

**METHOD AND APPARATUS FOR MACHINE VISION CLASSIFICATION
AND TRACKING**

Field of the Invention

5 This invention relates generally to systems used
for traffic detection, monitoring, management, and vehicle
classification and tracking. In particular, the invention
is directed to a method and apparatus for classifying and
tracking objects in images provided by real-time video
10 from machine vision.

Background of the Invention

 With the volume of vehicles using roadways
today, traffic detection and management has become ever
15 important. For example, control of intersections,
detection of incidents, such as traffic accidents, and
collection of data related to a traffic scene are all
integral to maintaining and improving the state of traffic
management and safety. Since the 1950s, point detection
20 devices, such as in-ground inductive loops, have primarily
been used for intersection control and traffic data
collection. The in-ground inductive loops basically
consist of wire loops placed in the pavement, detecting
the presence of vehicles through magnetic induction.

25 Many limitations exist with point detection
devices such as the inductive loops. Namely, the
inductive loops are limited in area coverage for each
individual loop, expensive to install, requiring a roadway
to be dug up for their installation, and are difficult to
30 maintain. Further, such point detectors possess
substantial limitations in their ability to accurately
assess a traffic scene and extract useful information
relating to the scene. While point detection devices can
detect the presence or absence of vehicles at a
35 particular, fixed location, they cannot directly determine
many other useful traffic parameters. Rather, they must
determine such parameters through multiple detection and
inference. For instance, to calculate the velocity of a



vehicle, a traffic management system employing point detection devices requires at least two detection devices to determine the time between detection at two points, thereby resulting in a velocity measurement. Other
5 methods of detection, such as ultrasonic and radar detection also possess similar limitations.

A traffic scene contains much more information than point detection devices can collect. While a point detection device can provide one bit of data, a video
10 image can provide a 300,000 byte description of the scene. In addition to the wide-area coverage provided by video images, the image sequences capture the dynamic aspects of the traffic scene, for example at a rate of 30 images a second. Therefore, advanced traffic control technologies
15 have employed machine vision, to improve the vehicle detection and information extraction at a traffic scene. These machine vision systems typically consist of a video camera overlooking a section of the roadway and a processor that processes the images received from the
20 video camera. The processor then attempts to detect the presence of a vehicle and extract other traffic related information from the video image.

An example of such a machine vision system is described in U.S. Patent No. 4,847,772 to Michalopoulos et al., and further described in Panos G. Michalopoulos,
25 *Vehicle Detection Video Through Image Processing: The Autoscope System*, IEEE Transactions on Vehicular Technology, Vol. 40, No. 1, February 1991. The Michalopoulos et al. patent discloses a video detection
30 system including a video camera for providing a video image of the traffic scene, means for selecting a portion of the image for processing, and processor means for processing the selected portion of the image.

The Michalopoulos et al. system can detect
35 traffic in multiple locations, as specified by the user, using interactive graphics. The user manually selects detection lines, which consist of a column of pixels, within the image to detect vehicles as they cross the



detection lines. While the manual placement of the detection lines within the image obviates the expense of placing inductance loops in the pavement as well as provides flexibility in detection placement, the

5 Michalopoulos et al. system still roughly emulates the function of point detection systems. The system still detects vehicles at roughly fixed locations and derives traffic parameters by induction, using mathematical and statistical formulae. For example, the system classifies

10 a vehicle based on its length and calculates velocity of a vehicle based on the known distance between detection locations divided by average travel time. Further, if a vehicle crosses through an area within the image where the user has not placed a detection line, the system will not

15 detect the vehicle. Thus, the system does not automatically detect all vehicles within the image.

Before a machine vision system can perform any traffic management capabilities, the system must be able to detect vehicles within the video images. The

20 Michalopoulos et al. system detects vehicles by analyzing the energy, intensity or reflectivity of every pixel in the predefined detection lines and comparing an instantaneous image at every pixel with a threshold derived from analysis of the background scene without the

25 presence of any vehicles.

Other systems have utilized edge detection for detecting vehicles. These systems often perform "blob analysis" on the raw image, which constitutes a grouping of elements. The goal of such an analysis is determining

30 which pixels belong together, based on pixel location, intensity and previous grouping decisions. The basic process may be described as region growing. First, the system picks a center pixel that it determines belongs in a grouping. Then, the system looks to neighboring pixels

35 and determines whether to include the pixels in the grouping. This process continues for each included pixel. Blob detector of this type have run into difficulties because all the decisions are interdependent. Once the



system has made initial decisions to include or exclude pixels, subsequent decisions will be based on the decisions already made. Thus, once the system makes an incorrect decision, future decisions are often also incorrect. This series of incorrect decision making may lead to failure of proper convergence. The same is true of edge detection based systems which rely on sequential decision processes.

A further desirable capability of machine vision systems is the capability to track the detected vehicles. Systems that track vehicles usually share some common characteristics. First, the system must identify the starting point of the track. The system may do this by detecting the vehicle by comparing an input image with a background image and judging objects having an area within a predetermined range as vehicles. Other systems perform motion detection to initiate the tracking sequence. Those systems using motion alone to initiate tracking are prone to errors because they must set some baseline amount of motion to initiate tracking. Thus, it is always possible for systems to fail to track slow moving or stalled vehicles.

After identifying a starting point, the systems perform a searching sequence. The systems have a current vehicle location, initially, the starting point. Then they look for potential displacement locations. The systems compare the potential displacement locations and select the location with the greatest suitability. They determine suitability by extracting a subimage region surrounding the current track location. Then, they displace the entire subimage region to potential new locations on the subsequent image frame. Thus, the systems perform a displacement of location and time. The systems perform a pixel-by-pixel correlation to determine which location's image best "matches" the previous location's image. This type of correlation runs into limitations because the system treats the background pixels the same as the pixels of the moving vehicle,



thereby causing problems with matching. Further, since all pixel intensities are weighted equally in importance, large areas of uniformity, such as the hood of a vehicle, are redundant. In such areas of uniformity, the system will be able to match a majority of pixels, but still may not line up the boundaries of the vehicle. While the edges of the vehicle constitute a minority of the pixels, they are the pixels that are most important to line up.

Traffic detection, monitoring and vehicle classification and tracking all are used for traffic management. Traffic management is typically performed by a state Department of Transportation (DOT). A DOT control center is typically located in a central location, receiving video from numerous video cameras installed at roadway locations. The center also receives traffic information and statistics, from sensors such as inductive loop or machine vision systems. Traffic management engineers typically have terminals for alternately viewing video and traffic information. They scan the numerous video feeds to try to find "interesting scenes" such as traffic accidents or traffic jams. It is often difficult for traffic management engineers to locate a particular video feed which has the most interesting scene because they must perform a search to locate the video line containing the video feed with the interesting scene. Current traffic management systems also generate alarms based on inferred trends, which tell the traffic management engineers the location of a potentially interesting scene. Because the systems infer trends at a location, the systems require time for the trend to develop. Thus, a delay is present for systems which infer trends. After such delay, the traffic management engineers can then switch to the correct video feed.

Summary of the Invention

To overcome the limitations in the prior art described above, and to overcome other limitations that will become apparent upon reading and understanding the



present specification, the present invention provides a method and apparatus for classifying and tracking objects in an image. The method and apparatus disclosed can be utilized for classifying and tracking vehicles from a plurality of roadway sites, the images from the sites as provided in real-time by video cameras. The images from the real-time video are then processed by an image processor which creates classification and tracking data in real-time and sends the data to some interfacing means.

The apparatus of the present invention includes a plurality of video cameras situated over a plurality of roadways, the video cameras filming the sites in real-time. The video cameras are electrically interconnected to a switcher, which allows for manual or automatic switching between the plurality of video cameras. The video is sent to a plurality of image processors which analyze the images from the video and create classification and tracking data. The classification and tracking data may then be sent to a workstation, where a graphical user interface integrates the live video from one of the plurality of video cameras with traffic statistics, data and maps. The graphical user interface further automatically displays alarm information when an incident has been detected. The classification and tracking data may further be stored in databases for later use by traffic analysts or traffic control devices.

The present invention provides for a method for classifying vehicles in an image provided by real-time video. The method first includes the step of determining the magnitude of vertical and horizontal edge element intensities for each pixel of the image. Then, a vector with magnitude and angle is computed for each pixel from the horizontal and vertical edge element intensity data. Fuzzy set theory is applied to the vectors in a region of interest to fuzzify the angle and location data, as weighted by the magnitude of the intensities. Data from applying the fuzzy set theory is used to create a single vector characterizing the entire region of interest.



Finally, a neural network analyzes the single vector and classifies the vehicle.

After classification, a vehicle can further be tracked. After determining the initial location of the classified vehicle, potential future track points are determined. Inertial history can aid in predicting potential future track points. A match score is then calculated for each potential future track point. The match score is calculated by translating the initial location's region onto a potential future track point's region. The edge elements of the initial location's region are compared with the edge elements of the future track point's region. The better the edge elements match, the higher the match score. Edge elements are further weighted according to whether they are on a vehicle or are in the background. Finally, the potential future track point region with the highest match score is designated as the next track point.

Brief Description of the Drawings

In the drawings, where like numerals refer to like elements throughout the several views;

Figure 1 is a perspective view of a typical roadway scene including a mounted video camera of the present invention;

Figure 2 is a block diagram of the modules for an embodiment of the classification and tracking system of the present invention;

Figure 3 is a flow diagram of the steps for classifying a vehicle;

Figure 4 is a top-view of a kernel element used in the classification process;

Figure 5 is a graphical representation of an image of a scene, illustrating the placement of potential regions of interest;

Figure 6A and 6B are graphs used to describe the angle fuzzy set operator;

This Sheet Has Been Replaced



Figure 7A is a top-view of a location fuzzy set operator as used in a preferred embodiment;

Figure 7B illustrates the placement of location fuzzy set operators with respect to the region of interest;

Figure 8 illustrates the process of organizing information from location and angle fuzzy set theory in matrix form;

Figure 9 illustrates the placement of icons on classified vehicles;

Figure 10 is a flow diagram of the steps for tracking a vehicle;

Figure 11 is a graphical representation of an image of the scene, illustrating the placement of potential future track regions;

Figure 12 is a diagram of a preferred embodiment of the system of the present invention;

Figure 13 is a graphical representation of an image displayed by the monitor of a graphical user interface; and

Figure 14 is a graphical representation of an image displayed by the monitor of a graphical user interface, illustrating an alarm message.

Detailed Description of the Preferred Embodiment

In the following detailed description of the preferred embodiment, reference is made to the accompanying drawings which form a part hereof; and in which is shown by way of illustration of a specific embodiment of which the invention may be practiced. It is to be understood that other embodiments may be utilized and structural changes may be made without departing from the scope of the present invention.

The fundamental component of information for machine vision systems is the image array from a scene of a specific section of roadway as provided by video. Figure 1 illustrates a scene where video camera 2 is positioned above roadway 4 viewing scene 6. Scene 6



contains various stationary items such as trees 7, barrier 8, light poles 9 and position markers 10. Scene 6 also may contain moving objects such as vehicles 12. Video camera 2 is electrically coupled, such as by electrical or fiber optic cables, to electronic processing equipment 14 located locally, and further transmits information to centralized location 16. Video camera 2 can thereby send real-time video images to centralized location 16 for use such as viewing, processing, analysis or storage.

Image information in the form of digitalized data for each pixel of an electronic video image of scene 6 is processed according to the flow diagram as illustrated in Figure 2. For example, the image array may be a 512x512 pixel three color image having an integer number defining intensity with a definition range for each color of 0-255. If image information is not in digitized form, video image preprocessor module 12 will digitize the image information. The camera image input is subject to environmental effects including roadway vibration, wind, temperature change and other destabilizing factors. To counter the undesirable effects of camera motion, video image preprocessor module 12 electronically performs image stabilization. Reference markers 10 are mounted within the view of video camera 2. Using frame to frame correlation with respect to reference markers 10, compensating translation and rotation is calculated. The appropriate warping of the image may be performed in real time by machines such as Datacube Corporation's (Danvers, Massachusetts) "Miniwarper". Video image preprocessor 12 may then calibrate the stabilized video image information by mapping the real world measurements of the image to the pixel space of the image. Video image preprocessor 12 may further perform background subtraction, eliminating any image information not associated to a vehicle. Thus, the image is segmented into vehicle related pixels and nonvehicle/background pixels. A preferred embodiment of a method and apparatus for background subtraction for use with the present invention is described in commonly-



assigned U.S. patent application entitled "Method and Apparatus for Background Determination and Subtraction for a Monocular Vision System" and identified by attorney docket number 49805USA1A, filed on even date herewith and now U.S. Patent Application Serial Number 08/163,422.

Vehicle identification and classification module 24 and vehicle tracking module 22 then process the stabilized video image information. Vehicle identification and classification may be used for vehicle tracking or may be directly output as data for further analysis or storage. The results of vehicle classification module 24 and vehicle tracking module 22 are consolidated in traffic interpreter 26, which may include a user interface and the system central processing unit. Results of the image processing are then available for other modules, such as traffic data module 27, control module 28, or alarm information module 29 which signals abnormalities in the scene.

Figure 3 is an image processing flow diagram of the data flow for vehicle classification. At module 40, the video image is digitized and stabilized, eliminating noise due to vibration and other environmental effects for a discrete image array a at time t_a . The discrete image array a may consist of a matrix of numbers, such as a 512 x 512 pixel image having a integer defining the intensity of each pixel, with a definition range for each color of 0-255. Successive image arrays would be $a+1$ at time $t_{(a+1)}$ etc,

At edgel definition module 41, each pixel of the image array output from the stabilized image is evaluated for the magnitude of its edge element (edgel) intensity. Edgel intensity indicates the likelihood that a given pixel is located on some edge having particular orientation and contrast. The greater the contrast between a particular pixel and the pixels surrounding it in a particular orientation, the greater the edgel intensity. An edge differs from an edgel in that an edge is a more global phenomena involving many edgels. Edgel



definition module 41 takes the data of the image array and produces two edgel images for each pixel. A first edgel image represents the likelihood that each pixel lies on a horizontal edge, or the degree of horizontal edgeness at each pixel, x_edgel , calculated according to equation 1.

$$x_edgel_{ij} = \sum_{u,v=-k}^{+k} \left(\frac{\text{sign}(v)}{2\pi\sigma^2} e^{-(u^2+v^2)/2\sigma^2} \right) * I_{(i+u)(j+v)}$$

Equation 1

Within equation 1, $\text{sign}(v)$ is +1 when v is positive and -1 when v is negative, $I_{(i+u)(j+v)}$ are the pixel intensities surrounding pixel (i,j) and the kernel is of size $2k+1$ by $2k+1$ where k is equal to two in one embodiment of the system. The second image represents the likelihood that a pixel lies on a vertical edge, or the degree of vertical edgeness, y_edgel , calculated according to equation 2.

$$y_edgel_{ij} = \sum_{u,v=-k}^{+k} \left(\frac{\text{sign}(u)}{2\pi\sigma^2} e^{-(u^2+v^2)/2\sigma^2} \right) * I_{(i+u)(j+v)}$$

Equation 2

Within equation 2, $\text{sign}(u)$ is +1 when u is positive and -1 when u is negative. Therefore, edgel detection module 41 determines the likelihood that each pixel within the image array lies on a horizontal or vertical edge.

Figure 4 shows a plot of a sample 8×8 kernel used to calculate edgel intensities. Edgel detection module 41 successively applies kernel 59 to each pixel within the image array to perform convolution. The convolution takes into account the pixels surrounding the pixel in question, thereby determining the likelihood that the pixel in question is located on an edge. Edgel detection module 41 replaces the original pixel value with

This Sheet Has Been Replaced



the outputs from applying kernel 59 twice in two orientations, resulting in two new integers representing both the horizontal and vertical edgeness of the pixel.

Edgel definition module 41 sends forward the horizontal edgeness data (x_edgel) and vertical edgeness data (y_edgel) of the array a to geometric transform module 45. Geometric transform module 45 converts the discrete pixel data pairs (x_edgel,y_edgel) from degree of horizontal edgeness and vertical edgeness values to a vector with direction and magnitude for each pixel (i,j). The direction may be expressed in angle format while the magnitude may be expressed by an integer. In a preferred embodiment, the angle is between 0 - 180 degrees while the intensity may be between 0 - 255. The transform of data is analogous to transforming rectangular coordinates to polar coordinates in a Euclidean space. The geometric transform is performed, in a preferred embodiment, according to equations 3 and 4. The magnitude value is a calculation of total edgel intensity (sum_edgel) of each pixel, and is calculated according to equation 3.

$$sum_edgel_{ij} = |x_edgel_{ij}| + |y_edgel_{ij}|$$

Equation 3

The angle value α developed in geometric transform module 45 for pixel i,j is calculated according to equation 4.

$$\alpha_{ij} = 0 \quad \text{if} \quad y_edgel_{ij} = 0$$

$$\alpha_{ij} = \frac{90}{1 + \frac{x_edgel_{ij}}{y_edgel_{ij}}} \quad \text{if} \quad \frac{x_edgel_{ij}}{y_edgel_{ij}} \geq 0,$$



$$\alpha_{ij} = 180 - \left(\frac{90}{1 + \frac{x\text{-edgel}_{ij}}{y\text{-edgel}_{ij}}} \right) \quad \text{if} \quad \frac{x\text{-edgel}_{ij}}{y\text{-edgel}_{ij}} < e$$

5

Equation 4

10 The angle and magnitude data is then sent to region selection module 46. Geometric transform module 45 also may send forward the magnitude data representing the pixel intensity to a vehicle tracking module.

15 Region selection module 46 must define potential regions of interest, or candidate regions, for classification. Referring to Figure 5, image 30 includes roadway 4 and roadway boundaries 31. Region selection module defines candidate regions 35. Each vehicle class is assigned a set of appropriately sized and shaped regions. In a preferred embodiment, candidate regions are substantially trapezoidal in shape, although other shapes may be used. Candidate regions may be dynamically generated according to prior calibration of the scene. Regions of interest are overlapped within the image space. In a preferred embodiment, overlapping regions are confined to the lower half of the image space. In the illustrated example, the candidate regions of interest are truck sized trapezoids, with some centered in the traffic lanes and others centered on the lane marker overlapping the traffic lanes. The number of overlapping trapezoids may vary, depending on the coverage preferred. Another set of overlapping trapezoids 36 could be, for example, car sized. Classification regions may be predefined.

30 Region selection module 46 scans the predefined regions of interest to determine presence of vehicles, selecting regions for analysis based on a minimum threshold of change in the average edgel value within a candidate region. The minimum threshold value only need be high enough such that changes in average edgel intensity will not rise above the threshold value due to noise within the region. Region selection module sums the



individual sum_edgel values for all of the pixels for each candidate region and calculates an average sum_edgel value. Region selection module 46 then compares the candidate regions of array a with the regions of frames a-1 and a+1, which are the previous and subsequent frames. The average sum_edgel value is compared to the like average sum_edgel value for the previous and subsequent frames. When comparison indicates that a local maximum for the sum_edgel value has been reached, or local minima, if some different criteria is used, the candidate region becomes a region of interest. Once a region of interest has been identified, the region priority module 47 selects the region of interest and sends it forward for vehicle identification by class.

While the data available would be sufficient for classification purposes, the amount of data is too voluminous for real-time processing. Further, there is great redundancy in the data. Finally, the data has too little invariance in both translation and rotation. Therefore, the system of the present invention reduces the amount of data, reduces the redundancy and increases the invariance of the data by applying fuzzy set theory. Vectorization module 51 converts the geometrically transformed data to vector data by applying fuzzy set theory to the transformed data. Vectorization module 51 may apply separate fuzzy set operators on both the location and the angular characteristics of each geometrically transformed edgel. Vectorization module 51 determines a vector which characterizes the entire region of interest, the vector which contains sufficient information for classification of any vehicle within such region.

In a preferred embodiment of the classification process, vectorization module 51 applies ten overlapping wedge-shaped operators 66, as shown in Fig. 6, to the angle components of the geometrically transformed edgels for angle fuzzification. Each wedge-shaped operator 66, has a width of 36 degrees and a unit height. When



overlapping wedge-shaped operators 66 are applied to the angle components of each edgel, operators 66 determine to what extent the angle should "count" toward each angle set. Each wedge-shaped operator is centered at a particular angle. For each angle component not falling directly under the apex of a wedge, i.e. at an angle that a wedge is centered at, then the angle will be counted toward both wedges it falls under, taking into account which angle that each wedge is centered at is more like the angle component as well as the magnitude of the intensity of the edgel. For example, an angle centered exactly at 18 degrees will have full set membership in the angle set centered at 18 degrees. On the other hand, for the 20 degree angle shown in Figure 6B, the angle will count more towards the angle set centered at 18 degrees and less towards the angle set centered at 36 degrees. After applying this fuzzy logic operator to all the angle components of the edgels in the region of interest, the resulting output is ten angle sets representing all the angle components of the geometrically transformed edgels.

While wedge-shaped operator 66 fuzzifies the angle characteristics of the edgels, the location of the angles of each edgel still must be taken into account. Thus, a location fuzzy set operator must be applied to the region of interest to determine the general location of the angle components of the edgels. In a preferred embodiment, tent function 60, as shown in Figure 7A, is used as the fuzzy set operator for location. The tent function 60 has a unit height. Tent function 60 performs a two dimensional fuzzification of the location of the edgels. The input variables (i,j) represent the pixel's location within the region of interest. Figure 7B illustrates placement of tent functions 60. For each region of interest 200, nine tent functions 60 are placed over region of interest 200. Center tent function 202 is centered over region of interest 200 and entirely covers region of interest 200. Top-right tent 204 is centered on the top-right corner of center tent 202. Top-center tent



206 is centered on the top-center edge of center tent 202. Right-center tent 208 is centered on the right-center edge of center tent 202. Top-left tent, left-center tent, bottom-left tent, bottom-center tent, and bottom-right tent are similarly placed. (Not shown in Figure 7B). At any point within region of interest 200, the sum of the heights of all overlapping tents is the unit height.

5 The tent functions are applied to each edgel within the region of interest. For each tent function a histogram is produced which records the frequency that a range of angles occurs, as weighted by the intensity of the edgels, within a particular tent function. For edgel 10 210, the histograms for center tent 200, top-center tent 206, top-right tent 204 and right-center tent 208 will all 15 "count" some portion of edgel 210, as determined by the height of each tent over edgel 210. After determining what proportion of edgel 210 is located under each tent, the angle sets associated with edgel 210, as determined by the wedge-shaped fuzzy set operators, determine which 20 angle range within the histogram the edgel value must be counted. Then, the edgel value is weighted by the magnitude of the edgel intensity, as calculated in geometric transform module 50, and counted in each histogram associated with the tent functions the edgel 25 falls under. This process is repeated for all edgels within region of interest. The result of the location fuzzification is nine histograms, each characterizing the frequency that ranges of angles are present, as weighted by magnitude, for a general location. Each histogram is a 30 vector of dimension 10. The nine location histograms are then strung together to form a vector of dimension 90, which is output from vectorization module 51.

35 In another embodiment of the present invention, a 10 x 9 matrix may be used instead of histograms for the fuzzification process. In this embodiment, each pixel is evaluated and proportionally allocated to positions within the matrix. In Fig. 8, matrix 250 is shown. Matrix 250 has ten rows, representing ten angle sets of 18 degrees



each, and nine columns, representing nine tent locations, such as center tent, top-right tent, etc. A pixel at location (i,j) is evaluated for its angle 260 and its magnitude 270. After the location of pixel (i,j) is subjected to the tent functions 262 for location fuzzification, the location is allocated to up to four columns of matrix 250, depending on the number of tent functions pixel (i,j) falls under. Angle 260 is then weighted by magnitude 270 of pixel (i,j) and subjected to wedge shaped operator 272 for angle fuzzification. The angle, as weighted by magnitude, is allocated up to two adjacent rows of matrix 250, depending on the number of wedge-shaped operators angle 260 falls under. This process is repeated for all pixels and the nine columns are placed end-on-end to create a vector of dimension 90, which is output from vectorization module 51.

Vehicle classification module 52 includes a neural network. The neural network is trained to recognize predetermined classes of vehicles by vehicle learning module 53. The vector of dimension 90, as developed in vectorization module 51 is evaluated by the neural network. The neural network then classifies the vehicle based on the vector and may generate a signal indicative of the classification of the vehicle. A preferred embodiment of a neural network for use with the present invention is described in commonly-assigned U.S. patent application entitled "Facet Classification Neural Network" and identified by attorney docket number 49597USA4A, filed on even date herewith and now U.S. Patent Application Serial Number 08/163,825.

A unique icon identifying the vehicle class may be assigned to a classified vehicle in icon assignment module 54. The icon assignment further may be output to the tracking module for facilitating visual tracking. Assignment of definitive classification icons to the individually classified vehicles provides unique visual output for oversight of system operation and specific identification for tracking. Figure 9 is one frame of the



visual image scene 70 cross shaped icon 72 located on car image 74 and star shaped icon 76 on truck image 78. Any icon may be chosen for identification, such as the characters "T" for trucks and "C" for cars. These icons will move with progression of the vehicle as the track progresses over time through multiple frames until the vehicle is no longer identifiable in the scene.

After detection and the associated classification of a vehicle, the portion of the video image containing the identified vehicle takes on a unique identity within the array a for time t_a . The unique identity includes dynamic region-of-interest boundaries, pixel data computations within those boundaries and assignment of the appropriate icon. A display of the video image may include the icon on the video screen of the user interface, if the user desires, and will be displayed centrally located within the assigned boundaries. In subsequent images, arrays $(a+1), (a+2) \dots (a+n)$ for times $t_{(a+1)}, t_{(a+2)} \dots t_{(a+n)}$ will develop the track of the vehicle. The tracking sequence will continue until the vehicle reaches a predetermined row of the array.

Referring to Figure 10, a flow diagram for vehicle tracking is shown. Selected modules of the tracking sequence receive information from modules from the classification process, as shown in Figure 4, including: edgel intensities from geometric transform module 45 and vehicle identification and classification data 55.

The tracking process starts with identification of a vehicle within the scene. Thus, at the initiation of the process or after a lull in traffic, existing vehicle module 81 will have no identified vehicles in memory. If the classification process has identified a vehicle in frame a at time t_a , it is fed forward as a new vehicle 55 to decision node 82. When decision node 82 receives information relating to a new vehicle, it generates a new tracking sequence for the new vehicle. Once a tracking



sequence has been established for a vehicle, decision node 82 processes the identified vehicle's track according to information from existing vehicle module 81. When decision node 82 receives new vehicle information, as
5 opposed to existing vehicle information from existing vehicle module 81, a new vehicle initialization format is fed forward to vehicle match score module 84. Alternatively, where decision node 82 feeds forward information on an existing vehicle, the existing vehicle
10 format is matched with information from traffic supervisor module 100.

Traffic supervisor module 100 maintains oversight of the scene, including keeping a vehicle log 102 of all vehicles currently within the scene, including
15 vehicles' associated track histories, inertial history 104 of the scene, accumulated and normalized over time, and potential future track positions of vehicles from candidate generator 106. Once a previous track point, or source track point, has been identified for either a new
20 vehicle or an existing vehicle, candidate generator 106 generates a range of potential future track points, or target track points. While target track points can include any areas in the scene, including areas in front of, to the side of and behind the source track point, in a
25 preferred embodiment, candidate generator 106 takes into account the inertial history of the vehicle, as received from inertial history module 104, to help predict the vehicle's next location. The target regions overlap, the centers of the regions in close proximity, preferably
30 lying only a pixel apart.

An example of the predicted future positions, or target regions, is graphically shown as image 120 in Figure 11. In the lower part of image 120, source region 130 is shown in time frame t_a 126 which includes
35 identified vehicle 122 marked with the appropriate icon 124. In the upper part of image 120, predicted future track points 140a, 140b and 140n, are shown for time frame



$t_{(a+1)}$ 128 with proportionately shrunken track regions 142a, 142b and 142n.

Data reduction is desirable to increase the speed of processing the information for tracking.

5 Reducing the resolution of an region of interest facilitates data reduction. Pixel averaging followed by subsampling is used to reduce resolution. For example, if a 2x2 kernel is used, averaging the pixel intensities of an image over the four pixels in the kernel, the
10 resolution of the target region is reduce by one-quarter. A multiresolution pyramid, with layers of images of decreasing size and resolution, can be created with multiple applications of the kernel. Thus, the target region can be searched for at a lower resolution image to
15 identify areas where the target region is likely located before searching in the same areas in a higher resolution image.

After the target regions have been identified, a match score is calculated for each target region. The
20 source region is overlaid over each target region. The match score takes into account edge elements from the translated source regions that match edge elements from a target region. The better the edge elements match, the higher the match score. Then, edge elements are weighted
25 according to evidence of motion, as determined by the amount the contrast is changing with respect to time at a given pixel, such that vehicle edge elements are maximized and background edge elements are minimized. The target region with the largest match score is the next track
30 point and becomes the new source region. Vehicle match score generator 84 takes as input frame subtraction data 83 which subtracts the magnitude of edgel values from consecutive frames, thereby creating a difference image. Vehicle match score generator 84 also receives target
35 regions from candidate generator 106 in addition to the source region from decision node 82. Vehicle match score generator 84 then calculates a match score for each target region using Equation 5.

This Sheet Has Been Replaced



This Sheet Has Been Replaced



$$M = \sum [\sqrt{S_k(R_k) * R_{s,(u,v)}} * (255 - |R_{s,(u,v)} - S_k(R_c)|)]$$

Equation 5

In Equation 5, R_c is the matrix of the current region's
 5 magnitude of edge element intensities. $R_{s,(u,v)}$ is the
 matrix of the subsequent region's magnitude of edge
 element intensities with translation (u,v) of the region's
 center. $S_k(R)$ is one of several shrinks of a region R . R_s
 is the current track-region in the difference image. The
 10 algorithm proceeds by finding a k and a displacement (u,v)
 such that the solution M is maximized. Note that the
 notation is interpreted such that subtraction,
 multiplication, square root, and absolute value are all
 components-wise over the matrices operated on. $\sum$ simply
 15 sums all the elements on a given matrix.

Maximum score decision node 87 compares results
 obtained from match score generator 84. The maximum score
 from this comparison is identified and the target region
 may be designated as the new location of the vehicle by
 20 new track point module 89. New track point 91 is supplied
 to traffic supervisor 100 and to the user
 interface/traffic interpreter 92. In one embodiment of
 the invention, intermittent tracking can be employed such
 that new track point module 89 only provides a new track
 25 point if a tracked object has moved significantly. New
 track point module 89 compares the current target region
 with a reference of the same region from a previous time
 frame and determines if the difference between them
 exceeds a threshold value. If it does, then new track
 30 point module 89 supplies a new track point at the new
 location of the vehicle and the reference buffer is
 updated to hold the contents of the current region of
 interest. Intermittent tracking increases the signal-to-
 noise ratio of the data used for tracking, and facilitates
 35 tracking in slow or congested traffic conditions. User

interface/traffic interpreter 92 can be any of a number of systems, either at the location of the scene or remotely located at a traffic control center. Traffic interpreter 92 can provide output data, traffic control information, or alarms for storage, traffic control, or use by a traffic control system of a manual user.

Referring to Figure 12, a vehicle classification and tracking system will now be described. While a vehicle classification and tracking system can be used at independent roadway scenes for providing data, traffic control information, or operational alarms, it is particularly useful in monitoring traffic from many roadway sites and integrating the analysis of a large volume of vehicles as they pass through the numerous sites. Video cameras 210A, 210B, 210C, . . . , 210N simultaneously monitor traffic conditions from several roadway sites. Video from the video cameras can be transmitted in a variety of ways such as a commercial network 212, dedicated cable 214, multiplexed in a microwave network 216 or a fiber optic network 218. The video is then processed. In a preferred embodiment, each video camera 210 sends video to a separate image processor 240, although a single image processing means could receive and process the video from all the cameras. Image processors 240A, 240B, 240C, . . . , 240N process the video in real-time and create classification and tracking data. A traffic interpreter may also reside in image processors 240, or in workstation 250, for further generating alarm data, traffic statistics and priority data in real-time. Traffic statistics such as volume, velocity, occupancy, acceleration, headway, clearance, density, lane change count or other characteristics may be computed. Further, traffic statistics may be computed according to lane, vehicle class or both, as well as per interval or over a time period. Image processors 240 then send the data to workstation 250. Video from video cameras 210 is also received by workstation 250 and the live video displayed on display 220. The live video can also be integrated



with associated traffic information for a highway site in a single screen display. The live video can be switched from scene to scene with video switcher 222 under manual control of a user watching display 220 or may
5 automatically be switched by the command of workstation 250 in response to alarm data or priority data generated by image processors 240. Video, traffic statistics and other traffic related data may all be stored in databases in workstation 250 or sent to a control center 260 for
10 storage and later use, such as studying traffic patterns.

A graphical user interface of the video image processing system for vehicle classification and tracking is shown in Figure 13. The graphical user interface aids traffic engineers and analysts by presenting live video of
15 a traffic site and all relevant traffic information for that site in a single integrated display. The traffic engineers may choose any one of a plurality of video feeds to view, along with its associated traffic information. Further, when image processors detect incidents, the users
20 are automatically notified through the graphical user interface. Video screen image 300 includes a live video window 310 of one scene in the highway system. The graphical user interface is window-based and menu-driven. Users make requests by selecting options from menus or by
25 selecting push buttons. Besides video window 310, traffic statistics, maps, or other information may be shown in a windows-type format and may be selectively chosen, sized and arranged, at the discretion of the traffic engineer. System map 320 shows location of video cameras with
30 respect to the roadway layout, and summary traffic information 330 for the current scene being reviewed is displayed graphically in real-time. The total display can further include specific operating switches 340 or operating menus within the same screen display. Real-time
35 alarm messages, based on alarm information generated by image processors or the traffic interpreter, are automatically generated and displayed as shown in Figure 14. Alarm conditions may include, for example, off-road



vehicles, excessive lateral vehicular acceleration, wrong way vehicle or excessive weaving. The video screen 350 of the graphical user interface can be overwritten with alarm window 360 with appropriate switches 362.

5 Although a preferred embodiment has been illustrated and described for the present invention, it will be appreciated by those of ordinary skill in the art that any method or apparatus which is calculated to achieve this same purpose may be substituted for the specific configurations and steps shown. This application is intended to cover any adaptations or variations of the present invention. Therefore, it is manifestly intended that this invention be limited only by the appended claims and the equivalents thereof.

15

20

25

30

This Sheet Has Been Replaced



35



CLAIMS:

1. A machine vision system comprising:

5 a) image acquisition means for acquiring images
from three-dimensional space;

 b) processing means for receiving and
processing said images electrically interconnected to said
10 image acquisition means, said image processing means
programmed to execute the steps of:

 i) determining the magnitude of vertical and
horizontal edge element intensity components of
15 each pixel of said image;

 ii) computing a first vector with magnitude of
total edge element intensity and angle for each
said pixel within said image;

20 iii) applying fuzzy set theory to said first
vectors in regions of interest to create a
second vector characterizing each said region of
interest;

25 iv) interpreting said second vector of each
said region of interest with a neural network
whereby said neural network determines a
classification of said object based on said
30 second vector; and

 v) generating a signal indicative of said
classification of said object; and

35 c) interface means electrically interconnected
to said image processing means.

This Sheet Has Been Replaced





2. The machine vision system according to claim
1, wherein said image processing means further comprises
traffic interpreting means for analyzing signals
indicative of classification generated by said image
5 processing means and for converting said signals to alarm
and statistical information.

3. The machine vision system according to claim
1, wherein said image processing means further comprises
10 means for mapping pixel space to real world measurements.

4. A machine vision system comprising:

a) image acquisition means for acquiring images
15 from three-dimensional space;

b) processing means for receiving and
processing said images, electrically interconnected to
said image acquisition means, said processing means
20 programmed to execute a program for tracking objects
having a first track point in a first image, said program
comprising the steps of:

i) determining a magnitude of vertical and
25 horizontal edge element intensity components of
each pixel of said first image and a second
image;

ii) computing a magnitude of total edge element
30 intensity for each said pixel within said first
image and said second image;

iii) creating a difference image by determining
the difference between the magnitudes of total
35 edge element intensity for each said pixel
within said first image and said second image;

This Sheet Has Been Replaced





iv) forecasting a plurality of potential subsequent track points from said second image;

5 v) assigning a value for each said potential subsequent track point, said value being determined by comparing the edge elements from a first region surrounding said first track point with those in a region surrounding each said potential subsequent track point;

10

vi) selecting one of said potential subsequent track points as said subsequent track point based on said values assigned; and

15

vii) generating a signal indicative of said subsequent track point; and

c) interface means electrically interconnected to said processing means.

20

5. The machine vision system of claim 4, wherein said processing means creates said first track point by executing a program comprising the steps of:

25

a) computing a first vector with a magnitude of total edge element intensity and angle for each said pixel within said first image;

30

b) applying fuzzy set theory to said first vectors in regions of interest to create a second vector characterizing each said region of interest;

35

c) interpreting said second vector of each said region of interest with a neural network whereby said neural network determines a classification and a location of said object based on said second vector; and





d) generating a second signal indicative of said location of said object, said location constituting said first track point.

5 6. A method for classifying objects in an image, comprising the steps of:

 a) determining the magnitude of vertical and horizontal edge element intensity components of each pixel
10 of said image;

 b) computing a first vector with magnitude of total edge element intensity and angle for each said pixel within said image;
15

 c) applying fuzzy set theory to said first vectors in regions of interest to create a second vector characterizing each said region of interest;

20 d) interpreting said second vector of each said region of interest with a neural network whereby said neural network determines a classification of said object based on said second vector; and

25 e) generating a signal indicative of said classification of said object.

30 7. The method for classifying objects in an image according to claim 6, further comprising the steps of:

 a) dividing said image into overlapping regions; and

35 b) selecting said regions of interest from said overlapping regions by selecting regions which meet predetermined criteria.

This Sheet Has Been Replaced





8. The method for classifying objects according to claim 7, wherein selecting said regions of interest which meet predetermined criteria comprises the steps of:

- 5 a) calculating an average total edge element intensity of a candidate region based on the average of said magnitude of total edge element intensity of all pixels within said candidate region;
- 10 b) comparing said average total edge element intensity of said candidate region with a baseline average edge element intensity value; and
- 15 c) selecting said candidate region to be a region of interest if said total edge element intensity of said candidate region exceeds said baseline edge element intensity value.

20 9. The method for classifying objects according to claim 7, wherein selecting region of interest which meet predetermined criteria comprises the steps of:

- 25 a) calculating an average total edge element intensity of a candidate region based on the average of said magnitude of total edge element intensity of all pixels within said candidate region; and
- 30 b) selecting said candidate region to be a region of interest if said average total edge element intensity of said candidate region is a local maxima.

35 10. The method for classifying objects in an image according to claim 6, wherein determining the magnitude of vertical and horizontal edge element intensity components of each pixel of said image comprises the steps of:

- a) defining a first kernel for convolution;





b) defining a second kernel for convolution;

5 c) applying said first kernel to each pixel within said region of interest, thereby computing said horizontal edge element intensity value;

10 d) applying said second kernel to each pixel within said region of interest, thereby computing said vertical edge element intensity value; and

e) assigning each pixel within said region of interest said horizontal and vertical edge element intensity values.

15 11. The method for classifying objects in an image according to claim 6, wherein applying fuzzy set theory to said first vectors in each said region of interest comprises:

20 a) applying fuzzy set theory according to the location of each said pixel in said image;

b) applying fuzzy set theory according to the angle of each said pixel; and

25 c) weighting the angle associated with each said pixel according to said magnitude of intensity.

30 12. The method for classifying objects in an image according to claim 11, wherein applying fuzzy set theory according to the location of each said pixel in said image comprises the steps of:

35 a) placing overlapping tent functions over each said region of interest, said tent functions substantially pyramidal in shape and each said tent function weighting each said pixel beneath said tent function according to

This Sheet Has Been Replaced





the height of said tent function at the location of said pixel;

5 b) apportioning said intensity of each said pixel to generalized locations corresponding to the locations of said tent functions, said apportionment according to said weighting of said pixels.

10 13. A method for determining a subsequent track point in a second image of an object having a first track point in a first image, comprising the steps of:

15 a) determining a magnitude of vertical and horizontal edge element intensity components of each pixel of said first image and said second image;

20 b) computing a magnitude of total edge element intensity for each said pixel within said first image and said second image;

25 c) creating a difference image by determining the difference between said magnitudes of total edge element intensity for each said pixel within said first image and said second image;

 d) forecasting a plurality of potential subsequent track points from said second image;

30 e) assigning a match value for each said potential subsequent track point, said value being determined by analyzing said difference image and said magnitudes of total edge element intensity for each said pixel within said first image and said second image; and

35 f) selecting one of said subsequent track points as said subsequent track point based on said match values assigned.

This Sheet Has Been Replaced





14. The method for determining a subsequent track point of an object having a first track point in a first image according to claim 13, wherein assigning said match value for each said potential subsequent track point comprises the steps of:

5 a) comparing pixel intensities from said difference image located in a first difference region corresponding to said first region surrounding said first track point with pixel intensities from said difference
10 image located in a second difference region corresponding to said second region surrounding each said potential subsequent track point to produce a weighting value;

15 b) comparing edge elements from said first region surrounding said first track point with those in said second region surrounding each said potential subsequent track point to determine a match of total edge
20 element intensity of pixels within said first region and said second region; and

c) assigning said match value for each said potential subsequent track point, said match value determined by weighting said match of total edge element
25 intensity of pixels within said first region and said second region with said weighting value.

30

35

This Sheet Has Been Replaced

