2014 22nd International Conference on Pattern Recognition

Vision-Based Road Bump Detection Using a Front-Mounted Car Camcorder

Hua-Tsung Chen[1], Chun-Yu Lai[2], Chun-Chieh Hsu[2], Suh-Yin Lee[2], Bao-Shuh Paul Lin[1], Chien-Peng Ho[3]
[1] Information & Communications Technology Lab, National Chiao Tung University, Hsinchu, Taiwan
[2] Department of Computer Science, National Chiao Tung University, Hsinchu, Taiwan
[3] ICL/Industrial Technology Research Institute, Hsinchu, Taiwan
{huatsung, cylai, hsucch, sylee}@cs.nctu.edu.tw, bplin@mail.nctu.edu.tw, cpho@itri.org.tw

***Abstract*—Advanced vehicle safety is a recently emerging issue, appealed from the rapidly explosive population of car owners. Increasing driver assistance systems have been designed for warning drivers of what should be noticed by analyzing surrounding environments with sensors and/or cameras. As one of the hazard road conditions, road bumps not only damage vehicles but also cause serious danger, especially at night or under poor lighting conditions. In this paper we propose a vision-based road bump detection system using a front-mounted car camcorder, which tends to be widespread deployed. First, the input video is transformed into a time-sliced image, which is a condensed video representation. Consequently, we estimate the vertical motion of the vehicle based on the time-sliced image and infer the existence of road bumps. Once a bump is detected, the location fix obtained from GPS is reported to a central server, so that the other vehicles can receive warnings when approaching the detected bumpy regions. Encouraging experimental results show that the proposed system can detect road bumps efficiently and effectively. It can be expected that traffic security will be significantly promoted through the mutually beneficial mechanism that a driver who is willing to report the bumps he/she meets can receive warnings issued from others as well.**

Keywords—intelligent vehicle; driver assistance system; pattern recognition; signal processing; motion analysis

I. INTRODUCTION

The rapid expansion of car ownership worldwide necessitates the development of traffic surveillance systems and driver assistance systems. Incorporating techniques of sensing, communication, control, and video processing, these systems are intended for monitoring traffic conditions, improving safety, enhancing mobility, and so on. As an important component in Intelligent Transportation System (ITS), computer-assisted traffic surveillance enables manifold functionalities. With the deployment of fixed sensors or cameras along roads or highways, real-time traffic information, such as traffic flow and average driving speed, can be provided to drivers for avoiding the congestion area [1], [2]. Also, traffic accidents or violations can be detected to promptly inform the police or rescue workers [3].

The research on driver assistance (DA) systems is another trend for ensuring security and preventing accident. Since a driver may not maintain a high level of concentration for a prolonged period of time, most DA systems are designed primarily for calling attention to potential dangers. Playing an essential role in DA systems, lane detection/tracking aims at locating the lane markings or boundaries on the road surface [4]-[6]. Not only the information about the lane marking patterns (solid or dash, single or double, etc.) and road types (straight or curve) can be provided, but the position of the vehicle within a lane can also be continually monitored. Thus, lane detection leads to many applications such as lane departure warning, lane change assistance, route direction guidance, and so on.

Object detection is another crucial element in developing intelligent vehicles. Many research works utilize cameras or multiple sensors for detecting/sensing vehicles, pedestrians, or obstacles nearby so as to warn drivers of the vehicles in the blind-spot areas [7], assist drivers in noticing that there are pedestrians or obstacles around the vehicle [6] or give alarms when being too close to the front vehicle [8]. As for collision avoidance at night, the detection of taillights and brake lights is of vital importance [8], [9]. Some works endeavor to achieve collaborative pre-collision warning through inter vehicle communication (IVC) and vehicular ad-hoc network (VANET) [10], [11].

Most DA systems focus on enhancing safety by informing drivers of what should be noticed, while some systems are intended for providing convenient and realistic functionalities to ease driving efforts, such as adaptive cruise control [12], automatic parking [13], and traffic sign recognition [14]. Gaining more attention than ever, the detection or monitoring of driver inattention, distraction, or drowsiness is another nascent topic in DA applications [15].

DA systems analyze surrounding environments with sensors and cameras for warning drivers of what should be noticed, while traffic surveillance systems usually survey wider areas for providing global information. However, the spatial coverage of the static traffic surveillance systems is still limited. Consequently, the concept of community sensing is brought up [16], which harnesses the wealth of mobile *probe data* via the cameras/sensors of DA systems to obtain real-time or historical information. The resources of a potential extensive and

This research is supported in part by projects ICTL-103-Q528, ATU-103-W958, NSC-102-2221-E-009-031, and NSC-101-2221-E-009-087-MY3.

1051-4651/14 $31.00 © 2014 IEEE
DOI 10.1109/ICPR.2014.776

distributed set of mobile cameras/sensors can be leveraged to enable large-scale sensing [17]. Orchestrating both the static and mobile cameras/sensors of TS and DA systems, more explicit and comprehensive data and inferences can be obtained.

Road bump detection has often been mentioned in the literature [17]-[19]. Road bumps not only damage vehicles but also cause danger to drivers or pedestrians. Especially at night or under poor lighting conditions, informing drivers of hazardous road conditions is of great importance. The task of road surface monitoring, such as bump detection, cannot be easily achieved by the static sensors/cameras deployed on roads, but demands *mobility* to accomplish. There are increasing community sensing works on road bump detection using the accelerometers in smartphones [18], [19]. However, a smartphones is not fixed on the vehicle, and may be put in a pocket or even be taken out of the vehicle. Primarily, a smartphone is someone's phone, which is very likely to be used in the vehicle, maybe not by the driver, but by a passenger. The detection results may be badly influenced by human motion. Hence, in this paper, we propose a vision-based road bump detection system using a front-mounted car camcorder. The rapid development and reduced cost of video capturing devices have made it economically feasible to deploy front-mounted car camcorders, which have stable power supply and can record the scenes along the way. Since these camcorders are fixed on vehicles, the proposed road bump detection system has the advantage of not being influenced by human motion. Furthermore, the images at the vicinities of the detected bumps can also be provided, displaying the road conditions and what cause the bumps.

The rest of this paper is organized as follows. In the next section, we give an architecture overview of the proposed road bump detection system. In Section III, the main processing modules, including time-sliced image generation, vertical motion estimation, and road bump locating are elaborated. Section IV presents the experimental results with comparison. Finally, we conclude this paper in Section V.

II. System Architecture

Based on the concept of community sensing, we outline a system framework for distributed road bump detection using front-mounted car camcorders, as illustrated in Fig. 1. First, a vision-based road bump detection algorithm is proposed, thereby allowing the vehicles equipped with car camcorders to be deemed as distributed sensors with high mobility. With stable power supply, the car camcorder keeps capturing the scenes all the way. Once a bump is detected, the location fix obtained from GPS is reported to a central server and is added into the bump location database as a *suspect bump*. The central server maintains a database of detections from multiple vehicles and clusters the detections based on location. If several suspect bumps are reported in the same vicinity, the region will be regarded as *bumpy*, and marked with a pushpin on a map. This mechanism facilitates filtering out the spurious detections which are not real road anomalies. The confidence and severity of the

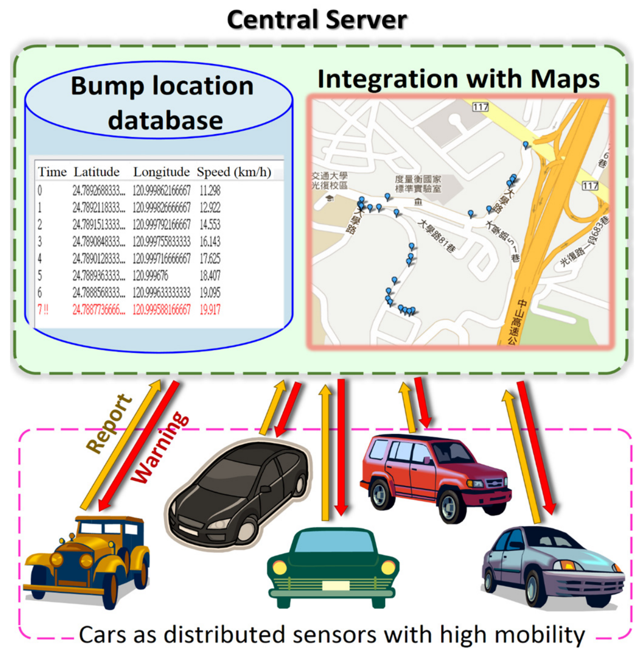

Fig. 1. Architecture of distributed road bump detection using front-mounted car camcorders.

bumpy region (corresponding to the report time and vibration amplitude, respectively) will also be recorded in the database. Hereafter, vehicles approaching the detected bumpy regions can receive warning signals from the server, informing the drivers of hazard road conditions. In this paper, we focus specially on the *sensing component*, which uses the front-mounted camcorder equipped on vehicles to detect road bumps.

The main contributions of this work are summarized as follows:

- To the best of our knowledge, we propose the first vision-based system for detecting road bumps using a front-mounted car camcorder, which has the advantages of high mobility and not being influenced by human motion.
- It is quite an inexpensive and practical way to detect abnormal road conditions by utilizing the car camcorders, which tend to be widespread deployed. The task of large-scale sensing is achieved merely by means of software operating on existing hardware. The proposed scheme requires no additional hardware and can avoid the dependence on infrastructure.
- A mutually beneficial mechanism is proposed that a driver who is willing to provide the information of abnormal road conditions which he/she meets can receive warnings issued from others as well. These warnings are especially of vital importance at night or under poor lighting conditions.

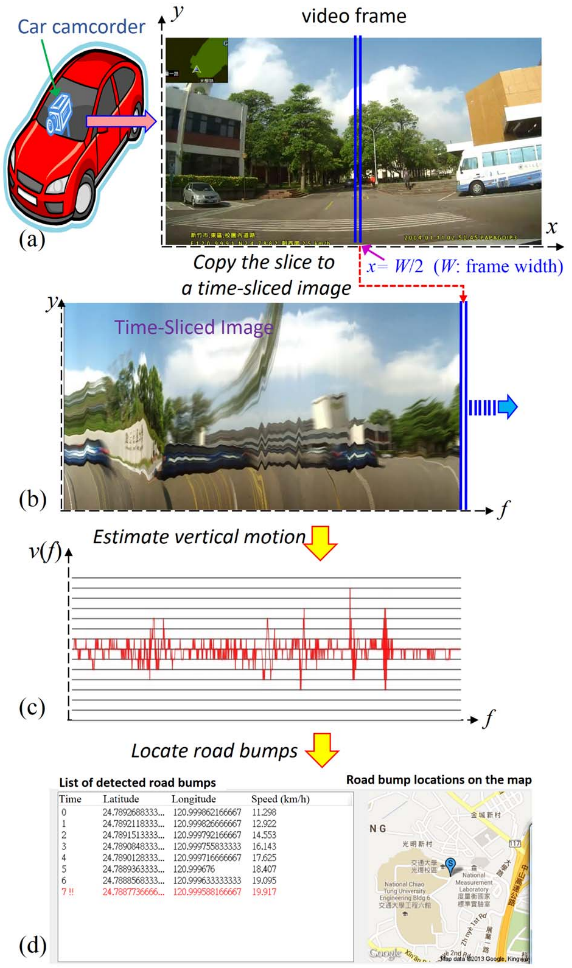

Fig. 2. Flowchart of the proposed vision-based road bump detection.

III. Road Bump Detection

A camera (the car camcorder) mounted frontward on a vehicle is used to capture the scenes along the path, as shown in Fig. 2(a). Therefore, we can infer the vehicle motion from the camera motion. There have been lots of researches on camera motion estimation in the literature by matching points, lines, and regions [20], tracing non-vertical edges [21], or utilizing image registration methods [22]. However, the iterative processing of video frames and statistics computation is time consuming. Moreover, these features which provide strong clues to camera motions are not always workable for the front-mounted car camcorder. Matching errors may also occur due to some problems such as repeated patterns and object occlusions.

Drawing the inspiration from the idea of time-sliced image in [5], we present a vision-based road bump detection algorithm via analyzing the vertical motion of a vehicle through a time-sliced image generated from the video captured by a front-mounted car camcorder. The generated time-sliced image contains only a very small fraction of data compared to the original video, and much redundancy in consecutive frames is discarded. Only one pixel line from each frame is extracted for motion estimation, so the proposed approach can be very efficient.

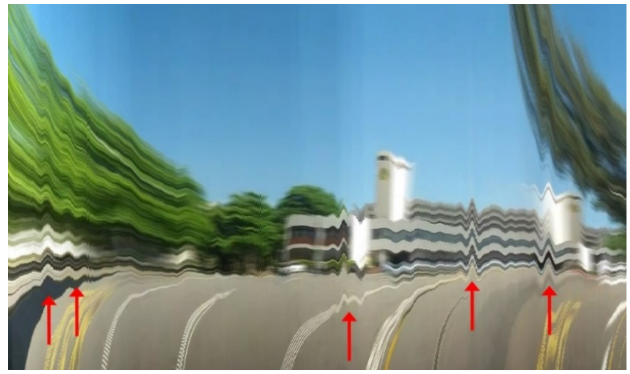

Fig. 3. Example of time-sliced image with jitters, indicated by red arrows.

The flowchart of the proposed road bump detection algorithm is illustrated in Fig. 2. Three major processing steps including time-sliced image generation, vertical motion analysis, and road bump locating are explained as follows.

A. Time-Sliced Image Generation

For road bump detection, the vertical motion of the camera is of our interest. Thus, we generate a time-sliced image by extracting a vertical slice (instead of a horizontal one as in [5]) at the x-coordinate $x = W/2$ (W: frame width) from each frame f, as shown in Fig. 2(a), and then paste it into a continuous image memory, as shown in Fig. 2(b), wherein the horizontal axis indicates the frame index and the vertical axis indicates the y-coordinate as in the video frame.

For noise smoothing, instead of directly extracting a one-pixel wide vertical slice from each frame, we extract an n-pixel wide slice and condense it into one-pixel wide by taking the horizontal average of each row. As a result, a video containing F frames with resolution of $W \times H$ (H: frame height) will produce a time-sliced image with size of $F \times H$.

B. Vertical Motion Estimation

An inherent property of the time-sliced image is that traversing along the horizontal axis is equivalent to tracing through different frames in the video or different road sections on the way. Hence, it is an efficient way to acquire the vertical motion of the vehicle through the generated time-sliced image.

In general, the time-sliced image should be smooth when the vehicle is moving on an even road. However, the time-sliced image would be jagged and waved once the camera suffers from vehicle shaking. Fig. 3 displays an example of time-sliced image with jitters, indicated by red arrows. Utilizing this property, we can analyze the vehicle motion for inferring road bumps.

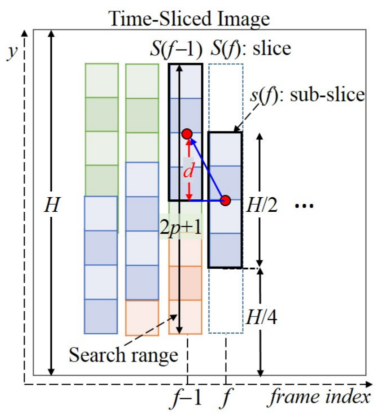

Fig. 4. Illustration of vertical motion estimation.

The illustration of vertical motion estimation is presented in Fig. 4. For each slice $S(f)$ in the time-sliced image (f: frame index), we attempt to compute its vertical displacement from the previous slice $S(f-1)$. It can be observed from Fig. 2(b) that most texture information exists in the central part and there tends to be less texture in the air and on the road surface, i.e., the upper and lower parts. For efficiency, we extract from each current slice $S(f)$ the central part of $H/2$ in height, termed a *sub-slice* $s(f)$, and then search the most similar sub-slice in the previous slice.

Let $C_k(f, y)$ be pixels in the time-sliced image, where $k \in \{r, g, b\}$ for three color channels. The difference between the two sub-slices can then be measured by their *Mean Square Error* (MSE), defined as:

$$MSE(f,d) = \frac{1}{H/2}\sum_{y=H/4}^{H\times 3/4-1}\sum_{k=r,g,b}\left(C_k(f,y)-C_k(f-1,y+d)\right)^2, \quad (1)$$

where d represents the vertical displacement, limited to a range $[-p, p]$. The goal of the search is to find a displacement d as the vertical motion $v(f)$ such that $MSE(f, d)$ is minimum:

$$v(f) = \{d \mid MSE(f,d) \text{ is minimum}, \quad d \in [-p,p]\}. \quad (2)$$

Finally, we can obtain the vertical motion $v(f)$ at each frame index f, as shown in Fig. 2(c).

For our application, the mean absolute difference is sufficient to find good match of sub-slices. Of course, this measure is by no means the only possible choice. Some other error measures, such as the *Mean Absolute Difference* (MAD) or simply *Sum of Absolute Difference* (SAD), may also be appropriate.

C. Road Bump Locating

As shown in Figs. 2(b) and (c), road bumps result in notable vertical jerks, which are expected to register in the signal of vertical motion $v(f)$. Also, since the vehicle shifts both up and down, there would be both peaks and dips in $v(f)$. Hence, we can infer the existence of road bumps by detecting both peaks and dips in $v(f)$ within a small time duration. Here, we apply a sliding-window scheme to accomplish this task, as illustrated in Fig. 5. In each iteration, we compute the difference between the *max* and *min* values of $v(f)$ in an observation window of w frames. If the difference is greater than a threshold δ, it can be judged that there exists a bump in the time duration of the observation window. Then, the window is moved forward, and the above process is iterated. The process of bump detection within an observation window can be formulated as:

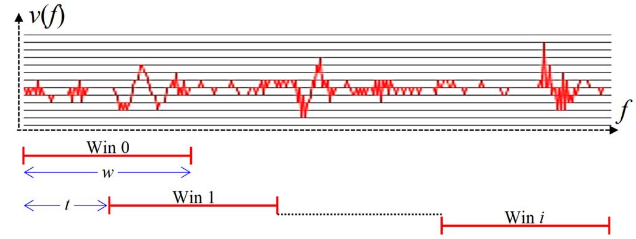

Fig. 5. Sliding window with some overlap, wherein w and t represent the window size and the step size between adjacent windows, respectively.

$$B(i) = \begin{cases} 1, if \max\limits_{f\in[t\cdot i,\ t\cdot i+w]} v(f) - \min\limits_{f\in[t\cdot i,\ t\cdot i+w]} v(f) > \delta \\ 0, otherwise \end{cases} \quad (3)$$

where $B(i)$ indicates whether a bump is detected or not, i is the window index, w and t are the window size and the step size between adjacent windows, respectively.

In our system, the reception rate of GPS data are received once per second, so the window size w and step size t are set to 45 and 30 frames, respectively, which involves 50% overlap between adjacent windows. Finally, when a bump is detected, the location fix obtained from GPS data is reported to the central server, marked with a pushpin on a map, as presented in Fig. 2(d), enabling the other vehicles to receive warnings when approaching the detected bumpy regions.

IV. Experimental Results

To evaluate the performance of the proposed road bump detection algorithm, we conduct the experiments on a video data set collected by a front-mounted car camcorder fixed on a sedan. The video resolution is 960 × 480 and the frame rate is 30 fps. We select 25 video clips (about 5 minutes per clip and 207810 frames in total) containing bumps under different road conditions, including highway, urban, campus, etc. The proposed system is implemented in C++ with OpenCV 2.46 libraries [23], and the experiments run on a notebook (AMD Athlon II X2 240 Dual-Core Processor @2.679 MHz, 8GB RAM, Windows 7 64-bit OS) show that the proposed system can achieve a very high processing speed of over 300 fps.

TABLE I. RESULTS OF ROAD BUMP DETECTION WITH DIFFERENT δ VALUES

δ	*#correct*	*#false*	*Precision*	*Recall*	*F-score*
1	338	656	0.34	1	0.51
2	336	504	0.4	0.99	0.57
3	333	441	0.43	0.99	0.6
4	325	288	0.53	0.96	0.68
5	292	131	0.69	0.86	0.77
6	**249**	**58**	**0.81**	**0.74**	**0.77**
7	186	40	0.82	0.55	0.66
8	142	29	0.83	0.42	0.56

(# of ground truth bumps = 338, *#correct*: number of correct detections, and *#false*: number of false alarms.)

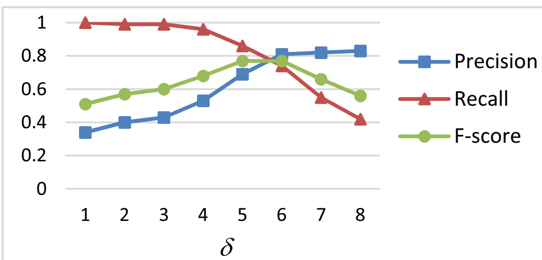

Fig. 6. Results of road bump detection with different δ values.

As mentioned in [18], [19], it is difficult to establish the ground truth for road bumps. Manually annotation is a typical way, but it is subjective and may be influenced by how the vehicle crosses over the bump and at what speed. In our experiments, the ground truth data are generated by inspecting the video clips and marking the occurring time of each bump manually. In total, there are 338 ground truth bumps manually marked in our testing data. A bump is said to be detected *correctly* if the time difference between the system output detection and the ground truth one is less than one second. The proposed system is evaluated in terms of *precision*, *recall*, and *F-score*, as defined by:

$$precision = \#correct\,/(\#correct + \#false), \quad (4)$$

$$recall = \#correct/\#gt, \quad (5)$$

$$F = 2 \cdot precision \cdot recall\,/(precision + recall) \quad (6)$$

where *#correct* and *#false* denote the numbers of correct detections and false alarms, and *#gt* represents the total number of ground truth bumps (i.e., *#gt* = 338). The results of road bump detection with different δ values, which is used in Eq. (3), are presented in Table I and Fig. 6. Based on the concept of community sensing, it is often the case that a bump mis-detected by one car can still be detected by others. Therefore, we prefer a high δ value to assure the precision. In our experiments, δ is set to 6 and the precision and recall rates are about 0.8 and 0.75, respectively. By inspection, we find that errors mostly occur when the car turning at a corner with complex scene variation or facing a scene with messy texture, as presented in Fig. 7, wherein the left column shows the video frame captured by a car camcorder and the right one shows a part of the time-sliced image. We have ongoing research to eliminate these error cases by applying more constraints.

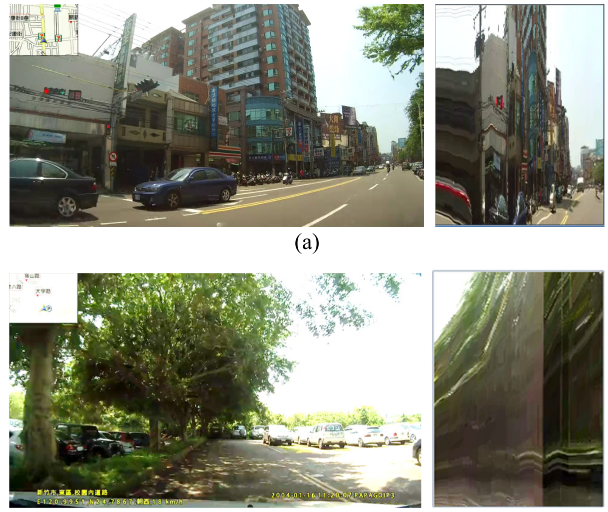

(a)

(b)

Fig. 7. Error cases: (a) turning at a corner with complex scene variation; (b) facing trees with messy texture. (Left: original video frame; right: part of a time-sliced image.)

TABLE II. RESULTS OF THE COMPARATIVE PHASE CORRELATION-BASED METHOD.

δ	*#correct*	*#false*	*Precision*	*Recall*	*F-score*
1	337	625	0.35	1	0.52
2	329	402	0.45	0.97	0.61
3	305	230	0.57	0.9	0.7
4	249	140	0.64	0.74	0.69
5	181	93	0.66	0.54	0.59
6	135	82	0.62	0.4	0.49
7	101	86	0.54	0.3	0.39
8	68	86	0.44	0.2	0.28

(# of ground truth bumps = 338, *#correct*: number of correct detections, and *#false*: number of false alarms.)

For performance comparison, we also implement another road bump detection method based on *phase correlation*, which is originally proposed for the registration of translated images [22]. To avoid the influence of the hood of the car and the overlaid map on frames (see Fig. 7), we perform phase correlation on only the center quarter region ($W/2 \times H/2$) of each frame and estimate the relative translation offset (T_x, T_y) between consecutive frames. The y-directional offset T_y is utilized for bump detection and other settings are the same as our proposed approach. The detection results of this phase correlation-based method are shown in Table II, and the comparison between the two approaches on precision and recall is presented in Fig. 8. By inspecting the error cases, we find that the phase correlation-based method is more likely to produce the false alarms resulted

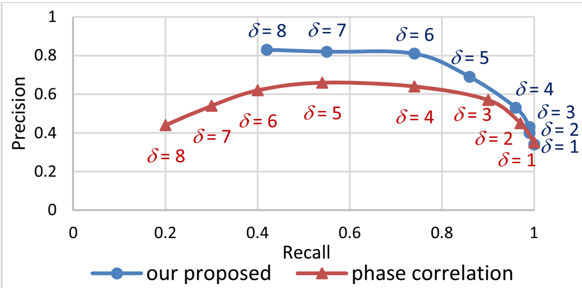

Fig. 8. Comparison between the phase correlation-based method and our proposed approach.

TABLE III. COMPARISON ON PROCESSING SPEED.

	Phase correlation-based	Our proposed
# of total frames	207810	207810
Processing time	3779 sec	571 sec
Frame per second	54.99 fps	363.94 fps

from the moving object(s) in frames. Table III shows the comparison on processing speed. To process a total of 207810 frames, the correlation-based method takes 3779 seconds, while our proposed approach takes only 571 seconds. Overall, our proposed road bump detection approach can achieve better results with much higher processing efficiency.

V. CONCLUSION

Bumps on the road may cause serious danger, especially when the lighting condition is poor or the driver is distracted. In this paper, we propose a vision-based system capable of detecting road bumps using a front-mounted car camcorder. Unlike the static traffic surveillance systems, car camcorders have the advantage of high mobility. Hence, it is an inexpensive and practical way to achieve the task of large-scale sensing by the widespread deployed car camcorders. Experiments on real video data captured by a front-mounted car camcorder show that our proposed system can effectively detect bumps with very high efficiency. As we believe, through the mutually beneficial mechanism that a driver who is willing to report the bumps he/she meets can receive warnings issued from others as well, traffic security will be significantly promoted.

REFERENCES

[1] B. Tian, Q. Yao, Y. Gu, K. Wang, and Y. Li, "Video processing techniques for traffic flow monitoring: A Survey," in *Proceedings of 14th International IEEE Conference on Intelligent Transportation Systems*, pp. 1103-1108, 2011.

[2] H. T. Chen, L. W. Tsai, H. Z. Gu, S. Y. Lee, and B. S. P. Lin, "Traffic congestion classification for nighttime surveillance videos," in *Proceedings of IEEE International Conference on Multimedia and Expo Workshops*, pp. 169-174, 2012.

[3] M. S. Shehata, J. Cai, W. M. Badawy, T. W. Burr, M. S. Pervez, R. J. Johannesson, and A. Radmanesh, "Video-based automatic incident detection for smart roads: the outdoor environmental challenges regarding false alarms," *IEEE Transactions on Intelligent Transportation Systems*, vol. 9, no. 2, pp. 349-360, 2008.

[4] J. C. McCall and M. M. Trivedi, "Video-based lane estimation and tracking for driver assistance: survey, system, and evaluation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 7, no. 1, pp. 20-37, 2006.

[5] A. Borkar, M. Hayes, and M. T. Smith, "A novel lane detection system with efficient ground truth generation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 1, pp. 365-374, 2012.

[6] M. M. Trivedi, T. Gandhi, and J. McCall, "Looking-in and looking-out of a vehicle: Computer-vision-based enhanced vehicle safety," *IEEE Transactions on Intelligent Transportation Systems*, vol. 8, no. 1, pp. 108-120, 2007.

[7] B. F. Lin, Y. M. Chan, L. C. Fu, P. Y. Hsiao, L. A. Chuang, S. S. Huang, and M. F. Lo, "Integrating appearance and edge features for sedan vehicle detection in the blind-spot area," *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 2, pp. 737-747, 2012.

[8] R. O'Malley, E. Jones, and M. Glavin, "Rear-lamp vehicle detection and tracking in low-exposure color video for night conditions," *IEEE Trans. on Intelligent Transportation Systems*, vol. 11, no. 2, pp. 453-462, 2010.

[9] D. Y. Chen, Y. H. Lin, Y. J. Peng, "Nighttime brake-light detection by Nakagami imaging," *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 4, pp. 1627-1637, 2012.

[10] T. Taleb, A. Benslimane, and K. Ben Letaief, "Toward an effective risk-conscious and collaborative vehicular collision avoidance system," *IEEE Transactions on Vehicular Technology*, vol. 59, no. 3, pp. 1474-1486, 2010.

[11] M. L. Sichitiu and M. Kihl, "Inter-vehicle communication systems: a survey," *IEEE Communications Surveys and Tutorials*, vol. 10, no. 2, pp. 88-105, 2008.

[12] S. Li, K. Li, R. Rajamani, and J. Wang, "Model predictive multi-objective vehicular adaptive cruise control," *IEEE Transactions on Control Systems Technology*, vol. 19, no. 3, pp. 556-566, 2011.

[13] H. G. Jung, Y. H. Lee, and J. Kim, "Uniform user interface for semiautomatic parking slot marking recognition," *IEEE Transactions on Vehicular Technology*, vol. 59, no. 2, pp. 616-626, 2010.

[14] X. Baró, S. Escalera, J. Vitrià, O. Pujol, and P. Radeva, "Traffic sign recognition using evolutionary adaboost detection and forest-ECOC classification," *IEEE Transactions on Intelligent Transportation Systems*, vol. 10, no. 1, pp. 113-126, 2009.

[15] Y. Dong, Z. Hu, K. Uchimura, and N. Murayama, "Driver inattention monitoring system for intelligent vehicles: a review," *IEEE Transactions on Intelligent Transportation Systems*, vol. 12, no. 2, pp. 596-614, 2011.

[16] A. Krause, E. Horvitz, A. Kansal, and F. Zhao, "Toward community sensing," in *Proceedings of 7th IEEE International Conference on Information Processing in Sensor Networks*, pp.481-492, 2008.

[17] T. Das, P. Mohan, V. N. Padmanabhan, R. Ramjee, and A. Sharma, "PRISM: platform for remote sensing using smartphones," in *Proceedings of 8th ACM International Conference on Mobile Systems, Applications, and Services*, pp. 63-76, 2010.

[18] P. Mohan, V. N Padmanabhan, and R. Ramjee, "Nericell: rich monitoring of road and traffic conditions using mobile smartphones," in *Proceedings of 6th ACM Conference on Embedded Network Sensor Systems*, pp. 323-336, 2008.

[19] J. Eriksson, L. Girod, B. Hull, R. Newton, S. Madden, and H. Balakrishnan, "The pothole patrol: using a mobile sensor network for road surface monitoring," in *Proceedings of 6th ACM International Conference on Mobile Systems, Applications, and Services*, pp. 29-39, 2008.

[20] A. Rav-Acha, G. Engel, and S. Peleg, "Minimal aspect distortion mosaic," *International Journal of Computer Vision*, vol. 78, no. 2-3, pp. 187-206, 2008.

[21] J. Y. Zheng, Y. Bhupalam, and H. T. Tanaka, "Understanding vehicle motion via spatial integration of intensities," in *Proceedings of 19th International Conference on Pattern Recognition*, pp. 1-5, 2008.

[22] B. Zitova and J. Flusser, "Image registration methods: a survey," *Image and Vision Computing*, vol. 21, no. 11, pp. 977-1000, 2003.

[23] OpenCV, 2013. Available: http://opencv.org/