

Vision-based Pedestrian Detection: will Ants Help?

M. Bertozzi, A. Broggi, A. Fascioli, and P. Lombardi

Abstract— This work presents the vision-based system for detecting pedestrians in road environments implemented on the ARGO prototype vehicle developed by the University of Parma. The system is aimed at the localization of pedestrians in various poses, positions and clothing, and is not limited to moving people.

Initially, attentive vision techniques relying on the search for specific characteristics of pedestrians such as vertical symmetry and strong presence of edges, allow to select interesting regions likely to contain pedestrians. Then, such candidates areas are validated verifying the actual presence of pedestrians by means of a shape detection technique based on the application of autonomous agents.

Keywords— pedestrian detection, machine vision, autonomous agents, symmetry, ants.

I. INTRODUCTION

The detection of pedestrians is a mandatory requirement for future driving assistance systems, since having the capability of avoiding crashes with pedestrians is essential for effectively aiding the driver in urban environments.

This work presents the vision-based system for detecting pedestrians implemented on the ARGO vehicle. ARGO is an experimental autonomous vehicle developed by the University of Parma, equipped with a vision system, a processing engine, and automatic steering capabilities [1]. The main target of the ARGO Project is the development of an active safety system which can also act as an automatic pilot for a standard road vehicle.

Vision-based pedestrian detection in outdoor scenes is a challenging task even in the case of a stationary camera. In fact, pedestrians usually wear different clothes with various colors that, sometimes, are barely distinguishable from the background (this is particularly true when processing black and white images). Moreover, pedestrians can wear or carry items like hats, bags, umbrellas, and many others, which give a broad variability to their shape.

When the vision system is installed on-board of a moving vehicle additional problems must be faced, since the observer's ego-motion entails additional motion in the background and changes in the illumination conditions. In addition, since Pedestrian Detection is more likely to be of use in a urban environment, also the presence of a complex background (including buildings, moving or parked cars, cycles, road signs, signals, ...) must be taken into account.

Widely used approaches for addressing vision-based Pedestrian Detection are: the search of specific patterns or textures [2], stereo vision [3, 4], shape detection [5, 6, 7], motion detection [8, 9, 10], neural networks [11, 12]. The great part of the

This work was supported by the European Research Office of the U. S. Army under contract number N68171-01-M-5857

M. Bertozzi, A. Broggi, and A. Fascioli are with the Dip. di Ingegneria dell'Informazione, Università di Parma, ITALY. E-mail: {bertozzi,broggi,fascioli}@ce.unipr.it.

P. Lombardi is with the Dip. di Informatica e Sistemistica, Università di Pavia, ITALY. E-mail: lombardi@vision.unipv.it

research groups use a combination of two or more of these approaches [13, 14, 15]. Anyway, only a few of these systems have already proved their efficacy in applications for intelligent vehicles.

In this work the strong vertical symmetry of the human shape is exploited to determine specific regions of interest which are likely to contain pedestrians. Subsequently, an agent-based algorithm is used to verify the presence of pedestrians in those areas through an innovative application of a shape detection technique. This method allows the identification of pedestrians in various poses, positions and clothing, and is not limited to walking people.

This paper is organized as follows. Section 2 introduces the attentive algorithm aimed at the determination of the areas of interests, and shows its results. Section 3 describes the agent-based algorithm used for the validation of the candidates and presents the preliminary results of this shape detection technique. Section 4 ends the paper with some final remarks.

II. ATTENTIVE VISION

As a first processing step, attentive vision techniques are applied to concentrate the analysis on specific regions of interest only. In fact, the aim of the low-level part of the processing is the focusing on potential candidate areas to be further examined at a higher-level stage in the following steps.

The areas considered as candidate are rectangular bounding boxes which:

- have a size in pixels deriving from the knowledge of the intrinsic parameters of the vision system (angular aperture and resolution); in other words, once defined the size and distance of a pedestrian in the 3D world (e. g. $1.8\text{ m} \times 0.6\text{ m}$ at 20 m), simple perspective considerations give the size in pixel of its projection in the image (see figure 1.a);
- enclose a portion of the image which exhibits the low-level features that characterize the presence of a pedestrian, i. e. a strong vertical symmetry and a high density of vertical edges. These bounding boxes will be then checked against a human shape model, taking into account the contour of the object they enclose, in order to be validated.

The search for candidates would require an exhaustive search in the whole image (see figure 1.b). However, the knowledge of the system's extrinsic parameters, together with a flat scene assumption (see figure 1.c), is exploited to limit the analysis to a stripe of the image (hereinafter referred to as *search area*). The displacement of this stripe depends on the pedestrian's distance, while its height is related to the pedestrian's height (see figure 1.d). Since by definition a pedestrian is a human shaped road participant, the flat world assumption becomes an assumption on the road slope, which is anyway a loose hypothesis in a road environment, particularly in the area immediately ahead of the vehicle. Besides the obvious advantage of avoiding false detections in wrong areas, the processing of the search

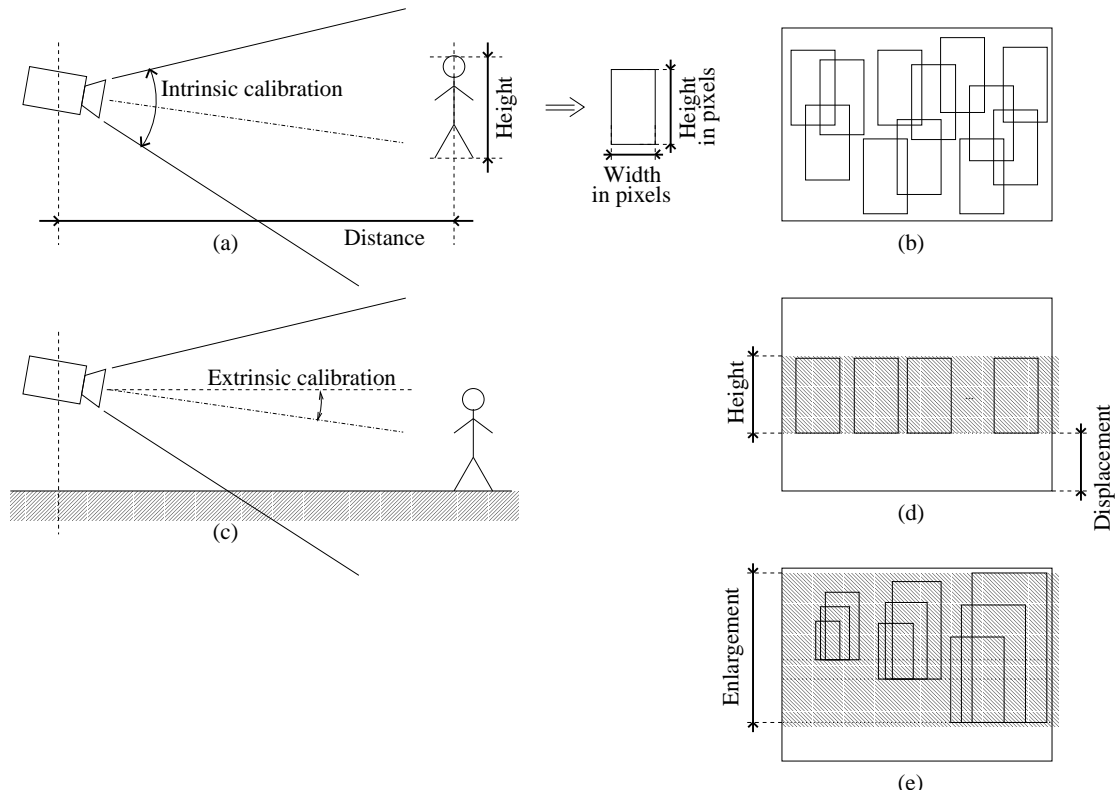


Fig. 1 (a) Computation of the bounding box size given the intrinsic parameters and the size and distance of a pedestrian; (b) exhaustive search for candidates in the whole image; (c) the search area can be limited to a stripe given the extrinsic parameters and a flat scene assumption; (d) the displacement and height of the stripe depend on the pedestrian distance and height, respectively; (e) the search area is enlarged to explore a range of distances and heights.

area only reduces the computational time. Indeed, the analysis cannot be limited to a fixed size and distance of the target and a given range for each parameter is in fact explored (e. g. $1.6 \div 2.0 \text{ m} \times 0.5 \div 0.7 \text{ m}$ at $10 \div 30 \text{ m}$). The introduction of these ranges generates two further degrees of freedom in the size and position of the bounding boxes. In other words, the search area is enlarged to accommodate all possible combinations of height, width, and distance (see figure 1.e).

The analysis proceeds in this way: the columns of the image are considered as possible symmetry axes for bounding boxes. For each symmetry axis different bounding boxes are evaluated scanning a specific range of distances from the camera (the distance determines the position of the bounding box base) and a reasonable range of heights and widths for a pedestrian (the corresponding bounding box size can be computed through the calibration).

However, not all the possible symmetry axes are considered: since edges are chosen as discriminant in most of the following analysis, a pre-attentive filter is applied, aimed at the selection of the areas with a high density of edges. In particular, for each axis the count of edge pixel is computed in a portion of the search area centered on the axis itself and as wide as the maximum bounding box width. Axes centered on regions which contain a number of edges lower than the average value are then dropped.

For each of the remaining axes the best candidate area is selected among the bounding boxes which share that symmetry

axis, while having different position (base) and size (height and width). Vertical symmetry has been chosen as a main distinctive feature for pedestrians. Symmetry edge maps, e. g. the Generalized Symmetry Transform (GST) [16], have already been proposed as methods to locate interest points in the image prior to any segmentation or extraction of context-dependent information. Unfortunately, these methods are generally computationally expensive. Alternatively, two different symmetry measures are performed: one on the gray-level values (G) and one on the gradient values, considering only edges with a vertical direction (E). The selection of the best bounding box is based on maximizing a linear combination of the two symmetry measures, masked by the density of edges in the box (D), as shown in the following equation: $S = (k1 \times G + k2 \times E) \times D$. The weights $k1$ and $k2$ were experimentally determined analyzing a large number of images. Figure 2 shows the original input image, the result of a clustering operation used to improve the detection of edges, a binary image containing the vertical edges, and a number of histograms representing the maximum (i) symmetry of gray-levels, (ii) symmetry of vertical edges, and (iii) density of vertical edges among the bounding boxes examined for each axis. The histogram in figure 2.g represents the linear combination of all the above. The histograms are actually computed only for the regions of the search area with a high density of edges, however in figure 2 they are completely displayed for a better understanding. It is evident that, using the density of vertical

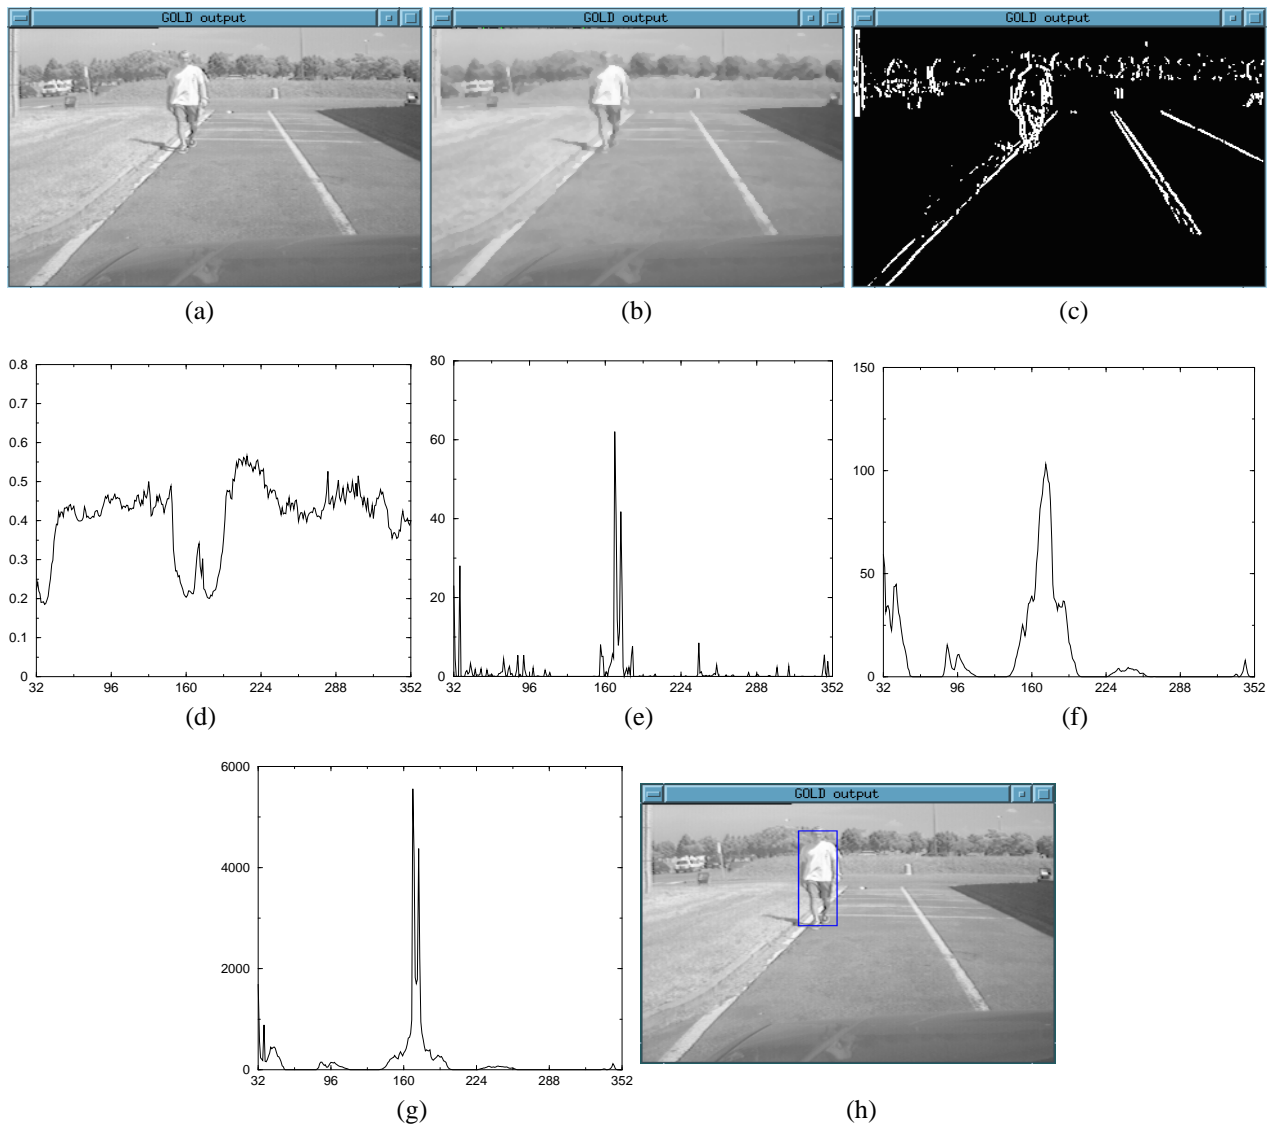


Fig. 2 Intermediate results leading to the localization of bounding boxes: (a) original image; (b) clustered image; (c) vertical edges; (d) histogram representing grey level symmetries; (e) histogram representing vertical edges symmetries; (f) histogram representing vertical edges density; (g) histogram representing the overall symmetry S for the best bounding box for each column; (h) the resulting bounding box.

edges as a mask, interesting areas present high values for both the symmetry of gray-levels and symmetry of vertical edges. The resulting histogram is therefore thresholded and its over-threshold peaks are selected as representing candidate bounding boxes.

An adjustment of the bounding boxes' size is yet needed. In fact, when comparing the gray-level symmetry of different bounding boxes centered on the same axis, larger boxes tend to overcome smaller ones since pedestrians are generally surrounded by homogeneous areas such as concrete underneath or the sky above (this is true for other objects, too). Therefore, the bounding box which presents the maximum symmetry tends to be larger than the object it contains because it includes uniform regions. For this reason, given a peak of the overall histogram representing a selected symmetry axis, the exact height and width of the best bounding box are actually taken as those

possessed by the box which maximizes a new function among the ones having the same axis. This function is computed as the product of the symmetry of vertical edges (E) and density of vertical edges (D) only. Figure 3 summarizes the overall candidate generation process.

The result of this low-level processing is a list of candidate bounding boxes which is fed to the following stage, whose task is their validation as pedestrians, based on higher-level characteristics.

A. Results of low-level attentive vision

The algorithm has been tested on a large number of images acquired in different situations ranging from simple uncluttered scenes to complex scenarios. As an example, figure 4 shows the result of the selection of candidate bounding boxes in three different situations. In figure 4.a a correct detection of two pedes-




Fig. 3 (a) Selection of the best bounding box for each symmetry axis; (b) selection of the best symmetry axes; (c) selection of the best candidates for each selected axis by choosing the bounding box which maximizes the symmetry and density of vertical edges.

trians is displayed. Figure 4.b presents a complex scenario in which only the central pedestrian is detected; two other pedestrians are missed because (i) the first is confused with the background, and (ii) the second is only partially visible; moreover, the high symmetry of a tree has been detected. In figure 4.c the two crossing pedestrians have been localized, but other symmetrical areas are highlighted as well; the first is due to the symmetry of the area between two pedestrians, and the other is due to the presence of symmetrical road infrastructures between two trees.

Some general considerations can be drawn on the behavior of this candidate selection procedure. In situations in which pedestrians are sufficiently contrasted with respect to the background and completely visible (i. e. not hidden by other pedestrians or objects) the localization of pedestrians based on symmetry and edge density proves to be robust. Thanks to the use of vertical edges the width of the bounding boxes enclosing pedestrians is generally determined with a good precision. On the other hand, a lower accuracy is obtained for the localization of the top and bottom of the bounding box. A refinement of the bounding box height is under development.

Symmetrical objects other than pedestrians may happen to be detected as well. In order to get rid of such false positives a number of filters have been devised which rely on the analysis of the distribution of edges within the bounding box. These filters, which are still under evaluation, show promising results regarding the elimination of both artifacts (such as poles, road signs, buildings, and other road infrastructures) and symmetrical areas given by a uniform portion of the background between two foreground objects with similar lateral borders (see figure 4.c).

III. SHAPE DETECTION USING AUTONOMOUS AGENTS

This section describes the shape detection technique.

Different edges are selected and connected, where possible, in order to form a contour. Thanks to the way the contour is built, it will represent the shape of the pedestrian body. Matching techniques may be used in regions of the bounding box that lie in the correct position in relation to the formulated hypothesis of the pedestrian.

Essentially, the process consists in adapting a deformable coarse model to the bounding boxes. Thanks to its roughness

the model is sufficiently general and can be adapted to a variety of postures. Anyway, it is limited to standing pedestrians. The model adjustment is done through an evolutionary approach with a number of independent agents acting as edge trackers. The agents explore a feature map displaying the edges contained in a given bounding box and stochastically build hypotheses of a feasible contour of a human body. The idea is taken from the Ant Colony Optimization (ACO) metaheuristic devised to solve hard combinatorial optimization problems, originally inspired by the communication behavior of real ants [17]. The system proposed here is a transposition to image analysis of one of the first ACO algorithms, the *AS-cycle*.

In nature, when ants look for food, they communicate the path and the outcome of their exploration to other ants by marking their path with a pheromone trail, its intensity depending on the distance of the food from the nest, and on its quality and quantity. Other ants are attracted by strong pheromone trails, thus the path to an abundant food source close to the nest is marked again and again until it becomes more frequented and even more attractive.

This concept can be applied to the analysis of an image by creating a colony of artificial ants that looks for an optimal combination of edge pixels that maximizes the coherency of their position according to a given model. Each ant in turn traces a solution in a solution space made up of all the possible paths connecting two pixels in a matrix. The decisional basis for each step of an ant is provided by two factors: one is a local heuristic η_i that quantifies the attractiveness of pixel i for its intrinsic characteristics; the second is the information on that pixel made available by previous attempts of other ants, in the form of a quantity of pheromone τ_i .

Artificial ants explore a world which is a matrix of pixels derived by the resampling of the edge map of the bounding box under analysis. In our experiments, the normalized world-matrix is sized 20×45 pixels. Each pixel i is initialized with a binary value: 1 if it contains an edge, 0 if not. This value represents its intrinsic attractiveness η_i and is the basis for the heuristic research. All the pheromone τ_i is initialized at 0. The world-matrix is visited by M ants in parallel, and the process is repeated for C cycles. At the end of each cycle, new pheromone is deposited on the trails pursued by the ants, and some of that




Fig. 4 Result of low-level processing in different situations: (a) a correct detection of two pedestrians (b) a complex scenario in which only the central pedestrian is detected, and the high symmetry of a tree has been detected as well; (c) two crossing pedestrians have been localized, but other symmetrical areas are highlighted as well.

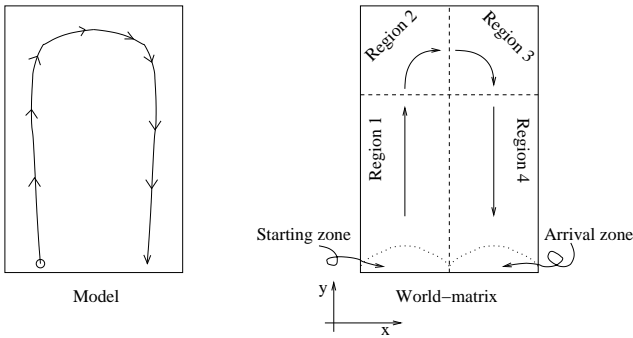


Fig. 5 Artificial ants move through the world-matrix starting from the left half of the lower border, and moving through regions 1, 2, 3 and 4 until they reach the arrival line.

accumulated evaporates. In this way, solutions built several cycles before, progressively loose their importance. On the other hand, pheromone on pixels that compose the path of frequently selected solutions grows, and eventually this information surpasses that given by the heuristic.

A crucial point to be understood is that artificial ants do not need to reach an optimal solution in the edge connection problem. Often, no real optimal solution exists even to a human inspector. The colony needs only to find a sufficiently valuable path that permits to continue the recognition, free of noisy edge pixels.

The ant system develops from a very elastic and deformable coarse model of a human body. The model is encoded in the progression rules that guide the ants through the solution space. The rules effectively restrict the whole space of possible solutions to a subspace that includes the searched shape. The system will then provide attempts to find a feasible path in this subspace, and each attempt will be evaluated by a confidence function as it is detailed in the following.

All ants start from the left half of the lower side of the world-matrix. The world is divided into four regions, as detailed in figure 5. In each region ants proceed of one step forward in the direction of one of the axis (y in regions 1 and 4, and x in regions 2 and 3), and can choose among a set of s pixels lying on the line or column in front of them. Additionally, in regions 2 and 3 ants

have the option of moving vertically, thus they can follow very steep edges as well as very flat ones.

The starting point of each ant is chosen randomly among the edge pixels lying in the starting region, i.e. the left half of the lower border. If no edge pixel is present, the starting point is set on a random point belonging to the region. The choice of starting from edge pixels does not pose a hard restriction to the exploration of the solution space: the bounding box usually comprises edges in the lowest line owing to the mechanism that determines its dimensions based on the edge density. Most of the times, the edges appearing on the lower line of a well-centered bounding box correspond to the feet of the pedestrian. Each ant stops its journey when it reaches the right half of the lower border of the world-matrix.

An ant is an independent pixel-sized agent; it has a local exploratory capability, limited to the set of pixels belonging to the scanning region N as described above, and of those lying on the following line as well. Figure 6 illustrates the situation for the scanning sets for each region of the world-matrix. Each pixel under consideration is associated to a quality measure that takes into account terms pertaining to both the feature map of the edges, and the pheromone deposited by previous ants. The quality of pixel j is expressed as $q_j = \alpha\tau_j + (1 - \alpha)\eta_j$ where η_j represents the binary heuristic information, τ_j is the quantity of pheromone accumulated at position j , and α is a parameter which determines the relative influence of the pheromone trail and the heuristic information.

Each ant always moves into one of the pixels of the nearest scanning line (line A, namely the shaded sets in figure 6), but the probability of transition combines the quality of each pixel in line A with that of a corresponding pixel in line B as indicated by the arrows in figure 6. Defining with l a pixel in line B corresponding to a pixel j in line A, the probability that ant k moves from position i to position j belonging to its feasible neighborhood N_i^k at step t is

$$p_{ij} = \frac{\frac{1}{d_j} \times [(1 - v) \times q_j + v \times q_l]}{\sum_{(j,l) \in N_i^k} \frac{1}{d_j} \times [(1 - v) \times q_j + v \times q_l]} \quad (1)$$

where v is a parameter in a range $[0, 1]$ indicating the ants




Fig. 6 Artificial ants move to one pixel of the shaded set (named *line A*) by calculating the quality of each pixel of line A and of the white region (named *line B*). The figures illustrate the set of pixels evaluated by ants when they cross region 1 (a), region 2 (b), region 3 (c), and region 4 (d).

field of view; for $v = 1$ the ant sees only line A pixels, for $v = 0$ the ant sees only line B pixels, while for intermediate values the ant focus of attention varies in between line A e B. d_j is the displacement of pixel j with respect to the central pixel of line A. The $1/d_j$ penalty favors straight trails in comparison with frequent small alternative leaps to the left and the right.

The system provides two different kinds of agents: purely stochastic ants and semi-deterministic ants. Both kinds choose their move with a uniformly distributed random rule, but the range of choice is different: purely stochastic ants have all the feasible neighborhood N illustrated in figure 6 at their disposal, while semi-deterministic ants choose only between the two pixels that have the highest p_{ij} . Both kinds of ants perform well on synthetic images; however, stochastic ants explore more widely the solution space but converge more slowly to a final solution than the semi-deterministic ants do. On the other hand, semi-deterministic ants follow well connected edges, but sometimes fail to find the best solution subspace in very irregular real images.

Once every ant has completed its tour, pheromone trails are updated through evaporation and reinforcement according to the following equation:

$$\tau_i(c+1) = (1 - \rho) \times \tau_i(c) + \rho \times \left(\sum_{k=1}^M \Delta\tau_i^k + \Delta\tau_i^d \right) \quad (2)$$

where ρ is the evaporation coefficient (ranging from 0 to 1), $\tau_i(c)$ is the quantity of pheromone present on pixel i at cycle c . Pheromone update $\Delta\tau$ is made up of two contributions. The first one is given by the sum of the pheromone deposited by each ant at the end of its tour. The second one is credited to the best trail according to an elitist strategy.

All ants are ranked according to the following rule: an ant obtains a high rank if it takes a long tour that passes through many edges or a low rank if it visits many pixels that are not edges. This rule is functional to the search of a good solution as it encourages ants to take the shortest path between two zones of connected pixels and does not pose any request on the total length of the trail.

The procedure described above is repeated for a number of cycles; experiments show that with 10 ants, 2 cycles are sufficient for a stable and reliable solution.

Finally, the output is the path of the ant of the highest rank in the last cycle.

A. Preliminary results of shape detection

Preliminary experiments were done mainly on synthetic images, like the one shown in figure 7.a. The performance of both purely stochastic and semi-deterministic ants were compared to deterministic edge trackers that follow the same movement rules as artificial ants, but proceed always on one deterministically chosen edge pixel. These trackers often failed to find a correct solution on both synthetic and natural images as edges often present bifurcations for which no deterministic decision rule could be conceived. The random decision rule of artificial ants, together with a high number of attempts, proved to perform better than deterministic trackers.

Only two cycles were sufficient to reach a feasible solution on synthetic images. Ten ants were running in each cycle. The evaporation coefficient ρ was set to 0 so that the second cycle would take into account the full information provided by the first. The second cycle proved useful in finding a stable solution, in the sense that the solution subspace detected by the system was kept the same over multiple attempts on the same image, in spite of the random nature of the algorithm.

Figure 7.d shows the path drawn by the best purely stochastic ant at the end of the second cycle on the normalized edge matrix of the synthetic pedestrian of figure 7.a. The ant correctly delineates the trunk and legs shape of the human shape, and, in this case, it cuts out the head. A future step in the recognition process will start from this shape and try to identify the head or possibly other parts of the body in their correct positions.

The result of figure 7.d was obtained using parameters $\alpha = 0.5$, $v = 0.2$, $\rho = 0.0$, $Q_a = 0.5$, $Q_d = 1.0$, $C = 2$, $M = 10$. Figure 7.b illustrates an image of the world-matrix on which artificial ants moved during the second cycle, where brighter pixels have a higher quality q_i . Figure 7.c shows the pheromone matrix τ_i . Again, the brighter the pixel, the higher its pheromone quantity.

IV. DISCUSSION

This work presents a vision-based system for detecting pedestrians in road environments.

Initially, low-level processing techniques are used to focus on few areas of interest which potentially contain pedestrians. Then, a subsequent higher-level processing is used to further analyze these areas of the image by means of autonomous agents: an ant-based matching with a human shape model is used for validating the presence of pedestrians.




Fig. 7 (a) Example of a synthetic image used for experiments with artificial ants; (b) map of the world-matrix on which artificial ants move; (c) map of the pheromone trails deposited by ants after two cycles have been completed; (d) example of the result obtained by the best purely stochastic ant after two cycles.

This algorithm suits a medium distance search area. In fact, large bounding boxes may contain a too detailed shape, showing many disturbing small details that would certainly make their detection extremely difficult. In other words, the presence of texture (not only caused by different clothing) and the many different human postures that must be taken into account, would make the detection hard. On the other hand, very small bounding boxes enclosing far away pedestrians feature a very low information content. In these situations it is easy to obtain false positives, since many road participants (other than pedestrians), other objects, and even road infrastructures may present morphological characteristics similar to a human shape. It is therefore imperative to define a range of reasonable-sized bounding boxes in which the detection may lead to sufficiently accurate detections. In this work the considered size is: 12 x 28 pixel for the smallest bounding box, and 42 x 100 pixel for the largest one. This choice removes the small errors caused by false detections of small objects, as well as inaccurate detection (or even missed detections) of large pedestrians. Indeed this choice leads to a limited detection area in front of the vehicle. The system was tested on the images acquired by the vision system installed on-board of the ARGO experimental vehicle. With the current setup the search area ranges from 10 to 30 m.

The candidate selection procedure based on vertical symmetry and edge density proved to be a robust technique for focusing the attention on interesting regions. From the first preliminary results, the ant-based processing appears to be a promising method for detecting the contour of a human shape. To extend the detection to a larger set of pedestrian postures, other models are currently under development.

ACKNOWLEDGMENTS

The authors gratefully thank Dr. C. Adams and Dr. M. Del Rose from U. S. Army TACOM and Dr. S. Sampath from USARDSG-UK for their support in the research.

REFERENCES

[1] A. Broggi, M. Bertozzi, G. Conte, and A. Fascioli, "ARGO Prototype Vehicle," in *Intelligent Vehicle Technologies* (L. Vlacic, F. Harashima,

- and M. Parent, eds.), ch. 14, pp. 445–493, London, UK: Butterworth-Heinemann, June 2001. ISBN 0750650931.
- [2] C. Curio, J. Edelbrunner, T. Kalinke, C. Tzomakas, and W. von Seelen, "Walking Pedestrian Recognition," in *Procs. IEEE Intl. Conf. on Intelligent Transportation Systems '99*, (Tokyo, Japan), pp. 292–297, Oct. 1999.
- [3] L. Zhao and C. Thorpe, "Stereo- and Neural Network-based Pedestrian Detection," in *Procs. IEEE Intl. Conf. on Intelligent Transportation Systems '99*, (Tokyo, Japan), pp. 298–303, Oct. 1999.
- [4] D. Beymer and K. Konolige, "Real-time Tracking of Multiple People using Continuous Detection," in *Procs. Intl. Conf. on Computer Vision*, 1999.
- [5] D. M. Gavrila, "Pedestrian Detection from a Moving Vehicle," in *Procs. of European Conf. on Computer Vision*, vol. 2, pp. 37–49, June–July 2000.
- [6] D. M. Gavrila, "Sensor-based Pedestrian Protection," *IEEE Intelligent Systems*, vol. 16, pp. 77–81, Nov.–Dec. 2001.
- [7] C. Papageorgiou, T. Evgeniou, and T. Poggio, "A Trainable Pedestrian Detection System," in *Procs. IEEE Intelligent Vehicles Symposium '98*, (Stuttgart, Germany), pp. 241–246, Oct. 1998.
- [8] R. Polana and R. C. Nelson, "Detection and Recognition of Periodic, Non-rigid Motion," *International Journal of Computer Vision*, vol. 23, pp. 261–282, June–July 1997.
- [9] S. J. McKenna and S. Gong, "Non-intrusive Person Authentication for Access Control by Visual Tracking and Face Recognition," *Lecture Notes in Computer Science*, vol. 1206, pp. 177–184, Mar. 1997.
- [10] R. Cutler and L. S. Davis, "Robust real-time periodic motion detection, analysis and applications," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 22, pp. 781–796, Aug. 2000.
- [11] C. Wöhler, J. K. Aulaf, T. Pörtner, and U. Franke, "A Time Delay Neural Network Algorithm for Real-time Pedestrian Detection," in *Procs. IEEE Intelligent Vehicles Symp. '98*, (Germany), pp. 247–251, Oct. 1998.
- [12] C. Wöhler, U. Kreßel, and J. K. Anlauf, "Pedestrian Recognition by Classification of Image Sequences – Global Approaches vs. Local Spatio-Temporal Processing," in *Procs. IEEE Intl. Conf. on Pattern Recognition*, (Barcelona, Spain), Sept. 2000.
- [13] C. Curio, J. Edelbrunner, T. Kalinke, C. Tzomakas, and W. von Seelen, "Walking Pedestrian Recognition," *IEEE Trans. on Intelligent Transportation Systems*, vol. 1, pp. 155–163, Sept. 2000.
- [14] V. Philomin, R. Duraiswami, and L. Davis, "Pedestrian Tracking from a Moving Vehicle," in *Procs. IEEE Intelligent Vehicles Symposium 2000*, (Detroit, USA), pp. 350–355, Oct. 2000.
- [15] L. Zhao and C. Thorpe, "Stereo and neural network-based pedestrian detection," *IEEE Trans. on Intelligent Transportation Systems*, vol. 1, pp. 148–154, Sept. 2000.
- [16] D. Reissfeld, H. Wolfson, and Y. Yeshurun, "Context Free Attentional Operators: the Generalized Symmetry Transform," *Intl. Journal of Computer Vision, Special Issue on Qualitative Vision*, vol. 14, pp. 119–130, 1994.
- [17] M. Dorigo and G. Di Caro, "The ant colony optimization meta-heuristic," in *New Ideas in Optimization* (D. Corne, M. Dorigo, and F. Glover, eds.), pp. 11–32, London, UK: McGraw-Hill, 1999.
- [18] M. Dorigo and L. M. Gambardella, "Ant Colony System: A Cooperative Learning Approach to the Traveling Salesman Problem," *IEEE Tran. on Evolutionary Computation*, vol. 1, pp. 53–66, Apr. 1997.