The Cyber-Physical Bike: A Step Towards Safer Green Transportation

Stephen Smaldone, Chetan Tonde, Vancheswaran K. Ananthanarayanan
Ahmed Elgammal, and Liviu Iftode
Department of Computer Science
Rutgers University

ABSTRACT

To improve road cycling safety, we present an approach that augments bicycles with video processing and computational capabilities. This Cyber-Physical bicycle system continuously monitors the environment behind the biker, automatically detects rear-approaching vehicles, and alerts the biker prior to the approach. In this paper, we *(i)* identify biker safety as a problem that can be addressed using mobile video sensing and processing, *(ii)* present the design of a Cyber-Physical bicycle system, which applies video processing techniques to perform automated vehicle detection, and *(iii)* demonstrate the feasibility of this system through the evaluation of our prototype implementation. Early results show that our approach operates with good accuracy at normal frame rates, and can perform detection in real time with reduced frame rates.

1. INTRODUCTION

Since their invention in 1817 [2], bicycles have proven to be a healthy and environmentally friendly mode of transportation for both enthusiasts and commuters alike. Although the bicycle has remained ubiquitous over time, the world has changed dramatically. Today, U.S. roadways are predominantly utilized by automobiles, which are convenient yet energy inefficient modes of transport. Unfortunately, these days, cyclists are considered second-class citizens by some motorists as they attempt to share roadways [29]. In fact, this has been the situation for most of the lifetime of the bicycle. According to a 1996 U.S. Department of Transportation report, one of the first automobile accidents ever to occur in the U.S. took place in New York City, during 1896, between an automobile and a bicycle, and proved fatal for the cyclist [3]. According to a more recent report (2007) in the U.S., over 700 bicyclists die annually in accidents with automobiles, while there are over 44,000 annually reported cases of injuries due to bicycle-automobile accidents [4].

A key concern for bikers is safety on roadways. This is primarily due to the unbalanced safety situation that exists during cyclist-motorist encounters. Cyclists are clearly more vulnerable to injury or fatality during unsafe encounters. Unfortunately, existing laws and approaches towards biker safety (e.g. bike paths) are only adequate. At best, laws proscribe remedies for incidents after-the-fact, and bicycle paths are only as good as the limited coverage they provide. Under most scenarios, a biker is forced to focus a substantial portion of her cognitive and physical capabilities on the task of maintaining situational-awareness, by continuously probing for the occurrence of rear-approaching vehicles. What is required is a preventative, biker-centric solution to the problem.

In this paper, we present an approach that addresses the problem by allowing the bicycle to maintain situational awareness for the cyclist. To support this approach, we enhance a standard bicycle with sensing and computational capabilities creating a Cyber-Physical bicycle system. The core goal of this system is to provide accurate and timely detection of rear-approaching vehicles alerting the biker of the pending encounter, through the cross-cutting application of mobile computing and novel computer vision techniques. To the best of our knowledge, our system is the first to extend pervasive computing concepts to bicycles by equipping them with computational capabilities for the purpose of improving road cycling safety. It is our position that this is an important problem, as we believe that roadway safety is a key limiter in the adoption of bicycling as a viable form of environmentally friendly transportation.

This paper makes the following contributions:

- It identifies biker safety as a problem that can be solved using mobile computing and computer vision techniques.
- It describes the design of an automated real-time safety system for bikers. Our approach augments a bicycle with a camera and computational capabilities, employs video processing techniques to detect when a motor vehicle is approaching from behind, and alerts the biker prior to each encounter.
- It demonstrates the feasibility of this system through experiments on our prototype implementation using real biker traces, which we collected. Our preliminary experimental results show that the system exhibits a good degree of accuracy at 30 frames per second. They also demonstrate that the system can operate within the real-time constraints of the problem scenario, at a reduced frame rate.

2. THE STATE OF BICYCLE SAFETY

In this section, we review the present state of biker safety. In short, a quick review of biker fatality and injury statistics shows consistent non-improvement [3, 4].

The Problem. Since bikers typically ride with the flow of traffic, a commonly occurring dangerous situation for them is when they are being passed from behind by a motor vehicle. To predict the occurrence of these situations, a biker must spend a substantial amount of her cognitive and physical ability to periodically scan for rear-approaching vehicles, reducing her capacity to handle the bicycle safely and maintain continual awareness for the forward situation.

In an attempt to reduce the cognitive and physical effects of looking back, many bikers employ either handlebar or helmet mounted rear-view mirrors. In either form, they distract attention even in the absence of cars, and have a limited and inconsistent range of view. More recently, products such as Cerevellum's [7] digital rear-view mirror attempt to provide a continuous video-based view of the situation behind the biker. Unfortunately, it does not automatically detect approaching motor vehicles, so it does little in the way of reducing the cognitive load on the biker.

HotMobile'11 March 1-2, 2011, Phoenix, Arizona, USA

Legal Solutions. Laws have been enacted in most states to force children to wear helmets while riding bikes, but they do not apply to adult bikers. Statistics show that the average age for cyclist fatalities is 40 years and for cyclist injuries is 30 years [4]. Although a helmet may effectively protect a cyclist's head for simple falls and lower energy accidents, it is unclear what protection they provide during higher energy accidents involving a motor vehicle, and they offer no protection to the rest of a cyclist's body. More recently, some states have even enacted laws to impose a three foot safe passing limit on motorists passing bikers [5].

Unfortunately, laws do little to *prevent* accidents due to insufficient enforcement and can only proscribe after-the-fact remedies to accidents. According to statistical studies, three out of four at-fault drivers are not even cited for hitting and killing cyclists, and 22% of fatal accidents involved "hit-and-run" drivers, who were never even found or charged. For example, in New York City, of the 92% of drivers who were at-fault for killing a cyclist, 74% did not even receive a traffic citation [26]. In short, laws may help penalize offenders, but are difficult to enforce.

Infrastructure Solutions. Certain cities, for example, Portland, OR, have built extensive networks of bicycle lanes [1] to promote safe cycling. Unfortunately, bicycle lanes are only as good as the coverage they provide, require a strong public commitment to install and maintain, and enforced legislation to ensure they are not improperly utilized (e.g., illegally parked cars).

Other Solutions. Various ways exist to make the biker more conspicuous to motorists. These include bike reflectors, flashing lights, reflective clothing, and even laser-based solutions [9]. Although they represent preventative measures, they do little to warn an unprepared biker to the presence of an approaching vehicle.

3. THE CYBER-PHYSICAL BICYCLE

In this section, we describe the design of a Cyber-Physical bicycle system, which utilizes video processing to automatically detect rear-approaching vehicles. The goals of our design are threefold. First, the system must accurately detect and track vehicles that approach from behind. Second, the system must be able to alert a biker with enough time that the biker can take evasive action. Third, the system should distinguish between vehicles that approach a biker in both safe and unsafe manners. There are a number of significant challenges to overcome when building a system to automatically detect rear-approaching vehicles. They are:

- *Limited Resources.* Bicycles have limited available power generation capabilities and obvious weight restrictions, which place a serious limitation on the computational resources that a bicycle can be equipped to carry.
- *Platform Instability.* A moving bicycle is subject to vibrational motion due to roadway conditions and rapid changes in direction enacted by the biker. The resulting video stream obtained from a bicycle-mounted camera is subject to large amounts of jitter and unpredictable changes in orientation.
- *Approaching Vehicle Directionality.* Distinguishing between rear and front approaching vehicles is critical, as an alert could be generated by the system for all vehicle encounters, reducing the effectiveness of the system.

3.1 Cyber-Physical Bicycle Overview

Figure 1 illustrates the key components of the Cyber-Physical bicycle and how it improves the safety of bikers. The central element is a normal bicycle augmented with video sensing and processing capabilities, and compute resources in the form of a bicycle computer. As shown in the figure, a camera faces backward from the bicycle's direction of forward motion, collects video data, and streams it to an embedded bicycle computer, which continuously processes the data utilizing computer vision techniques to perform rear-approaching motor vehicle detection.

As a biker rides along a roadway, her bicycle maintains situational awareness for her. When a motor vehicle approaches the biker from behind, the occurrence is detected by the bicycle computer and a notification is raised to the biker. Even after a notification has been raised, the system continues to track the approaching vehicle to determine the level of threat posed to the biker. To ascertain this, we define a *virtual safety zone* , around the biker (which is three feet in our case – see Figure 3).

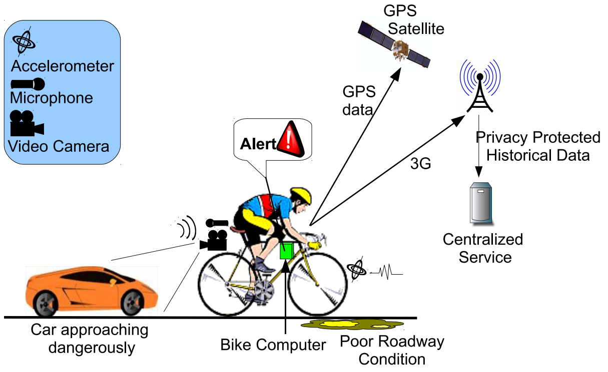

Figure 1: Cyber-Physical Bicycle

Whenever a vehicle crosses this perimeter, the Cyber-Physical bicycle detects this, considers it to be an unsafe situation for the cyclist, and produces a notification warning distinct from the early notification produced when the vehicle is first detected.

3.2 Video-Based Detection and Tracking

In this subsection, we describe the design of the video-based automated detection and tracking subsystem. Since approaching vehicles travel at relatively high speeds compared to bicycles, it is necessary to detect them early and accurately. Although we would like to leverage existing automobile driver assistance systems, which can detect the presence of vehicles in a driver's blind spot, we cannot. Several sensing technologies have been used in automotive assistant systems. We summarize them and others in Table 1.

To provide ample notification time to a biker, the Cyber-Physical bicycle system must detect rear-approaching vehicles while they are still very far. Once detection has occurred, the video-based subsystem performs real time tracking of these vehicles, to determine if each will pass the biker in either a safe or unsafe manner.

We base our design of the video-detection system on two observations. First, as a bicycle proceeds, all stationary objects behind it will appear to recede. Therefore, vehicles approaching from the rear will move counter to the other objects in the field-of-view (FOV). This is leveraged in the *Optical Flow Analysis* component.

Our second observation is that roadway traffic follows a predictable pattern. Since bikers ride with the flow of traffic, a rear-approaching vehicle occupies the same roadway lane as the biker. By identifying the natural segmentation of a roadway, we can reduce the image area that must be analyzed, focusing on areas of the FOV where a rear-approaching vehicle is likely to appear. We leverage this in the *Roadway Segmentation Analysis* component.

Optical Flow Analysis. Optical flow (OF) is the pattern of apparent motion of objects in a scene caused by the relative motion between an observer and the objects [21]. We use this to distinguish an approaching vehicle from all other objects in the FOV.

Figure 2 shows an example of optical flow computed from images captured by a rear-facing camera mounted on the back of a moving bicycle. In the figure, *(a)* and *(b)* depict a rear-approaching vehicle and its optical flow, while *(c)* and *(d)* depict the same for a rear-departing vehicle. Red pixels are moving towards the biker, while blue are moving away. A stronger color (i.e., more red or blue) indicates a faster relative speed. We observe that an approaching car exhibits a very distinct relative motion pattern compared to the rest of the scene (dominant red spot in the middle).

Technology	Examples	Description
Active Sensors	Laser, Lidar, Millimeter-Wave Radars	Capture the distance to objects based on time-of-flight of emitted. Not suitable for bikes since they require a lot of power to operate; they provide low spatio-temporal resolution and low scanning speed signals.
Metal Detectors	Honeywell HMC1022	Sense fluctuations in ambient magnet fields to determine when a metal body is in close proximity to the sensor. Very good for detecting close proximity, and functions for most environmental conditions, but does not provide necessary range or distance accuracy for timely vehicle encounter prediction.
Structured Light Cameras	Microsoft Xbox Kinect	Directly capture 3D depth information in real-time (e.g. for gaming applications). These sensors depend on reflecting structured light on objects, which is not suitable in our case because the structured light is diffused at the vehicle surface, besides interference with headlights, and they have limited range (based on our preliminary experimentation).
Optical Sensors	Video, IR, or Stereo Cameras	Provide very detailed information at sufficiently high speed (i.e. frame rate in this case) with lower power consumption, but requires more complexity in the computation for real-time detection determination, and may not be suitable for all environmental or spatial conditions.

Table 1: Vehicle Detection Technology

(a) Approaching (b) Approaching OF

(c) Departing (d) Departing OF

Figure 2: Optical Flow (OF) Processing

Roadway Segmentation Analysis. To reduce computational load, the video-based detector segments the roadway image along the existing visible natural boundaries based on its color, using a Gaussian Mixture Model in Hue Saturation Value (HSV) color space. To facilitate this, a fixed region in the bottom part of the image, which always contains the road, is chosen to train and update the model for every frame. Then, this roadway color model is utilized to segment the rest of the scene. Using the edge features of the image, we perform a Hough Transform to determine line-like features and lane markers to identify the vanishing point of the scene.

Vehicle Tracking. The system tracks detected approaching vehicles until they pass the biker. Figure 3 shows vehicle tracking for approaching vehicles (red box on vehicles). We use an appearance-based approach for tracking objects because it is more robust to noise from camera jitter. Our approach is based on the method proposed in Ross et al. [33], which uses Principle Component Analysis based appearance models to track the detected cars in affine subspace. We extend their algorithm to make it adaptive so that it can handle the rapid changes in pose, scale, and illumination (caused by the strong perspective effect), which commonly occur when tracking an approaching vehicle. Additionally, our algorithm can track multiple targets simultaneously by associating new observations with existing trackers through linear assignment [23] and simple tracker management using the location and appearance count.

As a tracked vehicle approaches a biker, the system continuously calculates the distance between them to determine if the vehicle crosses into the biker's virtual safety zone. In all cases, the biker is alerted to the presence of the approaching vehicle. This is visually demonstrated by the two images shown in Figure 3, with the yellow and white grid delineating the virtual safety zone. Figure 3(a) shows the safe car pass scenario. Since the vehicle in the image is passing the biker outside of the virtual safety zone, a yellow warning is signaled in the image (yellow box in the upper right corner). Alternately, in Figure 3(b), we see an example of an unsafe car pass. In this image, the car crosses the safety zone boundary and a red alert is visibly raised by the system.

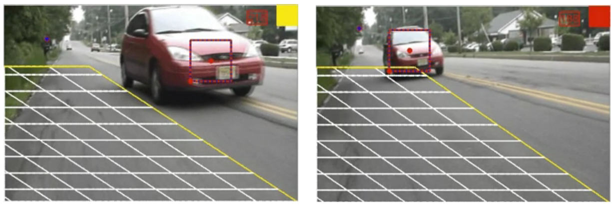

(a) Safe Approach (b) Unsafe Approach

Figure 3: Virtual Safety Zone

4. IMPLEMENTATION AND EVALUATION

We are developing a prototype Cyber-Physical bicycle system in C/C++ using the NVIDIA CUDA library (v2.3) and an open source optical flow library [35]. It runs on an HP Mini 311 netbook with an Intel Atom N280 1.67 GHz CPU, 3 GB RAM, NVIDIA ION GPU (16 CUDA cores and 256 MB memory), and an 80 GB SSD hard disk. We choose this as our development platform because the hardware is light-weight (3.26 lbs) and closely matches embedded hardware utilized in various multimedia applications. We consider the weight of this system to be a worst case because it includes an LCD display, keyboard, and other embedded peripherals that would be left out of an optimized form factor. Even so, it represents a less than 10% weight increase for the average commuter bicycle, which we believe to be in line with user requirements.

Our preliminary experimental evaluation of the Cyber-Physical bicycle system prototype addresses two questions:

- How accurate are video-based techniques in detecting rear-approaching vehicles? (Section 4.1)
- Can detection be performed in real-time? (Section 4.2)
- What are the power requirements of the system? (Section 4.3)

Experimental Setup. All experiments are executed using our prototype implementation on the hardware specified in Section 3. To support repeatability, we collected over 3 hours of real-world roadway cycling traces (over 10 GB). To gather these traces, we mounted a rear-facing digital video recorder (Sony Handicam DCR-SX40) to an ordinary road bicycle (Trek FX7.5). The camera collects video recordings, while a biker rides the bicycle along a set of roadway bike routes in central New Jersey. We manually annotated every interaction between the biker and a vehicle (approaching and departing) using the timestamps of each individual trace.

4.1 Video Detection Accuracy

In this section, we evaluate the accuracy of the video-based detector. Results are presented as a confusion matrix in Table 2. True positives represent the cases where a rear-approaching vehicle is correctly identified. True negatives represent the cases when the absence of a rear-approaching vehicle is correctly detected. False positives occur when something is mis-classified as a rear-approaching vehicle. A false negative occurs when a rear-approaching vehicle was not detected. Finally, we define accuracy as: $Accuracy = \frac{TP}{TP+FP+FN}$.

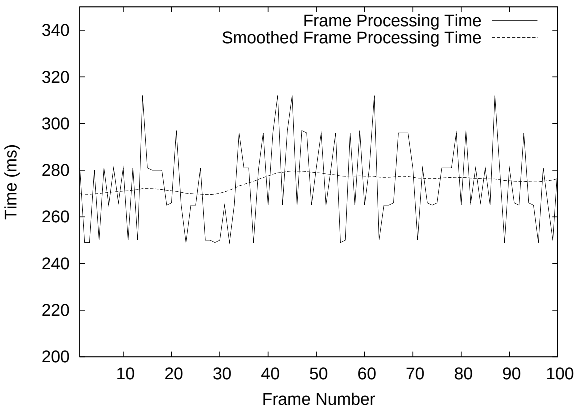

Figure 4: Optical Flow Performance.

From Table 2, we observe that in all tested cases only one rear-approaching vehicle was not detected (FN), while six false alerts were raised (FP) . Based upon these results, we calculate the overall accuracy of this method to be 73.1%.

	Positives	Negatives
True	19	N/A
False	6	1

Table 2: Detection Accuracy

To better understand the sources of incorrect classification, we reviewed the specific trace sequences that were incorrectly handled by the system. In one of the instances, a biker is being followed by another rider, who occludes a rear-approaching vehicle. This causes a false negative. A second example illustrates a false positive case, which occurs due to a sudden vibration of the system caused by a section of uneven roadway surface. The fast lateral motion of the camera causes the location of the car to shift by hundreds of pixels, confusing the tracker. From the accuracy results, and our manual inspection of the sources of errors, we conclude that the Cyber-Physical Bicycle system is feasible, exhibiting good accuracy, but there is still work to perform to improve the situation.

4.2 Real-Time Performance

In this section, we evaluate performance and timeliness of video-based detection. To better understand the contributing factors, we first measure the performance of the individual components of the video-based detector pipeline, then we measure the overall performance. Two components form the critical performance path: optical flow processing and vehicle tracking. We define timeliness as the average number of seconds warning the system provides to the biker prior to a vehicle encounter. We present the *potential timeliness* as the percentage of total possible time when a vehicle could be detected by the system, i.e., the difference between the first appearance of the vehicle and the time it passes the cyclist.

Optical Flow Performance. The first video performance experiment measures the latency of processing a frame using optical flow techniques. Figure 4 presents the results of this experiment. In the figure, the results are presented as the frame processing latency (in ms) for the sequence of frames from a portion of the roadway cycling traces. We observe that each frame in the experiment is processed in real-time. Although the latency for frame processing fluctuates between 250 to 315 ms, the frame rate never drops below 3 frames per second (FPS). We also observe, by focusing on the smoothed data (dashed line), that optical flow imposes a relatively constant and predictable processing cost.

Vehicle Tracker Performance. The second experiment examines the processing cost of vehicle tracking. The results are presented in Figure 5. The figure presents the instantaneous FPS rate as each frame is processed in the Vehicle Tracking component. We observe that the frame rate varies between 3.7 and 1.3 FPS. We also observe that performance is quite stable around those two values. The reason for this is that when the component is tracking vehicles, frame processing is more expensive, and we experience a subsequent drop in FPS rate. Two examples of this can be observed in the figure for frames 150-175 and again for frames 225-260.

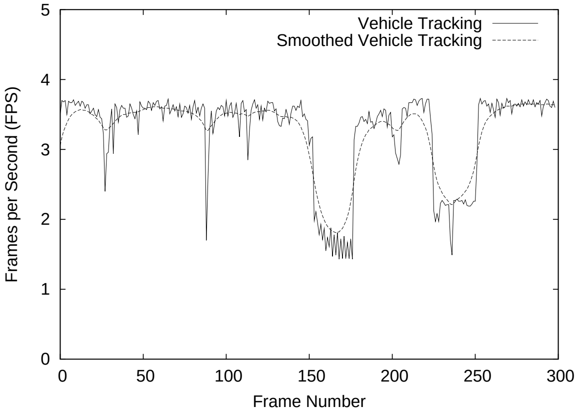

Figure 5: Vehicle Tracking Performance.

Component	Power Consumption Rate (Watts)
Idle	6.4
Cyber-Physical Bicycle	14.2

Table 3: Power Consumption Rates.

End-to-End Performance. Finally, we measure the overall performance (average FPS) of the optical flow and vehicle tracking, when executing our traces through the entire video detection pipeline. In this experiment, the optical flow component achieved an average frame rate of 3.6 FPS, while the vehicle tracker component achieved 3.3 FPS. When compared to the individual cases, both components exhibited slightly reduced average performance due to resource competition under concurrent processing conditions. Under these conditions, we measure the real-time video-based alerting to occur an average of 3.5 seconds prior to a vehicle encounter, which is 92% of the potential time. We conclude from the performance results that the current prototype system can execute the rear-approaching vehicle detection task within the time constraints of the problem, but at a reduced frame rate. This exposes a trade-off in the system between accuracy and timeliness. We intend to explore this trade-off as part of our continuing work on this project.

4.3 Power Requirements

To understand the energy burden the Cyber-Physical bicycle system imposes, we measure the absolute power rates compared to a baseline idle system. To capture the measurements, we utilize the BatteryBar utility. Table 3 presents the results of these measurements (in Watts). From the table, we observe that the video pipeline (Video) nearly doubles the power requirements over the idle system due to the inclusion of GPU processing. In fact, the video pipeline results are comparable to the requirements for a typical movie player, and we expect the battery lifetime while executing our system to be comparable to that of watching a movie (measured power requirements for the VLC Movie Player are 12.4 Watts). Based on these measurements, and assuming a linear battery discharge rate, we estimate the battery lifetime of our system to be approximately 5 hours. Of course, battery discharge rates are non-linear, but that is not likely to change the fact that the estimated lifetime matches on the order of typical roadway bicycle ride durations [6, 10]. Moreover, since a common application of most modern netbook users is to watch movies (operating on battery power), experience suggests that a typical netbook battery will support the Cyber-Physical bicycle system for the required duration.

5. FUTURE DIRECTIONS

In this section, we briefly describe a number of additional directions which we are presently exploring as a part our ongoing research project, as applied to the cycling domain. We also intend to broaden our research to include applications in the motorcycle and motor-scooter domains, since they share similar safety issues, while presenting a different set of challenges.

Audio-Based Detection and Integration. One potential method to improve the detection accuracy of the Cyber-Physical bicycle is to add additional sensing modalities. In particular, we believe that audio-based sensing provides a promising approach to augment our existing system. To inform audio-based automated detection subsystem, we make two observations. First, vehicles are clearly audible to a biker over the background wind noise, when she turns her head. Therefore, a wind shielded audio sensor should be able to detect vehicular sound. Second, since sound from a vehicle is directional, the system should be able to discriminate between rear and front-approaching vehicles. Furthermore, a vehicle that approaches from the rear will cause a longer sound because it will take more time to pass a biker than one approaching from the front.

To further improve the accuracy of detection, the Cyber-Physical bicycle can combine the results of both video and audio-based detection because both modalities can operate in parallel. This combination may allow the system to perform an additional level of comparison based upon the individual audio and video prediction results to further filter false positives and catch false negatives by leveraging the diverse characteristics of the different modalities.

In addition to higher-level combining, the system can also share predictions between modalities, as a form of cross-modal feedback. This will allow the introduction of adaptivity into the system. For example, under conditions of low visibility, the system could reduce the frame rate for video and rely more on audio detection, potentially optimizing both power (by reducing resource requirements) and accuracy (by reducing video false positives).

Roadway Hazard Detection. There are numerous highly-visible obstacles, which may appear along a biker's path such as tree branches or parked cars. Such obstacles are typically observable by a forward-facing biker and do not require any Cyber-Physical aid. However, there are less observable obstacles, which pose a more serious hazard to biker safety. For example, if a biker happens to be riding over a cluster of potholes or sandy patch of road while being passed by a vehicle, there is a greater risk of accident due to increased chances that the biker may lose control of her bicycle. Individually, neither the roadway condition nor the slow-moving vehicle poses a threat, but together they place aggravated risk on the biker. We intend to investigate the possibility of performing roadway surface sensing to automatically detect unsafe conditions and plan to use a combination of accelerometer, video, and audio-based sensing to collect and categorize the motion of a bicycle as it travels along the roadway. Accurate roadway sensing from a moving motor vehicle has proven difficult, yet tractable [18]. Compared to motor vehicles, bicycles pose additional difficulties due to their light weight and susceptibility to slight perturbations in motion.

CyberPeloton (Platoons of Bicycles). Typically, bikers ride in groups called *pelotons*. As bikers collect into groups, there is an opportunity for the Cyber-Physical bicycle system to take advantage of this proximity to improve safety, share processing load, and provide social functionality . For example, as a vehicle approaches from behind, the last bicycle in the group can pass detection alerts forward to other bikers in the peloton (using wireless communication). Similarly, a bicycle in front may perform roadway hazard detection and pass alerts back. Achieving this requires functionality to automate the formation of CyberPelotons whenever bikers are in close proximity to each other as well as support for intra-group signaling including high-priority messages.

CyberPeloton formation is challenging because Cyber-Physical bicycles must: *(i)* perform real-time proximity detection to determine when there are opportunities to form groups, *(ii)* perform fine-grained relative positioning[1], *(iii)* maintain group consensus to ensure that each member agrees that it is a member and has an accurate view of the group arrangement, and *(iv)* agree upon a load sharing schedule within the group. For example, in order to elect a member to perform roadway hazard alerting, the group must agree on the member that is at the front-most position in the group, and adapt with relative changes in position.

Automated Incident Detection. Besides accident avoidance, it is also important to consider biker safety from an accident response perspective. Except for the most fortuitous situations where no injury occurs, an accident that involves a bicycle and a motor vehicle will likely require immediate medical attention for the biker. To handle this, we plan to automate the detection of such situations and react accordingly. Additionally, since prosecuting hit-and-run motorists is difficult due to the inherent lack of actionable evidence (e.g., license plate number, make and model of car, etc.), we have the complimentary goal of gathering such evidence.

Since the Cyber-Physical bicycle system is equipped with wireless communication technology, notifying authorities in most situations is straight-forward, but we must also consider exceptional cases when traditional wireless communication fails (e.g., poor cellular signal strength). To handle this, we plan to investigate alternate emergency signaling methods, as part of this work. Finally, once the authorities have been contacted, the correct information must be communicated in a data sensitive manner that preserves the authenticity and integrity of the evidence.

Safe Route Planning. A proactive measure to accident avoidance is to incorporate safety metrics directly in route planning. The goal is to allow users to map out potential bike routes and then have the system quantify the safety of their route and suggest alternate paths based upon safety criteria. This would allow a user to make a direct quantitative comparison of the relative safety between different bike routes. Today, this can only be determined through biker experience, documented opinion, and anecdotal evidence.

None of the cycling-oriented web services [6, 10, 32] include safety as an automatically collected quantitative property for use in route planning. Cyclists may directly share qualitative data with each other regarding roadway safety and route planning. While one of the services, Cyclopath [32], attempts to perform quantitative analysis using qualitative user submissions, it can not provide the level of analysis that could be achieved from sensed data.

Power Generation. We have assumed that all power is supplied by the prototype platform's integrated battery. As we progress, we plan to consider possible sources of "green" power, to recharge the battery while riding. For example, specialized bicycle wheel hubs can capture power from a biker's pedaling motion [11], or from deceleration due to regenerative braking [30]. Such hubs can generate up to 6 Watts of continuous power, and are commonly used by bikers to power headlamps. Additionally, as a bicycle moves between 13 and 20 mph, an opposing wind force is generated, which can be used by small, light-weight, attachable, wind turbines [8] to generate power. Finally, since bicycling is an outdoors activity, solar energy is generally available.

6. RELATED WORK

The closest work to ours is the BikeNet [16] project, which utilized various sensors to collect environmental data for the purpose of improving biker fitness. Although this project was the first to apply sensors to bicycles, its application domain was biker fitness, whereas our goal is the equally challenging problem of biker safety. Similarly, the Copenhagen Wheel [30] project aims to add sensing and power-assist technology to bicycles to promote green transportation while enabling mobile urban sensing applications.

---

[1] Bikers frequently ride within a few feet of each other, well within the error range of common GPS.

Computer vision technologies have been successfully used to detect moving objects, such as people and vehicles, in many application domains. In particular, computer vision systems have been developed to assist automobile drivers [13, 14, 15]. For example, systems have been developed and deployed to assist automobile drivers to detect approaching cars in blind spots using cameras and radar sensors [31, 34], to detect and recognize traffic signs [19], crossing pedestrians [17, 20], to detect lanes and vehicles ahead in the same lane [12, 24, 25, 27, 36], and also for autonomous driving [22]. In the DARPA 2005 grand challenge, several team competed on developing autonomous vehicles to drive on 211 km desert roads [28]. Despite the effort spent in integrating vision systems in automobiles, almost nothing has been done to develop vision systems for bicycles. This is due to the challenges of developing accurate, light-weight, and power efficient vision systems under the severe computing constraints imposed by a bicycle. Basically, the detection problem from sensors mounted on bikes is much harder than that from cars and the computational constraints are tighter.

7. CONCLUSION

In this paper, we introduced the Cyber-Physical bicycle, a system that augments bicycles with video processing capabilities to perform automated rear-approaching vehicle detection. The purpose of this system is to directly improve the safety of bikers as they travel on roadways by reducing the cognitive overheads presently associated with roadway cycling. To the best of our knowledge, our system is the first *(i)* to identify biker safety as a problem that can be solved using mobile sensing and processing and *(ii)* to apply video processing techniques to perform automated vehicle detection from a bicycle. Finally, the preliminary results of an experimental evaluation of our prototype suggests the feasibility of this system, which is able to perform rear-vehicle detection from a bicycle with good accuracy at full input data rate, and can operate in real-time under conditions of reduced input data rate.

8. ACKNOWLEDGMENTS

We thank our shepherd Maria Ebling and the anonymous reviewers for their valuable comments and suggestions.

9. REFERENCES

[1] 11 Most Bike Friendly Cities in the World. `http://www.virgin-vacations.com/site_vv/11-most-bike-friendly-cities.asp`.

[2] A Quick History of Bicycles. `http://www.pedalinghistory.com/PHhistory.html`.

[3] Traffic Safety Facts 1996: Pedalcyclists. NHTSA Report, U.S. DOT, 1996.

[4] Traffic Safety Facts 2007: Bicyclists and Other Cyclists. NHTSA Report, U.S. DOT, 2007.

[5] 15 states require 3-foot clearance for bicycles - Biking Bis - Bicycle Touring and More web site., 2009. `http://www.bikingbis.com/blog/_archives/2008/3/5/3549263.html`.

[6] Bikely web site., 2010. `http://www.bikely.com`.

[7] Cerevellum web site., 2010. `http://www.cerevellum.com`.

[8] HYmini | miniSOLAR - Green Power is in your Hand - MINIWIZ S.E.Dev Ltd. web site., 2010. `http://www.hymini.com`.

[9] LighLane web site., 2010. `http://www.lightlanebike.com`.

[10] MapMyRide web site., 2010. `http://www.mapmyride.com`.

[11] Shimano DH-3N70 Dynohub - Peter White Cycles web site., 2010. `http://www.peterwhitecycles.com/shimano3n70.asp`.

[12] AMDITIS, A., BIMPAS, M., THOMAIDIS, G., TSOGAS, M., NETTO, M., MAMMAR, S., BEUTNER, A., MÖHLER, N., WIRTHGEN, T., ZIPSER, S., ETEMAD, A., DA LIO, M., AND CICILLONI, R. A situation-adaptive lane-keeping support system: overview of the safelane approach. *IEEE Transactions on Intelligent Transportation Systems 11* (September 2010), 617–629.

[13] BERTOZZI, M., BROGGI, A., CELLARIO, M., FASCIOLI, A., LOMBARDI, P., AND PORTA, M. Artificial vision in road vehicles. *Proceedings of the IEEE - Special issue on "Technology and Tools for Visual Perception" 90* (July 2002), 1258–1271.

[14] BISHOP, R. Intelligent vehicle applications worldwide. *IEEE Intelligent Systems 15* (January 2000), 78–81.

[15] BROGGI, A., IKEUCHI, K., AND THORPE, C. E. Special issue on vision applications and technology for intelligent vehicles: part i-infrastructure. *IEEE Transactions on Intelligent Transportation Systems 1* (August 2000), 69–71.

[16] EISENMAN, S. B., MILUZZO, E., LANE, N. D., PETERSON, R. A., AHN, G.-S., AND CAMPBELL, A. T. The bikenet mobile sensing system for cyclist experience mapping. In *SenSys '07* (2007).

[17] ENZWEILER, M., AND GAVRILA, D. Monocular pedestrian detection: Survey and experiments. 2179–2195.

[18] ERIKSSON, J., GIROD, L., HULL, B., NEWTON, R., MADDEN, S., AND BALAKRISHNAN, H. The pothole patrol: Using a mobile sensor network for road surface monitoring. In *Proceeding of the 6th international conference on Mobile systems, applications, and services* (2008), MobiSys '08.

[19] GAVRILA, D. Traffic sign recognition revisited. In *Mustererkennung 1999, 21. DAGM-Symposium* (1999).

[20] GAVRILA, D., FRANKE, U., WOHLER, C., AND GORZIG, S. Real time vision for intelligent vehicles. *IEEE Instrumentation Measurement Magazine 4* (June 2001), 22–27.

[21] HORN, B. K. P. *Robot Vision*. The MIT Press, Cambridge, Massachusetts, 1986.

[22] JOCHEM, T., POMERLEAU, D., AND THORPE, C. Vision-based neural network road and intersection detection.

[23] JONKER, R., AND VOLGENANT, A. A shortest augmenting path algorithm for dense and sparse linear assignment problems. *Computing 38*, 4 (1987), 325–340.

[24] JUNG, H., KWAK, B., SHIM, J., YOON, P. J., AND KIM, J. Precrash dipping nose (pcdn) needs pedestrian recognition. 678–687.

[25] KENUE, S. Lanelok: An algorithm for extending the lane sensing operating range to 100 feet. vol. 1388.

[26] KOMANOFF, C. Killed by Automobile. Death in the Streets in New York City 1994-1997. Right Of Way, 1999.

[27] MCCALL, J., AND TRIVEDI, M. An integrated, robust approach to lane marking detection and lane tracking.

[28] MONTEMERLO, M., THRUN, S., DAHLKAMP, H., AND STAVENS, D. Winning the darpa grand challenge with an ai robot. In *In Proceedings of the AAAI National Conference on Artificial Intelligence* (2006).

[29] MOSZCZYNSKI, J. N.J. Bicyclists May Get More Room on Road. `http://www.nj.com/news/index.ssf/2009/04/cyclists_hope_proposed_safety.html`.

[30] OUTRAM, C., RATTI, C., AND BIDERMAN, A. The Copenhagen Wheel: An innovative electric bicycle system that harnesses the power of real-time information an crowd sourcing. In *EVER Monaco International Exhibition & Conference on Ecologic Vehicles & Renewable Energies* (2010).

[31] POLYCHRONOPOULOS, M., MÖHLER, N., GHOSH, S., AND BEUTNER, A. *Title of Book*. System Design of a Situation Adaptive Lane Keeping Support System, the SAFELANE System, 2006, pp. 169–183.

[32] PRIEDHORSKY, R., JORDAN, B., AND TERVEEN, L. How a personalized geowiki can help bicyclists share information more effectively. In *Proceedings of the 2007 international symposium on Wikis* (2007), WikiSym '07.

[33] ROSS, D. A., LIM, J., LIN, R.-S., AND YANG, M.-H. Incremental learning for robust visual tracking. *International Journal of Computer Vision 77* (May 2008), 125 – 141.

[34] SUN, Z., BEBIS, G., AND MILLER, R. On-road vehicle detection: A review. *IEEE Transactions on Pattern Analysis and Machine Intelligence 28* (May 2006), 694–711.

[35] WERLBERGER, M., TROBIN, W., POCK, T., WEDEL, A., CREMERS, D., AND BISCHOF, H. Anisotropic Huber-L1 optical flow. In *BMVC'09* (London, UK, 2009).

[36] YU, B., AND JAIN, A. Lane boundary detection using a multiresolution hough transform. *ICIP'97* (1997).