Polynomial Curb Detection Based on Dense Stereovision for Driving Assistance

F. Oniga and S. Nedevschi

Abstract—**A real-time algorithm for curb detection in traffic scenes, based on dense stereovision, is proposed. Curbs are modeled as cubic polynomial curves. 3D points from stereovision are transformed into a Digital Elevation Map (DEM), in order to have a compact representation of the 3D space. Curb points are detected as the cells of the DEM that present a specific height variation. Only curb points that are temporally persistent and non-occluded are considered. Relevant cubic polynomials are computed from the set of curb points by a RANdom SAmple Consensus (RANSAC) approach. For each relevant polynomial, the curb patch is extracted by analyzing the DEM along the polynomial curve. Finally, the vertical location and height of each curb are computed based on the local elevation data.**

I. INTRODUCTION

CURBS are important delimiters of the drivable area: they must be taken into account for navigation, together with relevant obstacles. Urban lanes often have curbs on sides (without lane markings), therefore curbs can be used as features for lane detection algorithms. Curbs frequently separate the road from the sidewalk: they can be used for focusing attention in pedestrian detection algorithms.

Because curbs have heights much smaller than obstacles, dedicated approaches must be used for curb detection. The sensors used for curb detection are based on stereovision [1]-[4], lidar [5]-[8], time-of-flight camera [9] or a laser line stripper [10] combined with a camera [11].

One of the earliest approaches for curb detection is [1], based on stereovision. Candidate curbs are extracted as clusters of at least three parallel lines in the image space, by applying the Hough transform. Two planes are computed on the 3D data of the surrounding regions, and the height of the curb is estimated.

The approach presented in [2] is based on both range and photometry data. Image edge points are detected and the Hough accumulator is built. The scalar product of the elevation and brightness gradients of each edge point is used for voting in the Hough accumulator. One dominant straight curb (per scene) is extracted as the line with the maximum score in the accumulator. This method is improved later in [3]. The Hough accumulator is built from edge segments close to the ground plane, and the vote is weighted with the curvature index (based on the principal curvatures). Peaks of the Hough transform are considered candidate curb lines, which are further refined: regression is performed to the 3D points of the candidate line, and the ends of the curb are searched.

The 3D points, from dense stereo, are transformed into a Digital Elevation Map (DEM) in [4]. Edges are detected on the DEM. Candidate curb points are those edges that are persistent along successive frames. The Hough accumulator is built from the curb points. Curb segments are located along linear curbs (Hough peaks): the largest segment with specific height variation (orthogonal to the curb) is selected for each curb line. A scheme is proposed to extend iteratively each curb segment with additional segments, in order to provide a better representation for curved curbs (as chain of segments).

In [5], road curbs are extracted using a planar two-dimensional laser measurement system. Authors propose a robust filtering and segmentation of the laser range data using an Extended Kalman Filtering. Lines are then fitted to the segmented data by minimizing the sum of squared errors. The pair of lines oriented similar to the ego car that fits best the width of the road is selected as the pair of curbs delimiting the road. The method fails if the road does not have curbs on both sides or if one of the curbs is occluded by obstacles. In [6] the detection is improved by using a mixture of two-dimensional scanning laser radar (LADAR) and a charge-coupled device (CCD) monocular camera.

The method presented in [7] uses a four-layer laser radar sensor for range computation. A multi-stage curb detection algorithm is proposed: the initial curbs (beginning segments of curbs, having an orientation similar to the ego vehicle) are extracted from the range data; the initial curbs are extended using an extended Kalman filtering technique; curbs are tracked between frames in order to stabilize the results.

Curbs are detected based on a derivative approach, using a 64-beam lidar sensor [8] mounted on top of the car. The authors avoid the sparseness of the 3D Cartesian representation by processing the data in a sensor-centric Cartesian system (as a grid). Then, curbs are detected based on a filtering process: the grid edges are considered measurements and the curbs position estimation is performed based on the local road network available with the offline map. Ridges in the edge strength map, having the appropriate orientation and placed in the predicted locations, are considered curbs.

Manuscript received April 19 2010. This work was supported by CNCSIS – UEFISCSU, project number PNII – IDEI 1522/2008.

F. Oniga and S. Nedevschi are with the Computer Science Department, Technical University of Cluj-Napoca, Romania, email: florin.oniga, sergiu.nedevschi@cs.utcluj.ro

A Time-Of-Flight camera is used in [9] for curb detection. The authors propose a modified variant of the RANSAC algorithm called Connected Component-RANSAC (CC-RANSAC). The improvement is that the fitting score for a random sample is the largest number of inliers that fit the model and are connected to each other. CC-RANSAC is used to compute relevant planes in the scene and their sets of connected inliers (planar patches). Curbs are located at the boundary between the planar patches.

Curb detection was performed in [10] using a laser line striper (composed of a camera and a laser), and the system was integrated on a bus. The acquisition system is mounted in the front bumper, oriented laterally towards the curb. The bus-curb distance is detected for each frame. This system is simple and efficient, but it has a limited space of interest and field of view. Later in [11] the laser striper is augmented with a video camera, for extending the curb detection in front of the ego vehicle. Using information about the dynamics of the vehicle, past curb measurements are mapped into the current reference system. In this way, the shape of the curb is recovered, up to the current measurement. The curb is then extended in front of the ego vehicle, using the images from the forward-looking camera.

In this paper, we propose an algorithm based on stereovision that uses a DEM representation of the 3D data. We propose a polynomial curve model for curb, which is more realistic than the line or poly-line models used in existing stereovision based solutions. The steps necessary to extract the polynomial curbs in real-time are presented.

An overview of the proposed algorithm is presented in section II. The polynomial curve model used for curbs is presented in section III. Sections IV and V present the main steps of our algorithm. Results and future work are discussed in section VI.

II. Overview of the Algorithm

The curb detection algorithm (see Fig. 1) proposed in this paper takes as input 3D points computed with a dense stereovision system.

A DEM is generated from the set of 3D points and all the processing steps are performed on the DEM. The DEM is suitable for real-time processing: it provides explicit connectivity between adjacent 3D locations, and it reduces the processing space (only one relevant height is stored for multiple 3D points).

Edge detection is applied on the DEM to detect candidate curb points. A temporal/occlusion filtering is employed to filter out false/background curb points.

The map of curb points is processed using a RANSAC approach to extract the cubic polynomial that most likely contains a curb. Then, the curb extremities and additional curb features are searched for in the DEM. This will localize the curb patch in the bird-eye view (depth and lateral offset). The final step is the vertical localization of the curb patch based on the features extracted at the previous step.

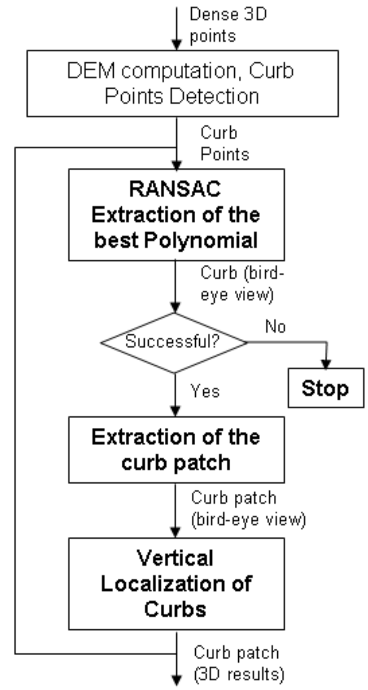

Fig. 1. Overview of the proposed algorithm

III. Cubic Polynomial Curb Model

Most of the stereovision-based methods for curb detection rely on a linear curb model [1]-[3] or a poly-line model [4] (chain of line segments). Curved curbs often appear in urban scenarios, and, even if they are estimated as poly-lines, the localization is less accurate (see Fig. 2).

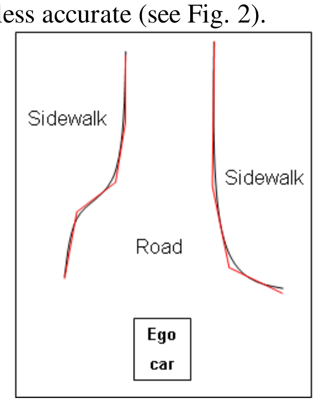

Fig. 2. Modeling of curved curbs (black) by poly-lines (red) is less accurate, especially near the poly-line vertexes.

A better model for curved curbs is a cubic polynomial. This model allows curbs to have curvatures and variation of curvatures, and it is in accordance to the widely accepted clothoidal model for the road (lane) boundaries.

The algebraic form of a cubic polynomial is:

$$y = ax^3 + bx^2 + cx + d\ . \qquad (1)$$

Fitting the cubic polynomial to n points, (x_i, y_i), $i=1..n$, is done by solving the system of n equations and 4 unknowns a,

b, c, and d :

$$\left\{ ax_i^3 + bx_i^2 + cx_i + d = y_i \,, \quad i = 1..n \right. . \tag{2}$$

The system can be written as a matrix equation:

$$AX = B \,, \tag{3}$$

$$A = \begin{bmatrix} x_1^3 & x_1^2 & x_1 & 1 \\ x_2^3 & x_2^2 & x_2 & 1 \\ \vdots & \vdots & \vdots & \vdots \\ x_n^3 & x_n^2 & x_n & 1 \end{bmatrix}, X = \begin{bmatrix} a \\ b \\ c \\ d \end{bmatrix}, B = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix}. \tag{4}$$

When the polynomial is fitted to a set of four points, the coefficients are computed by solving (3) (Cramer's rule, Gaussian elimination, etc). If the number of points is above four then the system is overdetermined and must be solved in a least square fashion. The solution that minimizes the sum of quadratic errors can be computed by solving the normal equation system:

$$(A^T A) X = A^T B \,. \tag{5}$$

Multiplying A^T by A during runtime is computationally intensive. It involves n multiplications and n-1 additions for each element of the output matrix. Considering the particular form of the matrix A, by performing explicitly the matrix multiplications, the system becomes:

$$\begin{bmatrix} \sum_{i=1}^{n} x_i^6 & \sum_{i=1}^{n} x_i^5 & \sum_{i=1}^{n} x_i^4 & \sum_{i=1}^{n} x_i^3 \\ \sum_{i=1}^{n} x_i^5 & \sum_{i=1}^{n} x_i^4 & \sum_{i=1}^{n} x_i^3 & \sum_{i=1}^{n} x_i^2 \\ \sum_{i=1}^{n} x_i^4 & \sum_{i=1}^{n} x_i^3 & \sum_{i=1}^{n} x_i^2 & \sum_{i=1}^{n} x_i \\ \sum_{i=1}^{n} x_i^3 & \sum_{i=1}^{n} x_i^2 & \sum_{i=1}^{n} x_i & n \end{bmatrix} \begin{bmatrix} a \\ b \\ c \\ d \end{bmatrix} = \begin{bmatrix} \sum_{i=1}^{n} x_i^3 y_i \\ \sum_{i=1}^{n} x_i^2 y_i \\ \sum_{i=1}^{n} x_i y_i \\ \sum_{i=1}^{n} y_i \end{bmatrix} . \tag{6}$$

This representation has a lower computational cost, because of two reasons:

- Powers of x_i can be computed apriori once for each point, therefore reducing the cost per matrix element to n-1 additions.
- Of the sixteen elements of the matrix $A^T A$ only seven are distinct.

The numerical stability of the polynomial is influenced by the set of points. Removal of the mean and variance normalization can be employed to improve stability. If the set of points is quasi-vertical (similar values of the x coordinate) then the polynomial cannot be computed (undetermined system). This can be avoided by rotating the reference frame by 90 degrees, or, in other words, by expressing x as a cubic polynomial in y. Therefore, for a given set of points two polynomials are computed, one for y=$P_1(x)$ and the other one for x=$P_2(y)$, and both are analyzed to select the optimal one (this will be explained in section V.A).

IV. Detection of Candidate Curb points

Building the DEM, detection of curb points and the temporal filtering are similar to the method described in [4]. For each DEM cell, the elevation of the highest point (within the cell) is stored, or the cell is marked empty (see Fig. 3). A DEM cell covers a 10x10 cm patch in the horizontal plane. Median filtering (5x5 window) is applied to reduce noise. Singular empty cells are filled with the neighbors' elevation, in order to improve the connectivity in the DEM.

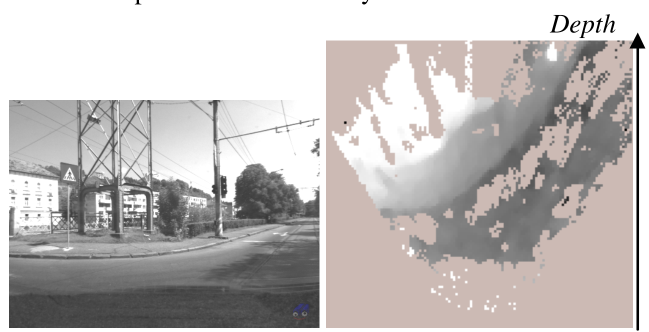

a. b.

Fig. 3. a. A turn left scenario in the urban environment, b. The DEM covering an area of 13x12 meters (bird-eye view): empty cells are marked with color, while valid cells store the elevation.

Edge detection (Canny) is applied to the valid cells of the DEM and only edges having the magnitude of the gradient (computed with a 3x3 Sobel filter) in an interval specific to curbs are considered candidate curb points (5..30 cm). Temporal filtering is performed by compensating the ego motion along successive frames and only persistent curb points are allowed for further processing (see Fig. 4.a).

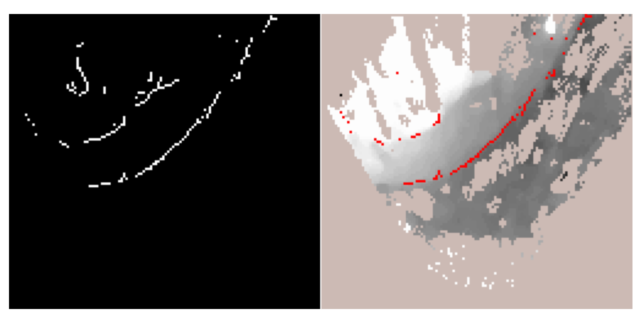

a. b.

Fig. 4. a. Candidate curb points with a temporal persistence of at least four frames, b. Persistent curb points that are not occluded by other curb points are drawn with red on the DEM.

We propose an additional filtering condition to the temporal filtering: occlusions removal. This condition is necessary in order to reduce the number of samples needed in the RANSAC polynomial fitting. Considering the geometry of the stereo system, a curb point C is most likely false if it is occluded by another curb point. The line segment between C and the stereo system is computed. C is considered occluded if the line segment intersects at least one curb point that is not part of the chain of curb points containing C. Only relevant curb points remain after filtering (see Fig. 4.b).

V. POLYNOMIAL CURB DETECTION

A. RANSAC Extraction of the Best Polynomial

The RANSAC approach requires the computation of the polynomial model for a number of K random samples from the data set, and chooses the model with the highest fitting score. The data set consists of the candidate curb points. A random sample is made of s=4 points, the minimum size required to compute the parameters of the cubic polynomial. The percentage w of curb inliers from the data set must be estimated. Considering that the data set contains only persistent and non-occluded points, a conservative value for w is 40% (although a normal value is around 70-80%). The number of samples required [12], assuming a success probability p=0.9999, is 355 samples.

The standard RANSAC approach involves selecting a number of K random samples of four candidate curb points. Two cubic polynomials are estimated for each sample ($y=P_1(x)$ and the other one $x=P_2(y)$), by solving the system of equations (3). The fitting score (the size of the consensus set / number of inliers for each polynomial) is computed for each estimated polynomial, and the polynomial having the best score S_B is selected. The best polynomial is validated if S_B>t (=10 inliers). Then, it is recomputed using a least square fit on its consensus set (see Fig. 6.a).

We enhanced the standard RANSAC approach with two improvements: the detection of bad samples in an early phase (thus reducing the computational complexity), and a different way of computing the fitting score.

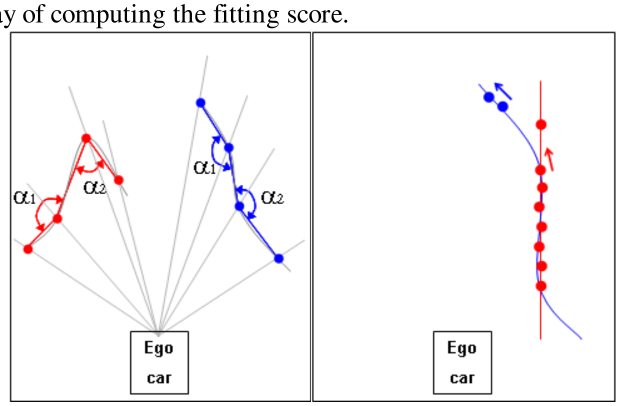

a. b.

Fig. 5. a. Two samples of four points, along with their lines of sight. The red sample has a sharp angle (α_2) and is not processed further, while the blue sample has only obtuse angles and is analyzed further. b. Noisy curb points (blue) can cause a false curvature for a straight curb (red points).

For each sample, four random points are selected. Some samples will have an undesired configuration of points, causing the polynomial curb to have high curvatures. These samples are less likely to represent a curb, and they can be rejected without the additional cost of computing the polynomial and its fitting score. More than half of the samples should be rejected on average scenarios. The criterion for rejection is estimated with the following steps (see Fig. 5.a):

1. The four points of the sample are ordered based on their lines of sight to the stereo system, counterclockwise.
2. Line segments are computed between adjacent points in the ordered set.
3. The angles between the associated segments are computed for each of the two inner points from the ordered set. If at least one of the angles is sharp, then the sample is considered bad.

The standard RANSAC approach computes the fitting score as the number of data points that verify the model (residuals below a threshold). This can cause an incorrect fitting when a straight curb is present together with some false curb points (see Fig. 5.b). For such a scenario, the cubic and quadratic terms of the best polynomial will allow it to "steer" towards outliers while still maintaining a good fit for the straight curb patch. The desired polynomial, that fits only the curb patch, has a smaller consensus set. To overcome this, we consider the consensus set C to be the largest sequence of inliers that have similar local orientations of the polynomial curve. The local orientation is estimated for each inlier, based on the first derivative.

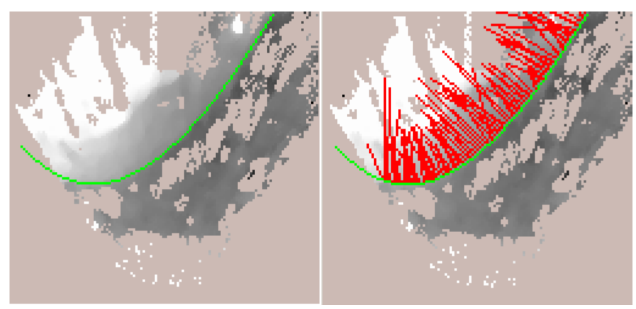

a. b.

Fig. 6. a. The best polynomial curve (green) detected on the DEM. b. The gradient is drawn for each point of the curve (orientation and magnitude).

B. Extraction of the Curb Patch

If a polynomial is computed successfully, then the extremities of the curb must be located. Additional features are computed from the DEM for each curb point: DEM gradient (magnitude) and the elevation of the roadside of the curb.

One way to detect the extremities of the curb on the polynomial curve is to find the extremes of the consensus set. However, we must take into account that the consensus set contains only curb points persistent along several frames. Due to lack of texture in previous frames, it is possible to have additional curb points on the current DEM that are not included in the consensus set.

The current DEM is reanalyzed along the polynomial curve in order to find the curb patch:

1. The DEM gradient is computed for each point of the curve (see Fig. 6.b).
2. Curve points having the magnitude in the curb expected interval [5..30] cm are considered potential curb points (see Fig. 7.a). Small gaps (1-2 points) of non-curb points are filled and small intervals of curb points are discarded. This is a compensatory measure for the lack of data and/or false elevations from dense stereo.
3. Intervals of consecutive curb points are extracted.
4. Adjacent intervals separated by a gap smaller than T (1 meter) are merged.
5. The largest interval is considered the curb patch (see Fig. 7.b).

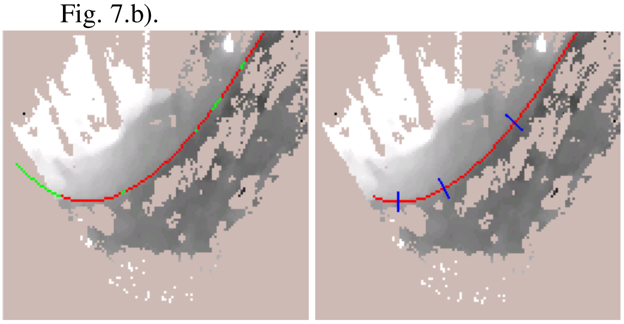

a. b.

Fig. 7. a. Potential curb points are marked (red) along the polynomial curve. b. The curb patch (red), and example of neighborhoods (blue) used to extract the roadside elevation for each curb point.

In order to locate the curb vertically, the roadside elevation is computed for each point of the curb patch. The road is the lower side of the curb. Thus, computing the road elevation for a curb point involves finding the minimum valid elevation inside a neighborhood around the point. A good choice for this neighborhood is a segment, orthogonal to the curb and centered on the curb point. A size of 5-7 DEM cells is acceptable, since it will contain several nearby road cells (see Fig. 7.b).

C. *Vertical Localization of Curbs*

Two features must be computed in order to describe completely the curb in the 3D reference system: the height of the curb and its vertical profile.

The height of the curb is computed as the mean value of the magnitude of the DEM gradients along the curb. A mean filter robust to outliers and to noisy measurements (see Fig. 8, top) is the Alpha-Trimmed mean filter. Gradient magnitudes are ordered, top α% and bottom α% (α=10) are discarded, and the mean of the remaining values is computed.

The vertical profile is estimated from the roadside elevations associated with the curb. This profile has both vertical slope and curvature for uphill/downhill roads. In order to model these features, a quadratic polynomial is fitted to the set of elevations, using a RANSAC approach (see Fig. 8, bottom). Then, the roadside elevation of each curb point is re-computed with the quadratic polynomial. This ensures that noise is filtered out, and that the roadside elevation is computed even for those curb points with no roadside elevation available in the DEM.

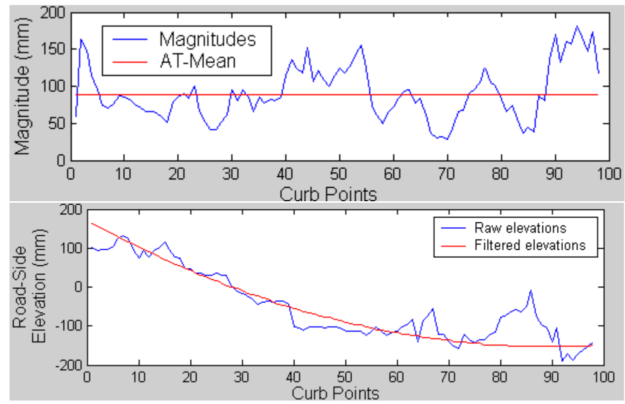

Fig. 8. Top: The noisy gradient magnitudes (blue) and the alpha-trimmed mean value (red) equal to 86 mm (the real curb height is around 90 mm). Bottom: The raw roadside elevations (blue) along the curb and the vertical profile of the curb (red, filtered elevations).

Now the curb patch is completely located in the 3D space. The curb is projected back onto the left image (Fig. 9), providing a first, visual way, to evaluate how robust the results are. This curb's inliers are removed from the set of candidate curb points and the whole process is repeated until no more polynomials are extracted.

Fig. 9. The 3D curb patch is projected onto the left image.

VI. Results and Future Work

The algorithm was tested both offline and online with an onboard dense stereo system. It runs in real-time, in less than 10 ms per frame. Implementation was done in C++, and the tests were carried out on a Pentium Dual-Core processor.

Curbs are robustly detected (see Fig. 11, on the next page) as long as 3D data is generally available along the curb, on each side.

The curb's height accuracy was evaluated on three different scenarios, with curbs having real heights of 7, 11 and 14 centimeters. The estimated curb height had an error of less than 5%. The curb vertical localization is accurate. Curbs were detected on a 10% uphill scenario (see Fig. 11, second from the top) and the slope of the curbs estimated with the proposed algorithm was 9.4%.

Future work is required for scenarios with poor 3D dense data. If sparse elevation data is available or the quality is poor, the polynomial curve is unstable. One such scenario can be seen in Fig. 10. A possible solution that will be investigated is the tracking of the results, in order to stabilize the detection along successive frames. False curbs might rarely appear. They can be removed by tracking or by increasing some of the internal thresholds from RANSAC (the minimum number of inliers, or the minimum segment length).

Fig. 10. Example of unstable polynomial curve caused by bad 3D data. Even though the real 3D curb is straight, the DEM curb "looks" undulated.

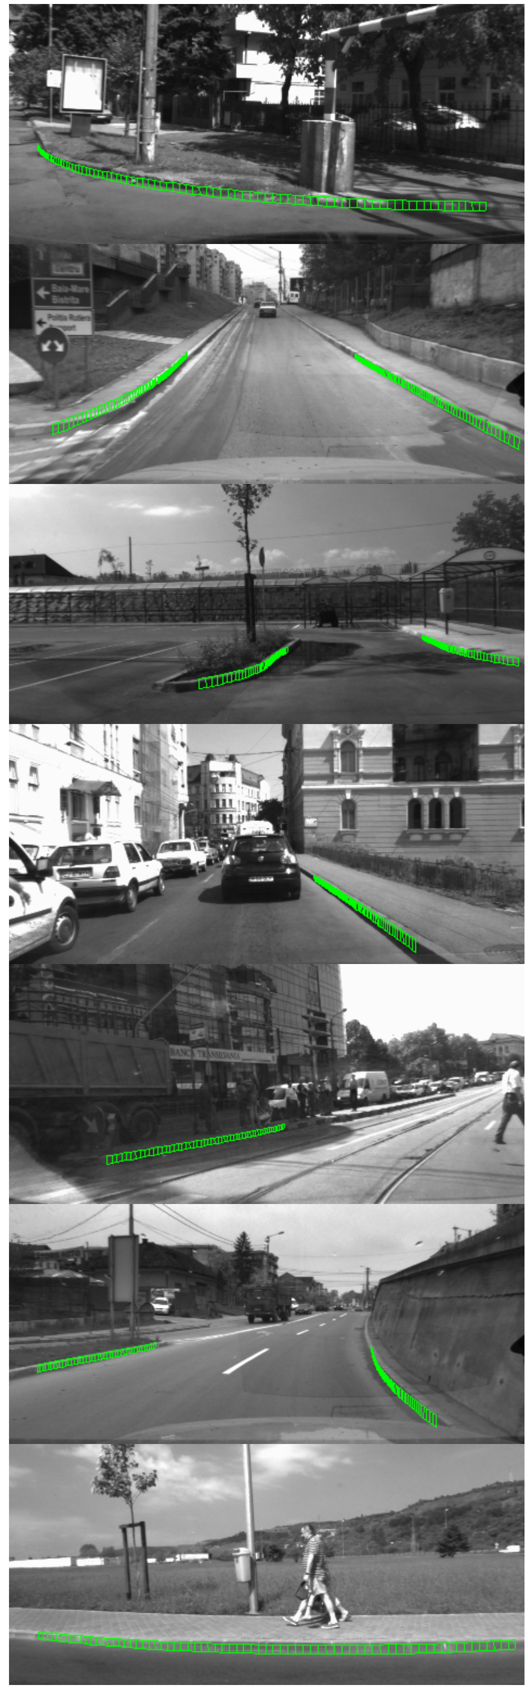

Fig. 11. Results of the proposed polynomial curb detection algorithm. Various road geometries are present, including a 10% uphill (second image from the top).

REFERENCES

[1] S. Se and M. Brady, "Vision–based Detection of Kerbs and Steps", *Proceedings of the British Machine Vision Conference*, BMVC 1997, pp. 1-10.

[2] R. Turchetto, R. Manduchi, "Visual Curb Localization for Autonomous Navigation", *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, Las Vegas, October 2003, pp. 1336-1342.

[3] X. Lu and R. Manduchi, "Detection and localization of curbs and stairways using stereo vision", *IEEE International Conference on Robotics and Automation* (ICRA '05), Barcelona, April 2005, pp: 4648-4654.

[4] F. Oniga, S. Nedevschi, M-M. Meinecke, "Curb Detection Based on a Multi-Frame Persistence Map for Urban Driving Scenarios", *Proceedings of the 11th International IEEE Conference on Intelligent Transportation Systems* (ITSC'08), 13-15 Oct. 2008, Beijing, China, pp. 67-72.

[5] K. R. S. Kodagoda, W. S. Wijesoma, and A. P. Balasuriya, "Road Feature Extraction Using a 2D LMS", *7th International Conference on Control, Automation, Robotics and Vision*, pp:453- 458, December 2002.

[6] K. R. S. Kodagoda, W. S. Wijesoma, and A. P. Balasuriya, "CuTE: Curb Tracking and Estimation", *IEEE Transactions on Control Systems Technology*, vol. 14, no. 5, pp. 951–957, Sept. 2006.

[7] C. Yu and D. Zhang, "Road Curbs Detection Based on Laser Radar", *8th International Conference on Signal Processing*, Volume 4, November 2006.

[8] J. Stuckler, H. Schulz, and S. Behnke, "In-lane Localization in Road Networks using Curbs Detected in Omnidirectional Height Images", *In Proceedings of Robotik 2008*, Munchen, June 2008.

[9] Gallo O., Manduchi R., and Rafii A., "Robust Curb and Ramp Detection for Safe Parking Using the Canesta ToF Camera", In IEEE Conference on Computer Vision & Pattern Recognition, Workshop on ToF-Camera based Computer Vision, June 2008.

[10] C. Thorpe et al., "Driving in Traffic: Short–Range Sensing for Urban Collision Avoidance", *Proceedings of SPIE: Unmanned Ground Vehicle Technology IV*, Vol. 4715, April 2002.

[11] R. Aufrere, C. Mertz, and C. Thorpe, "Multiple sensor fusion for detecting location of curbs, walls, and barriers", *IEEE Intelligent Vehicles Symposium*, IV2003, pp. 126–131.

[12] M. Fischler and R. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Graph. Image Process.*, vol. 24, no. 6, pp. 381–395, 1981.