Outdoor Context Awareness Device That Enables Mobile Phone Users to Walk Safely through Urban Intersections

Jihye Hwang[1], Younggwang Ji[2], Nojun Kwak[1] and Eun Yi Kim[2]
[1]*Graduate School of Convergence Science and Technology, Seoul National University, Seoul, Korea*
[2]*Internet & Multimedia Engineering, Konkuk University, Seoul, Korea*

Keywords: Guidance System, Mobile Phone User, Outdoor Context Awareness, Support Vector Machine, Multi-scale Classification.

Abstract: Research in social science has shown that the mobile phone users pay less attention to their surroundings, which exposes them to various hazards such as collisions with vehicles than other pedestrians. In this paper, we propose a novel handheld device that assists mobile phone users to walk more safely outdoors. The proposed system is implemented on a smart phone and uses its back camera to detect the current outdoor context, e.g. traffic intersections, roadways, and sidewalks, finally alerts the user of unsafe situations using sound and vibration from the phone. The outdoor context awareness is performed by three steps: pre-processing, feature extraction, and context recognition. First, it improves the image contrast while removing image noise, and then it extracts the color and texture descriptors from each pixel. Next, each pixel is classified as an intersection, sidewalk, or roadway using a support vector machine-based classifier. Then, to support the real-time performance on the smart phone, a multi-scale classification is applied to input image, where the coarse layer first discriminates the boundary pixels from the background and the fine layer categorizes the boundary pixels as sidewalk, roadway, or intersection. In order to demonstrate the effectiveness of the proposed method, some real-world experiments were performed, then the results showed that the proposed system has the accuracy of above 98% at the various environments.

1 INTRODUCTION

When walking through an urban area, we can easily observe pedestrians talking or reading and sending text messages on mobile phones. A number of studies on pedestrians crossing streets have shown that mobile phone users exhibit less safe behavior than other pedestrians, as their attention has been directed away from the external environment and they do not see the overall context (Institute of transportation engineers, 2004; Kate, 2010; Leena, 2012; Mark, 2009; Rovert, 2010; Tianyu, 2012).

The mobile phone users pay less attention to their surroundings, which exposes them to various hazards such as collisions with people, cars, and other obstacles. Among the numerous hazards outside, traffic intersections are the most dangerous areas for pedestrians. Approximately 20.3% of all traffic accidents occur at intersections and this pedestrian accident rate is one of the highest rates for all traffic accidents (Jack, 2007; Julie,2006).

During last decades, a number of solutions for safety crossing the traffic intersections has been proposed and implemented (Brabyn, 1933; Barlow, 2003; Dragan, 2011; Karacs, 2006; Volodymyr, 2008; Chen-Fu, 2012; James, 2013). Then, the initial research has been focusing the development of the safety system for the people with visual impairments. For example, the audible pedestrian signals (Barlow, 2003) and talking signs (Brabyn, 1933) have been developed in order to inform the visually impaired pedestrians to know when to cross intersections. However, although these solutions are being adopted more widely, they are still only available in limited places and additional devices should be installed or used.

As an alternative to such systems, some vision-based systems such as Crosswatch (Volodymyr, 2008), Zebralocalizer (Dragan, 2011) and Walksafe (Tianyu, 2012) have been proposed to prevent the pedestrians of the collisions with vehicles at urban traffics. Crosswatch is a handheld vision system that provides real-time feedback to the user about their

Hwang, J., Ji, Y., Kwak, N. and Kim, E.
Outdoor Context Awareness Device That Enables Mobile Phone Users to Walk Safely through Urban Intersections.
DOI: 10.5220/0005664705260533
In *Proceedings of the 5th International Conference on Pattern Recognition Applications and Methods (ICPRAM 2016)*, pages 526-533
ISBN: 978-989-758-173-1

orientation and location at the crosswalk; then, it can detect the crosswalk using edge information. Similarly, the Zebralocalizer identifies pedestrian crossings (i.e. zebra crossings) using line analyses and localization of crosswalks using a camera and 3D accelerometers. These vision-based methods function successfully on localizing crosswalks at intersections and guiding users to cross intersections safely.

Recently, some commercial products such as “Text and Walk,” and “Walk and Email,” which displays the road situation captured from back camera on application background, thereby making pedestrians users write SMS and e-mail while walking safety. However, it was shown in (Ophir, 2009) that the users may not be aware of dangers even if they are displayed as application background. In (Sivaraman, 2010), authors presented a car detection system based on Harr features that exploits the back camera of smartphone. It was well worked on the resource constrained smartphone, however, it can only detect cars when they were very close to the pedestrians, limiting the time for the pedestrians to react safety.

On the other hand, none of them have considered the current situations where a user stands on. In real scenarios, the people require different guidance solutions according to the context where they are currently located. For example, if a user is at an intersection, they want to locate the crosswalk. However, when the user is walking on the sidewalk, the system leads the user to walk on the far side from the road. Accordingly, the outdoor contexts where the user is located should be first recognized.

In this paper, a novel method for automatically recognizing a user’s current context is proposed in order to increase pedestrian safety, particularly for users who operate their mobile phone while walking. Here, the context refers to the type of place where a user is standing, which is classified as a sidewalk, roadway, or intersection. Among these types of contexts, the discrimination between a sidewalk and an intersection are more important than recognizing a roadway.

As a key in discriminating outdoor contexts, the orientation of the boundaries between sidewalks and roadways are used: horizontally oriented boundaries are found in images corresponding to intersections and more vertically oriented boundaries are observed in images corresponding to sidewalks. Therefore, localizing such boundaries from input images should be undertaken first. In order to separate the boundaries between sidewalks and roadways from other lines, and then to discriminate such boundaries as sidewalks or intersections, the color and texture properties of the images are considered and machine-learning based classifications such as a support vector machine (SVM) are used. Then, in order to improve the computation cost and accuracy, a multi-scale classification is adopted, where a coarse layer first classifies the boundary pixels from the background and a fine layer classifies the boundary pixels into one of the three contexts: sidewalks, intersections, or roadways.

In order to evaluate the effectiveness of the proposed system, numerous videos were collected from real environments, and they were used to measure the accuracy of the proposed system. From the experimental results, it was found that the average accuracy was 98.25%.

2 SYSTEM ARCHITECTURE

We propose a novel assistive device that aids mobile phone users walking and crossing roads more safety. The proposed system is implemented on smartphone and uses its back camera to detect users’ current context, and notifies the recognized results to users through sound and vibration from the phone.

Then, as a key element of discriminating between sidewalks and intersections, the orientation of the boundaries between the sidewalks and roadways are used. Fig 1 illustrates some sample intersections and sidewalks captured from outdoors, where the vertical and horizontal lines in the boundaries between sidewalks and roadways can be easily observed. The images corresponding to intersections have horizontal boundaries, as shown in Fig. 1(a), whereas the images corresponding to sidewalks have some boundaries that are close to vertical and non-horizontal lines (see Fig. 1(b)).

(a) (b)

Figure 1: Some of outdoor images (a) images categorized to intersections (b) images categorized to sidewalks.

Based on these observations, the proposed method was designed and developed. As seen in Fig. 1, it is critical to accurately localize the boundaries from the input images. For this, we use the color and textural properties, which are trained by machine learning algorithm.

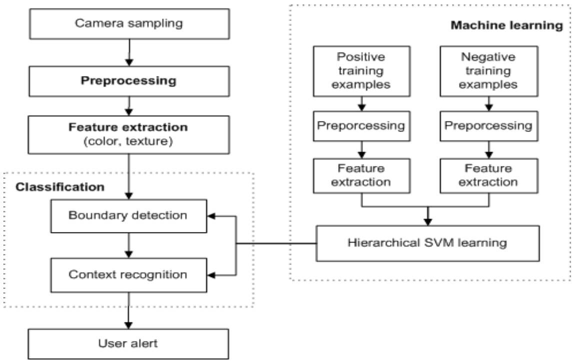

Figure 2: System overview.

The core algorithm to discriminate users' current contexts as sidewalk, roadway and intersections is based on computer vision techniques. Computer vision processing is a computational intensive process that can easily drain the computational resources and batteries of smartphone. To address this, the proposed system use the simple image features such as colors and textures, and use the learning-based recognition model that is first trained offline and then uploaded and used for the outdoor situation recognition, as shown in Fig 2.

To build a learning based recognition model (the right side of Fig 2), we first prepare a dataset containing positive and negative images, for example, images that show the sidewalk, roadway and intersection (see Fig 1). Such sets of images are first preprocessed to enhance the contrast while filtering some noises, then input to an algorithm that extracts characterizing features such as colors and textures and then used to build classifiers able to determine if a picture contains sidewalk or traffic intersection. The resulting classifiers are then used by the smartphone application running on smartphones, running the online context recognition in real-time.

In on-line scenarios, the proposed system is performed by three steps: preprocessing, feature extraction, and context recognition, which are described in the left side of Fig 2. In preprocessing, some noise and contrast are removed and improved using Gaussian smoothing and histogram equalization. Then, the useful visual features are extracted including the color and textures; saturation and intensity are used as color descriptors and a histogram of oriented gradient (HOG) is used to describe the textural properties. These features then become the input for the context recognition module, which separates the boundaries from the background and re-categorizes them as pixels corresponding to sidewalks or intersections. In this study, multi-scale classification is used to reduce the computational cost of the classifier and improve the accuracy.

3 FEATURE EXTRACTION

As a key element of discriminating the outdoor contexts, the orientation of the boundaries between sidewalks and roadways are used. In order to discriminate the boundaries between sidewalks and roadways from other lines, the visual characteristics such as color and texture are investigated. Here, the saturation and intensity are used to describe the color and a histogram of oriented gradient (HOG) is used to describe the texture.

The color information is used as inputs for classification problems. In this study, the RGB image is converted to the HSI color model using below equation, because the HSI color model represents colors in a similar manner to that which human eyes sense them.

Generally, the pixels that correspond to the roadway have distinctive color information, i.e. these regions have lower color saturations than others, as half of the color belongs to roadways that are marked with gray. Accordingly, the saturation and intensity are used to discriminate these contexts. Then, the saturation and intensity for every M×M sized sub-region is computed as the average of each quarter sub-region.

In order to describe the textural properties of sidewalks and intersections, various texture descriptors have been considered such as the histogram of oriented gradient (HOG), fast Fourier transform (FFT) coefficients, and wavelet transform. Through the experiment, the HOG method was selected as the most reliable method to describe the textural properties.

The HOG is a feature descriptor that counts the occurrences of a gradient orientation in the sub-regions of an image, and it is often used for object detection. This method is computed on a dense grid of uniformly spaced cells and uses overlapping local contrast normalization for improved accuracy.

In this paper, the contexts are defined as sidewalks and intersections; these two contexts have different characteristics in terms of the gradient orientations. The boundaries of sidewalks are oriented close to vertical, and the horizontally oriented boundaries were found in the images corresponding to intersections. However, other image parts corresponding to the same objects have a relatively uniform gradient orientation. Therefore, the HOG is effective in both separating the boundaries between sidewalks and roadways from other images, and categorizing the contexts.

For all M×M-sized sub-regions, the HOGs are calculated using the following procedures. First, the gradient magnitude and orientation of each pixel are computed using the Sobel operator; then, each pixel's gradient orientation is discretized into six histogram bins, and the pixels' magnitudes are accumulated in the corresponding HOG bin. Finally, the HOG is normalized.

Fig 3 presents HOG characteristics for some different types of regions. Fig 3(a) shows the HOG distributions for the sub-regions corresponding to the sidewalks and intersections, and Fig 3(c) shows those for parts of some objects. Interestingly, in such figures, distinctive distributions are seen according to their region type. Generally, for the sub-regions of the sidewalks and intersections, the distribution of the HOGs tends to be biased at specific histogram bins. However, other regions have relatively uniformly distributed HOGs. In other words, the first two histograms have larger variances than the last one.

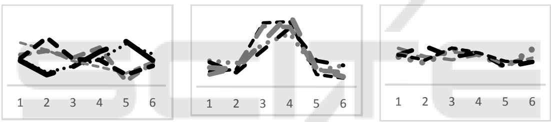

(a) (b) (c)

Figure 3: Examples of HOGs (a) and (b) HOGs distributions of sub-regions corresponding to the sidewalks and intersections, (c) HOGs distributions of other sub-regions. (row: HOG's bin).

4 CONTEXT RECOGNITION

In this module, the current context where a user stands is recognized based on the colors and textures. As mentioned above, the key elements for context recognition are the boundary orientations between sidewalks and roadways: horizontally oriented boundaries are found in images corresponding to intersections and more vertically oriented boundaries are observed in images corresponding to sidewalks. Accordingly, the boundary should be first discriminated from other natural lines; then, these pixels should be classified. In the proposed method, such classifications are accomplished using a multi-scale classification, which is illustrated in Fig 4.

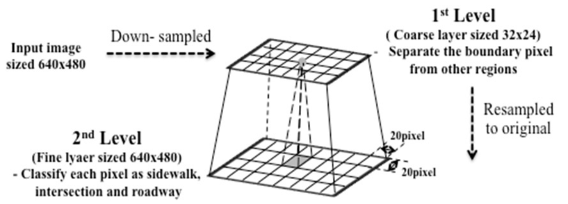

Figure 4: The multi-scale classification scheme.

The multi-scale classification applies the same operator to the input image while changing their scale and this method has been used in numerous applications including image segmentation and classification (Zoltan, 1996). The proposed multi-scale classification mechanism is composed of a coarse layer and a fine layer: the coarse layer discriminates the boundary pixels and the fine layer categorizes the pixels into sidewalk and intersection pixels. Then, in order to accomplish the different goals of the respective layers, the SVM is used as a classifier.

4.1 Locating the Boundary Pixel in the Coarse Layer

Generally, a boundary is detected using the edge operator and Hough transform. However, the boundaries detected using these methods include many sections that correspond to natural lines as well as boundary pixels, which increases the computation costs for the classification process. In the coarse layer, each boundary pixel is classified as a boundary pixel from other regions.

As shown in Fig. 4, the original image is first downscaled by 20×20, thus the average value of input pixels within a 20×20 block is assigned to the pixel in the down-sampled image. Then, for each pixel and its neighbors, the color and texture properties are extracted.

Thereafter, they input into the SVM-based classifier. The SVM was originally a binary classification algorithm developed by Vapnik et al. and it has been successively extended by numerous other researchers (Crammer, 2010).

Given a set of training examples that were classified as boundary or other, the SVM training algorithm builds a model that assigns new data into one category or the other. In this paper, the Gaussian radial basis function (RBF) was used according the experiment.

$$K(x_i, x_j) = exp\left(-r\|x_i - x_j\|^2\right), r > 0 \qquad (1)$$

For the experiment, the parameters were set as follows: gamma = 1 and cost = 1000.

Fig 5 shows the results of boundary extraction. For the input image in Fig 5(a), the result is shown in Fig 5(b), where the boundaries are denoted using the white color. As can be seen in the figure, the SVM training algorithm can successfully separate boundaries from other natural lines.

4.2 Determining the Context of the Boundary Pixel in the Fine Layer

In the fine layer, the original image sized at 640×480 is used. Accordingly, every boundary pixel expands to a block with a size of 20×20, each pixel of which has been categorized as intersection, sidewalk, or roadway by the SVM-based classifier.

As is well known, the SVM is used to classify objects into two class types. Accordingly, in order to to determine the current context from the three possible contexts, i.e. sidewalk, roadway, or intersection, two SVMs are used hierarchically, as follows:

$$O(x,y) = \begin{cases} 2 & \text{if} \quad SVM_s(x,y) > 0 \\ 1 & \text{else if} \quad SVM_i(x,y) > 0 \\ 0 & \text{else} \end{cases} \quad (2)$$

The first SVM classifies a pixel's class as sidewalk or other, and then the second SVM discriminates a pixel′s class as either an intersection or a roadway. Then, the same visual features that are used in the coarse layer are exploited.

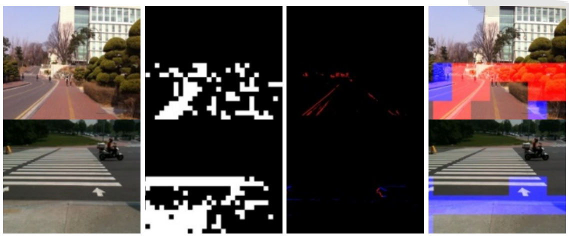

(a) (b) (c) (d)

Figure 5: Examples of the multi-scale classification results. (a) input image, (b) boundary extraction results, (c) context classification results, (d) final results after post-processing.

Fig 5 presents the outdoor context recognition results. For the input image in Fig. 5(a), the context classification result is shown in Fig. 5(c) where the red pixels and blue pixels are used to denote the pixel class as a sidewalk or intersection, respectively. As shown in Fig. 5(c), the classification results have some errors (noise) because the classification decision is performed locally on each pixel. Thus, smoothing is performed globally on the texture classification results in order to combine the individual decisions for a whole image. In order to solve this problem, a grid map was designed where each cell has a size of 80×80. The grid map is overlaid on the classification results, and then the classification decision for each cell is made. Then, the context of a cell is determined by the majority rule. For example, if most pixels in a cell are classified as intersections, the cell is also considered to be an intersection. Fig 5(d) shows the final classification result where the red block and blue block denote the sidewalk and intersection classes, respectively.

5 EVALUATION

In order to assess the practical validity of the proposed system, a number of real-world experiments were performed. The 30-minte walking experiments were asked to participants from 10:00 to 18:00 hours at six different locations. The smartphone was held by the participants standing at the streets. We recorded the captured by a phone, and counted the number of the respect context that appeared in the video as ground truth, which were used for objective evaluation of the proposed system.

Three females and two males participated in the evaluation and their average age was 25.6 years. In a day, all participants used mobile phone 6 hours on average and usually were used mobile phone on the road to sending text message, playing the mobile phone game and so on. In addition, three participants had the accident experience with person or poles that stay in front of the intersection while using a mobile phone on the road.

Each participant was given an introduction on how to operate the proposed system. Furthermore, the participants were asked to walk using the instructions provided by the proposed system. If a participant arrives at intersections on their way, they should stop and wait for the traffic signals. Then, we took participants to unfamiliar traffic intersections to test the system.

At the selected locations for real-world experiments, the accuracy of the proposed context recognition was evaluated. We compared the results with the ground truth. The average results are shown in Table 1. The overall accuracy was above 98%.

Table 1: The accuracy of proposed system.

Location	1	2	3	4	5	6	Total
Accuracy	100	98.4	100	94.4	100	96.7	98.25

Fig 6 some recognition results for various environments. Fig 6(a) shows the input images, where the images have time-varying lighting and the sidewalks have diverse patterns and colors. The input images were first enhanced in the preprocessing stage, and then the classifications were performed.

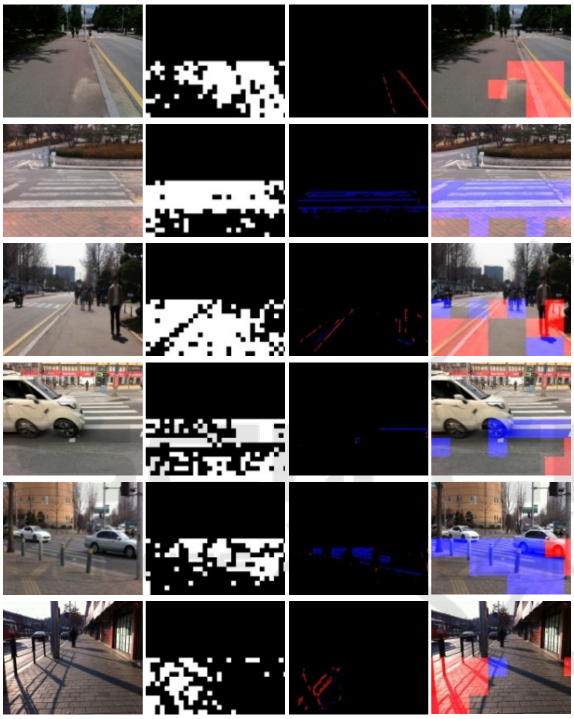

(a) (b) (c) (d)

Figure 6: Context recognition results. (a) Input image, (b) boundary extraction results, (c) classification results, (d) final results after post-processing.

As shown in Fig. 6(b), the boundaries between the sidewalks and roadways were correctly extracted, despite the diverse patterns on the sidewalks; however, they still included some false classifications. Fig. 6(c) shows the classification results, and Fig 6(d) shows the final results after post-processing. Although some cells were misclassified, the majority of the cells were classified as intersections; thus, its context was considered as an intersection. All examples recognized the correct context even though some pixels on the boundaries were misclassified. The results demonstrate that the proposed method has robust performance in the ground pattern and illumination type.

Fig 7 presents some examples of errors that occurred in complex scenes. As seen in Fig. 7, most boundary sections between the sidewalks and roadways were concealed by pedestrians and their shadows, so the boundaries in the fine layer classification could not be extracted (see the first row of Fig 7(b)). Thus, two images were misclassified as sidewalks. However, these errors can be easily resolved if the history of the place types between specified time frames are used.

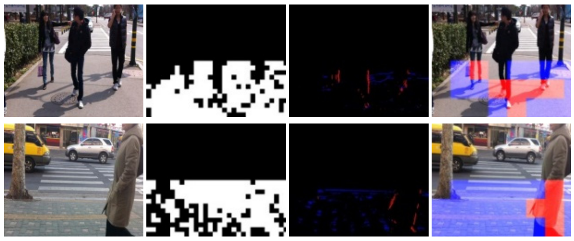

(a) (b) (c) (d)

Figure 7: Examples of misclassified images by moving walkers. (a) Input image, (b) boundary extraction results, (c) classification results, (d) final results after post-processing.

The primary purpose of the proposed system is to increase the safety of pedestrians who use their mobile phone while walking. For the system's practical use as an assistive device, real-time processing should be supported. The average processing time was approximately 304.16ms. As such, the proposed method can process more than four frames per second on resource-constrained smartphone. Consequently, the experiments demonstrated that the proposed method produces superior accuracy for context awareness, thereby assisting safer navigation for pedestrians in real-time.

In order to study the participants' satisfaction in using the proposed system, the participant satisfaction when using the proposed system was investigated. The following four questions were used to determine the participant satisfaction.

- E1: How helpful the system when you were walking?
- E2: How difficult was it to understand the system?
- E3: Is it comfortable to use the proposed system in the real context (How much it is comfortable)?
- E4: Do you want to use this system if it is developed as smart device app?

Fig 8 presents the results of the participant survey where the middle line represents the standard deviation. As seen in the figure, most participants were satisfied with the proposed system.

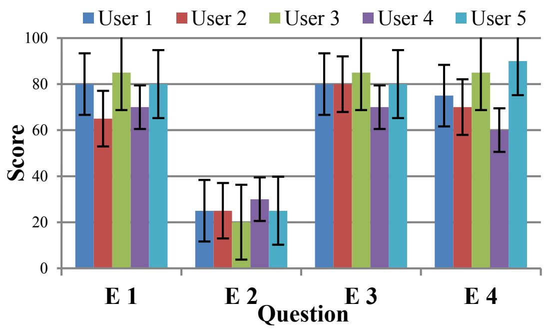

Figure 8: The user survey results.

In question E1, there was 76% satisfaction with the proposed system. Participant 5 stated, "I usually fall down while walking, but if I used this system, I could prevent a dangerous situation." Although the proposed system sometimes recognized incorrect information, the accuracy for detecting intersections was 96%, and participants could avoid most dangerous situations.

Question E2 referred to the ease of understanding and using the proposed system. The proposed system sent a warning alarm to the participant when they were approaching a traffic intersection and made them stop. Participant 1 said, "Because I was using the system for the first time, I was confused when I stopped. However, after practice I could use the system well without any effort." Most participants were slightly confused by the proposed system when they used it for the first time, but gradually their skill with and understanding of the proposed system improved with use without greater effort.

In question E3, all participants felt that it was convenient to use the proposed system. The participants received information that analyzed the current context using an audio or tactile interface. Sometimes they complained about the earphones, because they blocked other outside sounds. Participant 2 stated, "Usually I use earphones to listen to music, so I thought that the alert sound 'beep' was just noise."

In question E4, approximately 80% of the participants agreed that if the system were developed on a smart phone as multitasking, they would use it. Participant 5 said, "Actually, now I really need this system. I always use my smart phone while walking, and I don't see intersection quickly enough. So, if this system was developed as multitasking, I will download it." Consequently, in future research, the proposed system will be developed for multitasking smart devices.

Consequently, the experiments demonstrated that the proposed system produces superior accuracy for context recognition and that participants feel comfortable in using the proposed system.

6 CONCLUSION

In this study, an assistive device that increases the safety of the pedestrian, particularly for users who use their mobile phone while walking, was presented. The primary goal of the proposed system is to automatically recognize outdoor contexts such as intersections, sidewalks, and roadways, thereby reducing the number of traffic accidents involving pedestrians and vehicles. The proposed system is composed of three modules: pre-processing, feature extraction, and context recognition. In order to reduce the computational cost and improve the accuracy, a multi-scale classification was used where a coarse layer discriminated the boundary pixels and a fine layer categorized the pixels into sidewalks and intersections. In both layers, the support vector machine was used as the classifier.

In order to assess the validity of the proposed system, real-world experiments were conducted with five participants. The results demonstrated that the proposed method could recognize outdoor contexts with an accuracy of 98.25%, which proved the practical validity of the proposed system as mobility aids for pedestrians. To fully support the safety of pedestrians, obstacle detection and avoidance should be embedded into the current system. In future research, this will be extended to the current system.

ACKNOWLEDGEMENTS

This research was supported by Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Science, ICT & Future Planning(NRF-2013R1A1A3013445).

REFERENCES

Institute of Transportation Engineers U.S. Department of Transportation Federal Highway Administration,

2004, Pedestrian safety at intersection, *PEDESTRIANS*.
Kate seth, 2010, Walk on the safe side, *Traffic Technology International*.
Leena, Rauno, 2012, Accidents and close call situations connected to the use of mobile phones, *Accident Analysis & Prevention.*
Mark B et al., 2009, *Pedestrians, vehicles, and cell phones*, Accident Analysis & Prevention.
Robert et al., 2010, Association Between Roadway Intersection Characteristics and Pedestrian Crash Risk in Alameda County, California, *Transportation Research Record: Journal of the Transportation Research Board.*
Tianyu et al., 2012, WalkSafe: a pedestrian safety app for mobile phone users who walk and talk while crossing roads, *ACM HotMobile.*
Jack et al., 2007, Mobile telephones, distracted attention, and pedestrian safety, *Accident Analysis & Prevention.*
Julie, Susanne, 2006, The effects of mobile phone use on pedestrian crossing behaviour at signalized and unsignalized intersections, *Accident Analysis & Prevention*.
Brabyn J et al., 1933, Talking signs: a remote signage, solution for the blind, visually impaired and reading disabled, *Engineering in Medicine and Biology Society*.
Barlow J et al., 2003, Accessible pedestrian signals: Synthesis and guide to best practice, *National Cooperative Highway Research Program.*
Dragan et al., 2011, Zebralocalizer: identification and localization of pedestrian crossings, *ACM MobileHCI.*
Karacs K et al., 2006, Bionic Eyeglass: an Audio Guide for Visually Impaired. Biomedical Circuits and Systems, *BioCAS.*
Volodymyr et al., 2008, Crosswatch: a Camera Phone System for Orienting Visually Impaired Pedestrians at Traffic Intersections, *International Conference on Computers Helping People with Special Needs.*
Chen-Fu Liao, 2012, Using a Smartphone App to Assist the Visually Impaired at Signalized, *University of Minnesota.*
James M et al., 2013, Crosswatch: a System for Providing Guidance to Visually Impaired Travelers at Traffic Intersections, *Journal of Assistive Technologies.*
Barlow J et al., Tabor, 2003, Accessible pedestrian signals: Synthesis and guide to best practice, *National Cooperative Highway Research Program.*
Brabyn J et al., 1933, Talking signs: a remote signage, solution for the blind, visually impaired and reading disabled, *Engineering in Medicine and Biology Society*.
Dragan et al., 2011, Zebralocalizer: identification and localization of pedestrian crossings, *ACM MobileHCI.*
Ophir et al., 2009, Cognitive control in media multitaskers, *Proc. of the National Academy of Sciences.*
Sivaraman, Trivedi, 2010, A general active-learning framework for on-road vehicle recognition and tracking, *IEEE Transactions on Intelligent Transportation Systems.*
Zoltan et al., 1996, A hierarchical Markov random field model and multi-temperature annealing for parallel image classification, *GRAPHICAL MODELS AND IMAGE PROCESSING.*
Crammer et al., 2010, On the Algorithmic Implementation of Multiclass Kernel-based Vector Machines, *J. of Machine Learning Research.*