Phase III is a high level discussion of cat locomotion based on the trajectories and movements explored by phase II.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-75C-0643-0005.

This paper describes a three phase project concerning the viewing, analyzing, and description of the motions of a cat running on a treadmill. The first phase deals with the dynamic vision issues of finding and tracking the key points of the cat from a 16mm film. This was successfully done with a microscope attachment to a slow scan vidicon which viewed the film directly. The second phase was more a graphics problem of how best to display the points of phase one in order to be most useful and easily understood. Graphs of angular positions, angular velocities, and angular accelerations of all joints are available as well as footfall diagrams and both static and dynamic trajectories. The third phase is more in the realm of kinematics and/or biophysics as it deals with the hows and whys of locomotion. For this reason, the analysis is not as deep as the data will allow. A full and careful analysis can and should be undertaken, but that is not the purpose of this paper, and hence the analysis will not be as fully developed as is possible.

Two short 16mm films of a cat running on a treadmill were shown in the MIT AI playroom in January, and afterwards, the question was raised, "I wonder how long it will be before a machine can run like that?" To which the response was, "How long will it be until it understands that?" A final remark was then heard, "How long until a machine can watch that?"

This paper answers the last of these questions, namely that now a machine has "seen" the films, and has maybe even begun to "understand" cat locomotion, by graphically analyzing the data from the film. Whether or

not-cat locomotion will prove a suitable model for machine locomotion is an open question.

This project is really the combination of three distinct projects. The first is extracting the positions of key points on the cat. While this may seem an unexciting low level task, it is actually an interesting vision project, one which employs both tracking and recognition skills. It is by no means a simple task, because the cat's body at times obscured the key points that were being tracked, the film occasionally became slightly blurred or out of focus, some points were moving quite rapidly, plus all the usual problems associated with taking data from real world pictures. In spite of these problems, the first phase of this project is able to reduce the information in the hundreds of frames of thousands of pixels of 256 grey levels each, to hundreds of sets of key point positions, one set per frame of film viewed. The positions of these key points contain most if not all of the same information that the film contained, but is much easier to use.

The second phase takes these sets of points (eight points per set) and plots them into motions and trajectories, as well as displaying an animated cat that moves across the screen. The differences between galloping, trotting, and walking can easily be seen in the graphs and trajectories. Simple curves on the graph contains all the information of the many sets of points, and thus a further reduction of volume of information and increase in power of the information is obtained.

The third phase assimilates all the motions and trajectories viewed in phase II into a compact theory of motion. The differences in the

various gaits are understood and explained. The graphs and trajectories are analyzed and the reasons behind them are understood. At this stage can "The Computer" be said to have watched and understood what the cat was doing?

Note: The term "run" or "running" will be loosely used in this paper to mean moving in a walk, trot, or gallop. Eadward Muybridge states that "in its reference to quadrupedal movements, 'running' can be applied only as it is to a stream of water running down a hill, a locomotive running along a railroad, or an ivy plant running up a wall."

PHASE I -- EXTRACTING POINTS FROM ACTUAL FILM

In this section, the basic gray-level methods of finding the key points are described. The input to this section is the 16mm film, and the output is a disk file of key-point positions.

The method of viewing the film should first be described. An old 16mm movie projector was torn apart to obtain the hook and claw mechanism that pulls the film exactly one frame forward. This mechanism is placed on a light table, and a slow scan vidicon with a microscope attachment is used to view it from above. The resolution of the microscope lens is changable, but only two lenses proved useful, since the higher the resolution, the smaller the field of view, and the cat would soon drift out of range. The movie camera must be manually advanced after each frame, in order to view the next frame. See figure 1. The vidicon interfaces to a PDP 11/45 which has a DEC GT40 (graphics display terminal) as its primary output device. All software is written in PDP 11 Assembly language.

The films used were prepared at the University of Arizona by Dr. Mary C. Wetzel, to whom we are deeply thankful for a copy of the films. Gummed reinforcements (small circular items usually found on three ring notebook paper) were attached to the skin of the cat directly over the front right shoulder, elbow, and wrist; the back right hip, knee, and ankle; and two more near the tail. Since the film can be viewed from either side, the cat is always viewed so as to be traveling to the right of the vidicon. The side facing the viewer is termed the "right" side, and





FIGURE 1. Physical set-up.

the side away from the viewer shall be referred to as the "left" side. The points actually tracked are points on the skin which may move relative to the joints beneath them, but this relative motion is believed to be small. A treadmill was used to keep the cat relatively in one place in front of the camera. The points found represent only the right front legs, right rear legs, and rear end positions, and tell nothing of the left legs. In some of the later films, the rear end markers were not even tracked, as their positions were used little in the final analysis. The unmarked left legs were assumed to travel the same trajectories as the right ones, only with some phase lag. See phase II for further discussion of unmarked legs. The gummed reinforcements (hereafter referred to as markers) were not placed in the same position on the two different cats of the two films.

Thus the problem of finding the markers was one of finding a target

that looks like  on a basically dark but splotchy cat. This is a vastly simpler task than finding the joints without special markers, as was done by Speckert for the human figure (MIT Working Paper 118, January, 1976), and more general methods are applicable. The basic method employed was inspired by a method used by Kimae, Ballard, and Sklansky (1), and has no relationship to pattern matching. The contrast in the film was often small between the cat, the target, and the background, especially near the back of the cat, and this did cause some problems.

The gradient (magnitude and direction) of the gray level picture is computed at each cell. The magnitude of the gradient is the maximum difference between one neighbor of a cell and the "opposite" neighbor (through the center cell). The direction is one of the eight possible directions this can take (see figure 2). If the magnitude of the gradient is above a threshold, that cell is considered an edge point. In a second, initially zero array, an arc is "struck" about each edge point in the direction of the gradient at the known marker radius (the larger radius) by incrementing the cells that lie on this arc. This arc is a quarter of a circle in length. (See figure 3). Each cell in this array ends up as a count of the number of edge points that are a given radius away and whose gradient points toward that cell. The cell whose position is the same as the center of the marker should get incremented by all the marker edge points, and thus would be the maximum. A second array was also used, but with the inner radius of the target and reverse gradient directions used to strike the arcs (a gummed reinforcement has an outer and an inner radius). Both of these arrays are examined for local maxima, and a list of possible

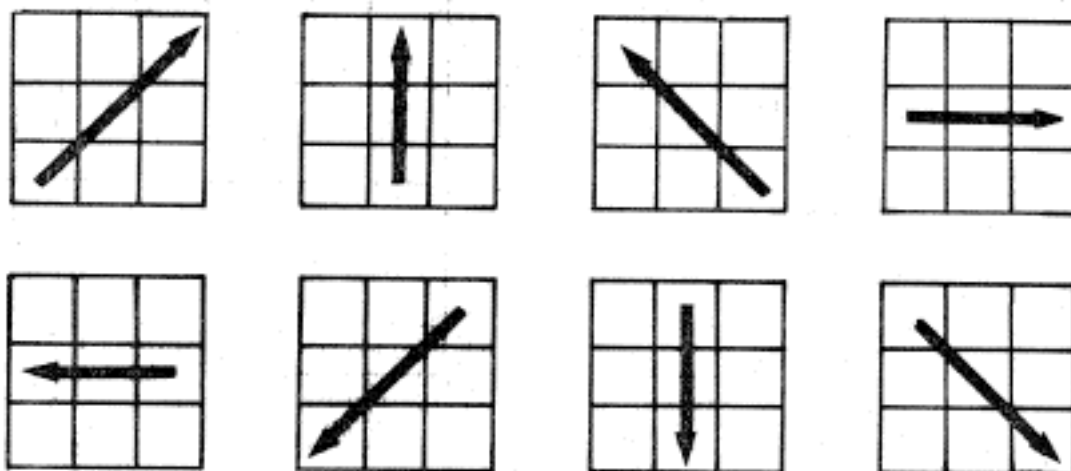


FIGURE 2. Eight possible directions of gradients.




FIGURE 3. Arcs are struck from each edge point. Maximum is near the center of hypothesized circle.

target locations is prepared. The target position is chosen from these possible locations based on other factors, such as predicted position (See below). In practice, this method worked fairly well when the picture was reasonably sharp and clear, and better at higher resolutions than lower ones. See figure 4 for an example of how this method works when everything is in order.

However, often the markers were blurred, obscured, or distorted, and this method did little good. Figure 5 shows a poor quality marker and its associated arrays.

The position of the target is first predicted based on its past history and if it is a "dependent" target, its position is also computed based on a relative position to the target on which it is dependent. A dependent target is one that is more or less rigidly attached to another target, for example, the elbow joint is dependent on the shoulder joint. The shoulder/hip joints are independent as are the two on the rear end. The elbow/knee and wrist/foot joints are dependent on the shoulder/hip and elbow/knee joints respectively. The shoulders were chosen as the independent targets rather than the feet due to their slower motion. Some joints must be independent, in order that their position can be predicted based only on their past histories, and not relative to other targets. These independent targets are found first, in order to give a relative position to their dependent targets. If the position of a target predicted from its past motions agrees with the prediction based on relative position to another target, the search area is constricted. If they vary widely,





FIGURE 4. Method works well for sharp pictures.

FIGURE 5. Method doesn't work as well for poor pictures.



the area searched is widened. A picture is taken centered about the best guess, and is handed to a routine along with a threshold for edge points, and a radius within which to search. This radius can be clearly seen in figures 4 and 5. The area outside of this radius is eliminated from consideration. In figure 5, some points were not even read from the vidicon, since even if they contained edge points, the arcs struck would not reach inside this radius of consideration. A list of possible target locations is prepared based on the picture. The "best" choice is chosen from this list based on the past history of this target and relative positions of other targets. If the picture is of good quality, there will be only one or two pictures in this list. In figure 5 above, many possible targets were returned. The information in the database which is used to pick the "best" choice includes target type, the past eight positions, the past three scores, the initial search radius, and if it is a dependent target, also the past three angles to the parent target and the distance to that target.

After all markers have been located, they are displayed while the operator advances the film. The operator may decide to correct any target location, by specifying the target number and scrolling it to the proper position. Before starting the next frame, all the positions are written out on disk to be used in phase II.

Just to get an idea of what the film is like, figure 6 shows every fifth frame of one stride of a cat trotting. The dynamic range has been greatly enhanced, and thus the picture appears to have much higher contrast than it actually does. Most of the sections of film viewed by Phase I were

two to six strides long. The time per frame required to locate the targets varies from 30 to 60 seconds, depending on the number of targets (6 vs 8) and the amount of intermediate display to the user. Almost half of the time is spent reading points from the vidicon.

For ease in discussion, the front shoulder joint shall henceforth be joint 0, the front elbow joint is joint 1, the front wrist is joint 2, the back hip is joint 3, the back elbow is joint 4, the back foot is joint 5, the hindmost rear marker is joint 6, the other rear marker is joint 8.

Now that the cat films have been reduced to disk files of points, we are ready to begin to analyze them. This is the task of phase II.



FIGURE 6. Cat trotting




FIGURE 6. Cat trotting

PHASE II Developing a graphics tool to view cat films

Once the point files have been created, one needs some method of viewing, examining, and analyzing them. For this purpose, a graphics system was developed. The inputs to this system are the point files from phase I, and the output is a variety of graphs, trajectories, and other "high-level" descriptions of the cat's movements, to be used in phase III to develop a theory of cat locomotion.

The graphic terminal is a DEC GT40 which has an 11/05 processor as well as a scope processor built into it. The 11/05 was used to communicate with the 11/45, perform minor computations, do bookkeeping, and modify the scope code.

The graphics routine developed (called CATDIS for CAT DISplay) has the capability to plot the angle, angular velocity, and angular acceleration at any joint or set of joints. The user specifies an angle by naming three joints. If only two different joints are specified (one joint specified twice), the angle that the line through these two joints makes with the horizontal is plotted. The angles which are examined in phase III include 0-1-2, 0-1-1, 1-2-2, 3-4-5, 3-4-4, and 4-5-5 (See end of phase I description to review joint numbers). The angular velocity is the difference in the last two positions, i.e. position (i)-position(i-1). However if this is too noisy, the user can request that the velocity be computed as position(i)-position(i-2), or even position(i)-position(i-3), etc. This effectively filters out some of the noise. Similarly, the acceleration can use a filter on the velocity. The file name is always

shown as OFICAL.N S:2 T:2, where the number after the S tells how many times the data has been smoothed, and the number after the T tells the time delay in points used to compute the velocity and acceleration, with T:1 corresponding to the velocity shown being position(i)-position(i-1). A value of S:2 T:2 was used in phase III to minimize noise. Figure 7 shows a plot of the rear knee joint angle for a trot.





FIGURE 7. Rear Knee Joint. Position, Angular Velocity, Angular Acceleration (trot).

An animated cat (built around the sets of points) can also be seen which will either run in place, or move across the screen as the successive sets of points arrive from the 11/45. This cat will disappear at the flick of a switch to allow closer inspection of the points themselves. The trajectories of the actual points can be made to stay in view as new sets of points arrive, thus showing a history of that joint, either as the cat

is moving across the screen, or as it is running in place (on a treadmill).

Eadward Muybridge (1887) did most of his analysis of motion by watching the patterns of footfalls, i.e. studying the order in which each leg hit and was lifted from the ground. He used a simple diagram that looked like



where \triangle is the front left foot, \blacktriangle is the front right foot, \circ is the back left foot, and \bullet is the back right foot. If these symbols are present in the diagram, that foot is on the ground, and if absent, that foot was not at that moment supporting the body. Thus the above diagram depicts an animal standing on all four legs, and



represent an animal with two feet on the ground and one that momentarily has no feet on the ground, respectively. These Muybridge type footfall diagrams can be seen in CATDIS, either dynamically as the cat moves, or as a series of static diagrams, showing the sequence and duration of footfalls. The solid symbols \bullet and \blacktriangle have been replaced by \odot and \triangleleft . A foot is considered as being on the ground if it is below a user specified Y value AND if it is moving backwards on the original film (where the cat is on a treadmill that is moving backwards). The front and rear ground levels can be set independently in the event that the markers are a different height above the ground on the cat. These diagrams are very




FIGURE 8. One complete stride of a walk.

useful in classifying gaits. Figure 8 shows one complete stride of a walk. The count below each diagram tells how many frames that "phase" lasted. Each frame is approximately 10 milliseconds, since the film was taken at 100 frames per second. Figure 9 shows a cat galloping with a dynamic diagram at the left (as the cat moves, the diagram changes).




FIGURE 9. Cat galloping.

For more detailed information on how to see what you want to see

from CATDIS, see Appendix 2, an operational guide to CATDIS.

Thus with all of these high level descriptions of the motions of the targets available, an understanding of cat locomotion is possible. This is the purpose of phase III.

Phase III Understanding Cat Locomotion

The input to this section is the graphs, diagrams, trajectories, and other high level information of phase II and the output is a theory of cat locomotion. The rotary gallop, trot and walk are compared.

These cat films are accurate enough and frequent enough that some understanding of the motion is possible, unlike earlier attempts of watching human figures in motion. (See Working Paper 118 Knowledge Driven Recognition of the Human Body by Glen Speckert) The films were taken at a rate of 100 frames per second, so each frame represents about 10 msec.

The walk is the slowest, most restful gait, and as can be seen by figure 10, the cat is supported alternately on two or three legs. As a



FIGURE 10. Walk.

general rule, whenever the animal is supported by two legs, if the moving legs are between the supporting ones, both are lateral pairs. If the moving legs are outside the supporting legs, both pairs are diagonals

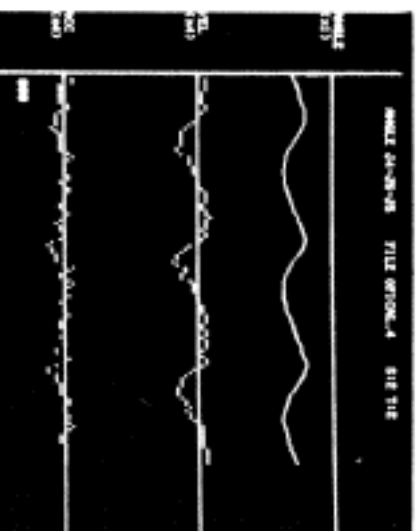







FIGURE 11. Complete set of graphs for walk.

(Muybridge). Thus the walk sequence consists of (refer to figure 10) starting on three legs and lifting the front right, dropping the back right, lifting the back left, dropping the front right, lifting front left, dropping the back left, lifting the back right, dropping the front left and recursing. This is a symmetric gait as the right legs are 180 degrees out of phase with the left ones. The front leg seems to lag the back leg on the same side by 90 degrees. Figure 11 shows the graphs of all the angles for a walk. Note the low accelerations and hence low energy expended by this gait.

The trot is a slightly faster gait, and consists basically of diagonal pairs working together (in phase) to alternately support the body. Muybridge gained great fame by showing that horses have a slight period of non-support between changing of pairs of supporting legs. This can be seen in a very short phase (about 10 msec) in figure 12. In all the trots (see



FIGURE 12. Trot.

figure 13), the cat stayed a long time (100-120 msec) on its diagonals and took about 40 msec to change to the other pair. The four events of lifting


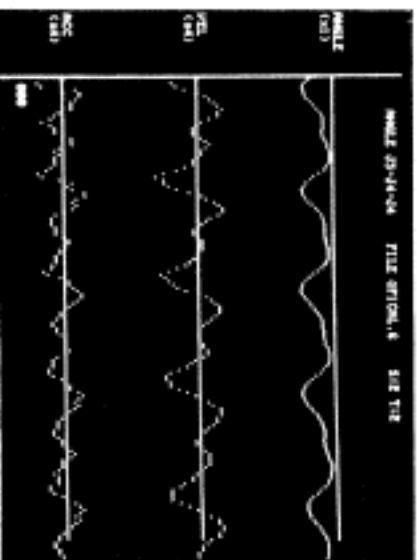






FIGURE 14. Complete set of graphs for trot.




FIGURE 13. More trots.

the two supporting feet and setting the two new supporting feet down happen close together in time, and their precise order is not always the same. However the front leg in the diagonal pair always leaves the ground sooner than its diagonal partner, and it also always lands slightly sooner. Figure 14 shows a complete set of graphs for all joints and angles of the trot. Note that the accelerations are higher than in the walk, but not as high as for the gallop. The trot is a symmetric gait with the right and left legs 180 degrees out of phase and also the front feet are 180 degrees out of phase with the rear ones on the same side.

The rotary gallop is the fastest gait of the cat. Except for one long period of unsupported transit, the cat alternates support on one or two legs. Starting in flight, the cat first touches down with its rear left leg followed by the right rear. The right rear then stands alone as the left rear shoves off and the cat stretches out. The right front touches down just before the right rear leaves and it is soon joined by the left front. The left front soon becomes the sole supporting leg of the body for









FIGURE 16. Complete set of graphs for Rotary Gallop.

a fairly long period before contributing a strong kick which sends the cat back into its long period of unsupported flight. See figure 15. When the



FIGURE 15. Rotary Gallop.

cat is in flight, all four of its legs are tucked beneath the body, and this phase lasts for slightly over one fourth of the stride. It can be easily seen why this gait is called the rotary gallop, as the footfalls proceed in a counter clockwise circle around the cat. If the roles of the right front and left front are reversed, this gait becomes the transverse gallop. The graphs of figure 16 can be seen to contain higher acceleration peaks than in either the walk or the trot, and thus the energy required is higher. This is not a symmetric gait, as the front left leg lags the front right one by 60 degrees and the back left leg leads the back right one by 60 degrees.

The velocities of these gaits can be found by finding the minimum offset needed to add to each frame such that the feet never appear to move backwards. This is done to produce the point trajectories of the cat as if it were running on stationary ground. The speeds are 25, 19, and 10 for

the gallop, trot, and walk respectively. By setting these velocities in CATDIS, the cat can be seen to move across the screen. The file length, period, phase lags, and velocities for all the data files are tabulated in Appendix 3. Again, further analysis is possible from the data, and is encouraged, but that is not the purpose of this research.

Thus it can be seen that a reasonable amount of kinematic analysis can be done on points which were visually located from an actual real world film. This shows the reliability of the dynamic vision system which the graphics system and kinematic analysis are based upon.

"The Computer" has seen the cat films, understood the basic mechanics of several types of motion, and now it is up to the mechanical engineers to enable it to imitate the cat, and someday be able to walk, trot, or gallop.

APPENDIX 1 Operational Guide to Phase I

This appendix describes how to use the first phase of this project, the phase concerned with converting 16mm film to disk files of point positions. To run the program, do a

`:CATHI`

This will ask then for a target radius (6 for 2 Plan scope probe, 9 for 3 Plan) and a threshold. The threshold is the value which determines whether a point has enough spread across it to be an "edge" point. 7 is a good value to give. At this point it will respond with

ENTER COMMAND

There are many commands available, and a `?<cr>` will list them. Probably though, you will want to do an

`EP<cr>` (enter picture mode)

first. This will load in PICTB7 into the GT40. This is a display program which will display one sector pictures. After loading in this display program, my program will ask you what you want to see displayed. You chose what you want to see, and indicate if it should wait for a character before

proceeding (thus allowing you to examine the picture at your leisure).
Items which you can view include the actual picture as it is taken from the vidicon, the array SCRBUF which is the array with the large arcs drawn on it, the array SCR2BF, the array that looks for the smaller center hole in the target, and a count of the number of "edge points".

At this point it will again ask you for a command. Probably at this point, you will want to build a database, or read one off disk. RD will read one in from disk, but you probably will want to do a

BD<cr> (build database)

This will ask you if the target is independent, or dependent on another target, if so, which one. Targets 0,3,6,7 are independent, and 1,2,4,5 are dependent on 0,1,3,4 respectively. It will also ask you for a "center distance max". This is the distance from the projected position that is searched. For closely spaced targets 6 and 7, a small distance like 6 should be used to avoid confusion between them. For the other targets, this is not a very important parameter, and anywhere from 10-30 is OK. Now you must scroll to within a target radius of the target. Whenever it wants you to scroll, the words "Please Scroll" are displayed. Use the arrow keys on the left of the keyboard, and hit space when you are done scrolling. Hitting "S" will change the scroll mode to single pixel, i.e. each time you hit the key, the cursor moves one pixel. Hitting "N" returns to normal scrolling mode. After you have scrolled, it will locate the target, and ask you

ANY MORE TARGETS?

Reply Y for yes, then supply the above information for each target. When you have informed it of all targets, reply N to the last question, and it will again ask you for a command. Now advance the film one frame, and type

B2 (step 2 in building database)

This will ask you to scroll to the new position of each target from the last position, as the targets may be moving rapidly, and it has no idea of each target's velocity. This gives it the velocities of all the targets. Now advance the film and type

B3 (step 3 in building a database)

This will ask you to again scroll to the targets, thus giving it an idea of the acceleration of each target. After you have built this database, you may want to write it on disk. It will only be useful if you are at the current frame in the film again. WD writes database on disk.

The normal sequence is to advance the film and type

NE (process next frame)

This will scan for all the targets, showing you what you specified you wanted to see on the GT40 screen, find all the points, display them and ask for your next command. Be sure to advance the film after each NE command. After 8 frames (including the first 3 used to build the database) you can do an

OO (open disk)

This opens a disk file, and each NE command then writes its points to disk before processing the next frame. When you are done you will want to do a

CO (close disk)

This closes this point file.

Sometimes the routine will incorrectly find a target. When this happens, it becomes necessary for the operator to scroll to the proper position. For this reason, after each NE command, all the points that are found are displayed. If one is incorrect, you should do a

CO (correct one)

This will ask for the target number and allow you to scroll to the proper position of that target. Since the positions are written out at the beginning of the NE command, it is this corrected position that is written out, when you do the next NE command.

Other commands which are useful include

DT display current set of targets
DH displays last 8 positions of a target (it asks for target #)
DA displays last 4 positions of all targets
AD add more targets to database
EP enter picture mode
LP leave picture mode
DO specify display options (what arrays to display on GT48)
RL re-open long point file at end
PA change parameters of radius and threshold
TI time of day
EX exit
? display all commands

APPENDIX 2 Operational Guide to Phase II

This appendix describes the use of the graphics routines which display the points found in phase I. To start the program, type

```
:CATDIS
```

This will load the display routine into the GT40, and when it is done loading, there should be a sketch of a cat and the words

```
ENTER FILE NAME
```

```
#
```

to which you should give the name of a file which is a data file from phase I. The data files which I have collected are on my disk and are named OFICAL.N where N is a decimal number. Since in general there is a phase difference between the first and last set of points in the file, there will be a discontinuous jump in motion at the "wrap-around" point, where the first set of positions follows the last set. However, the files are in ASCII, and thus can be edited to minimize this phase difference. I have done this for most of the files, and the modified files (usually shorter) are called OFICAL.1M, OFICAL.2M, with the 'M' meaning that it is a modified or matched file. After receiving a file name, phase II reads this file into core, and displays the following message:

BIT 15 HALVES CAT BIT 11 TURNS OFF CAT
BIT 14 CLEARS POINTS BIT 10 TURNS OFF TEXT
BIT 13 WIPES OLDEST POINT BIT 9 TO VIEW HIDDEN FEET
BIT 12 ADDS NEWEST POINT BIT 8 CHANGES SCALE OF GRAPHS
 BITS 0-7 DISABLE JOINT DISPLAYS
 ? TO DISPLAY 11/45 COMMANDS

There are three modes to CATDIS. These are normal, graph, and footfall diagram mode. The user starts out in normal mode. The above message tells what the GT40 bit switches do. The sets of points are sent from the 11/45 to the GT40 to be displayed. Some of the bit switches only have meaning when updating with the next set of points. The switches effective at this time are bits 15, 13, and 12. If bit 12 is up, a point will be displayed (in addition to any other points or pictures) at each joint position just received. If bit 13 is up, the oldest such set of joint position points will be flushed. Thus having bits 12 and 13 up will cause the number of points displayed to be constant, i.e. the last n positions will always be visible. If bit 15 is up, the cat will become half size (unless it already is, in which case it will stay small). When changing size, all points on the screen are flushed. With all these points on the screen, it often becomes confusing as to which points are in which joint trajectory. Thus by flipping bits 0-7 (anytime), the set corresponding to that joint nbt be displayed (it is not lost, and putting the bit down redisplayes them). Bit 14 will flush all points currently on the screen. Bit 11 turns off the

sketch of the cat (in order to view the points better), and bit 10 turns off the text. Bit 9 causes the left front and back feet to be viewed. The left feet are the same as the right ones, only delayed by a phase lag (see P command below). Bit 8 changes the scale of the next graph's angular velocity and angular acceleration. Scale goes in a cyclic 1-2-4-6-1 order, with change of state of bit 8 advancing the scale one position. Thus in this mode, the user can watch the point trajectories and locomotion. To get a set of points, the GT40 must interact with the 11/45. There are many other ways to interact with the 11/45. The following is a list of 11/45 commands available.

SPACE This sends over one set of points. Thus hitting the space bar will look like viewing successive frames of the film.

"A" This will send a set of points over every time the clock ticks the *n* times. *N* is initially 6 and can be set with the F command. This auto-mode will show the cat in motion.

"F" This sets the frame rate for the A command, i.e. determines how long to wait between successive frames. One clock tick is 1/60 of a second.

"M" This sets the Move distance, or the offset to be added to each frame. A nonzero value here will cause the cat to move across the screen. A good value to use is the smallest one that doesn't allow any

point to move backwards. See appendix 3 for this speed for each file.

"R" Asks for new file name. Used to read another file.

"P" Set phase lag (in points) between the right feet and the left feet. The left feet are simply the right positions delayed n points. P sets n. Bit 9 must be up to see left legs.

"G" Graph mode. This asks for three joint numbers, and computes the angle, angular velocity, and angular acceleration at each frame and plots this angle for all frames in the disk file. The same joint can be specified as two of the three joints, thus resulting in the angle from the horizontal. Catdis then goes into graph mode in which it displays a graph showing the angle, angular velocity, and angular acceleration of that joint for entire point file. A dotted line shows where on the graph the current position of the cat is. The user must also specify a time constant for computing the velocity and acceleration. If $T=1$, the velocity is the difference in the last two positions. If $T=2$, the velocity displayed is the difference of position i and $i-2$, etc. Only the A,L,S, and SPACE commands can be executed without returning to normal mode. Toggle bit 8 and do an "L" command to change scale.

"L" Review last graph viewed. Useful in returning to graph mode.

"S" Smooths data. The number of times data has been smoothed is given in S:n immediately after file name.

"C" Change scale of points. It asks for numerator and denominator. Needed because some of point files were obtained using the 3 plan microscope probe, and some were obtained using the 2 plan piece. Use 1/1 for the 3 plan files, and 3/2 for the 2 plan files. (see appendix 3 for list of which is which).

"T" Displays date, time, and current version number.

"I" Displays the bit switch information above.

"?" Displays a list of these commands and a brief description.

"D" Diagram mode. This command first asks if the dynamic Muybridge footfall diagram should be turned on. If so, it asks for a Y value for the front foot. Points which are below this value AND which are moving backwards in the original film are considered as touching the ground. By having bit 12 up, typing "A" to let the cat run for a stride or two, and then doing a "D", the user can more accurately place the ground line. The front and back are independently settable in the event that the front and rear markers are a different height above the ground. After having done this, the user is returned to normal mode, but a Muybridge

diagram is visible on the left of the screen. As the cat moves (either with another "A", or successive "SPACE" commands), this diagram will show which feet are on the ground. To turn off this diagram, type another "D" and answer "N" to the first question.

"O" Other Diagrams. O doesn't seem to stand for multiple footfall diagram mode, but most of the good letters were taken when this feature was implemented. This command asks for a start frame and a count. If there is a phase lag of 17, use a starting frame of at least 18 to avoid points from the end of the point file messing up the diagram. It displays the Muybridge footfall diagram for the whole series, and under each diagram gives a count of the number of frames in a row that that diagram was valid. 20 diagrams are the most it can show (different from 20 frames).

Just to get one started, try using the file GALLOP. When the bit switch information is displayed, type an A. The cat will appear to be running on a treadmill. Watch for a while, then put up bit 12. This will cause the points to remain in view. Flip bit 14 up and down to clear the points, and then put bit 13 up (the cat is running this whole time). Now it should appear as if there is almost a string attached to the joints of constant length. Playing with bits 8-7 will cause points streaming from each joint to not be displayed. Now turn on bit 15. The cat is half size, and there is just one point on each joint. By turning bit 13 down for a second, then up, a history of arbitrary length can be created. Set the frame rate to 4 ticks, the move distance to 24 units, and the phase lag to 6 points. Put up bit 9 and 15, type A, and watch the cat run across the screen. Type P to set the phase to 6 and -6 respectively.

APPENDIX 3 Some useful information

The following chart tells some useful information to anyone using or playing with CATDIS.

FILE NAME	DESCRIPTION	LENGTH (FRAMES)	PERIOD	PHASE LAG (FR, BACK)	SIZE	SPEED
OFICAL.1	ROTARY GALLOP	192	35	6,-6	3 plan	24
OFICAL.2	BAD TROT	82	42	20,21	3 plan	12
OFICAL.4	WALK	123	38	19,19	3 plan	10
OFICAL.6	TROT	213	32	16,16	2 plan	20
OFICAL.7	TROT	144	35	17,18	2 plan	18
OFICAL.8	TROT *	100	33	16,17	2 plan	18
OFICAL.9	TROT	84	33	17,16	2 plan	19
OFICAL.10	ROTARY GALLOP	66	36	6,-6	2 plan	26
OFICAL.11	WALK	145	43	22,21	2 plan	11

* Cat's rear legs stop when tail hit back wall of treadmill, (about frame 70) but front ones keep going until cat was against back of treadmill wall.

REFERENCES

Kimme, Ballard, and Sklansky, "Finding Circles by an Array of Accumulators," February 1975, Communications of the ACM, Volume 18 Num 2, p 128.

Eadward Muybridge, Animals in Motion, Lewis Brown, editor, 1959, Dover Publications, Inc. New York.

Eadward Muybridge, The Human Figure in Motion, 1955, Dover Publications, Inc. New York.

Glen Speckert, "Knowledge Driven Recognition of the Human Body," January, 1976 MIT Working Paper 118.

Dr. Mary C. Wetzel, Department of Psychology, College of Liberal Arts, University of Arizona (Produced the 16mm films).