Programmable Automotive Headlights

Robert Tamburo[1], Eriko Nurvitadhi[2], Abhishek Chugh[1], Mei Chen[2],
Anthony Rowe[1], Takeo Kanade[1], and Srinivasa G. Narasimhan[1]

[1] Carnegie Mellon University, Pittsburgh, PA, USA
[2] Intel Labs, Pittsburgh, PA, USA

Abstract. The primary goal of an automotive headlight is to improve safety in low light and poor weather conditions. But, despite decades of innovation on light sources, more than half of accidents occur at night even with less traffic on the road. Recent developments in adaptive lighting have addressed some limitations of standard headlights, however, they have limited flexibility - switching between high and low beams, turning off beams toward the opposing lane, or rotating the beam as the vehicle turns - and are not designed for all driving environments. This paper introduces an ultra-low latency reactive visual system that can sense, react, and adapt quickly to any environment while moving at highway speeds. Our single hardware design can be programmed to perform a variety of tasks. Anti-glare high beams, improved driver visibility during snowstorms, increased contrast of lanes, markings, and sidewalks, and early visual warning of obstacles are demonstrated.

Keywords: Adaptive headlights, reactive visual system, computational illumination.

1 Introduction

Traditional headlights consist of a small number of lamps with simple optics to direct a light beam onto the road. Starting with gas/oil lamps in the 1880s, research has been primarily geared towards developing headlights that can be electrically controlled, have a long working life, and are bright and energy efficient. The inventions of Halogen lamps, Xenon (HID) lamps [4], [5], and the more recent LED [1], [2] and Laser sources [9] have followed this research trend. These latest sources provide bright and comfortable color temperatures improving driving experiences. However, even with these new light sources the only control offered to a majority of drivers is to switch between high and low beams.

Low beams illuminate the road a short range in front of the vehicle while high beams have a longer range and wider angle. High beams are useful in a variety of situations providing better visibility farther down the road and along narrow, curvy roads. However, they cause significant glare to other drivers, bicyclists, and pedestrians. High beams also significantly reduce contrast in the presence of fog and haze, and cause bright distracting streaks during precipitation events. Even after 130 years of headlight development, more than half of vehicle crashes and fatalities occur at night despite significantly less traffic [13]. More than 300,000

D. Fleet et al. (Eds.): ECCV 2014, Part IV, LNCS 8692, pp. 750–765, 2014.

crashes and thousands of fatalities are caused by rain and snow at night annually [13]. Approximately 30% of drivers are stressed by glare causing hundreds of fatalities every year [12]. Thus, a headlight that adapts to the environment can be critical to improving safety on the road during poor visibility conditions.

Recognizing the limitations of traditional headlights, adaptive lighting systems have been developed to adjust their brightness in response to changing driving conditions. Some systems, e.g. Lincoln [6], Audi [2], Volkswagen [7], mechanically swivel the headlight based on the vehicle's turning radius allowing drivers to see around curved roads. Other systems use configurations of multiple LEDs, where individual LEDs can be automatically turned off toward the driving lane and/or the opposing lane to reduce glare, e.g., BMW [21], Audi [22], Mercedes [23], and Volvo [24]. In [9] swiveling LEDs spotlight pedestrians on sidewalks. These advanced systems have come a long way from traditional headlights, but fundamental issues remain: they are not versatile and are designed for one-off applications, they require mechanical components that reduce reliability, and their low-resolution and high latency limits them from adapting to many types of road conditions and poor visibility situations.

This paper presents a new computational illumination design for an automotive headlight that is flexible and can be programmed to perform multiple tasks at high speeds. The key idea is the introduction of a high-resolution spatial light modulator (SLM) such as the digital micro-mirror device (DMD) present in DLP projectors. A DMD divides a light beam into approximately one million beams that can be individually controlled to shape the collective beam for any situation. A sensor (camera) is co-located with the light source and a computer processes images to generate illumination patterns for the SLM. While the design may seem straightforward and follows many works on projector-camera systems in computer vision, there are many challenges in building such a system to serve as a headlight. The accuracy requirements can be high since small errors in beam positioning and flickering are easily perceived and can be more disturbing than standard headlights. High accuracy can be achieved by minimizing the time from when a camera senses the environment to when the headlight reacts (system latency). Low latency is also required to avoid the need for complex prediction algorithms to determine where an object will move next.

A prototype system was built with an ultra-low latency of 1 to 2.5 ms (variation due to factors explained in Section 4) and hence our system requires no prediction algorithm in most cases. We have conducted road demonstrations while traveling at usual traffic speeds to show the feasibility and effectiveness of the design (see website for videos [25]). Example applications include anti-glare persistent high beams, visibility improvement in snowstorms (shown with artificial snow), and illuminating roads with better contrast and lane definition. Results of providing early visual warning of obstacles can be seen at [25]. Our system is able to tackle all of these applications with a single hardware design and achieve higher light throughput than what is possible with any configuration of a small number of controllable LEDs available in current headlights.

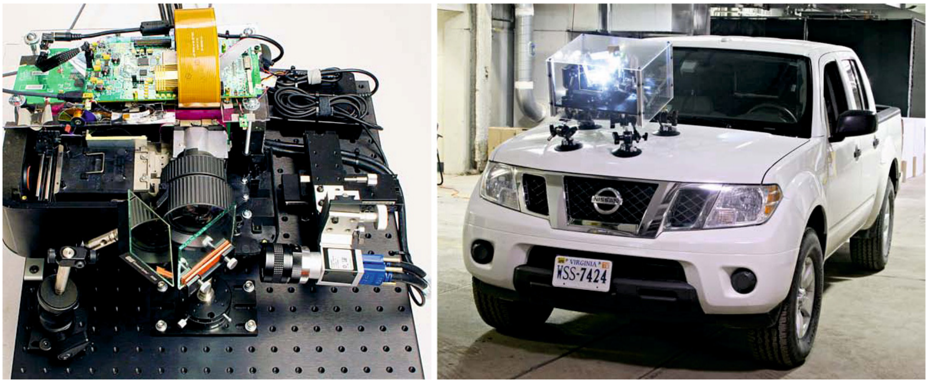

Fig. 1. Left: Prototype of our programmable automotive headlight design (computer not pictured). The camera, spatial light modulator, and beam splitter are firmly mounted to an optical breadboard. A mirror to the side of the beam splitter deflects reflected light from the light source upward. Right: Road tests were conducted by securing the prototype to the hood of a vehicle with a suction cup-based mount. An acrylic enclosure was constructed to protect components from dust, dirt, and moisture.

2 Overview of Programmable Headlight Design

Our programmable headlight design consists of four main components: an image sensor, processing unit, spatial light modulator (SLM), and beam splitter. The *imaging sensor* observes the road environment in front of the vehicle. Additional sensors such as RADAR or LIDAR can be incorporated into the design to complement the camera. The *processing unit* analyzes image data from the sensor and controls the headlight beam via a spatial light modulator. The *spatial light modulator* (e.g., digital micro-mirror device, liquid crystal display, liquid crystal on silicon, etc.) modifies the beam from a light source by varying the intensity over space and time in two dimensions. We use a DMD because its high working frequency and small pixel size permit high-speed modulation and fine illumination control, which makes it possible for our headlight to quickly react to objects as small as snowflakes and objects as large as vehicles.

The camera and SLM are co-located along the same optical line of sight via a *beam splitter*, which virtually places the image sensor and DMD at the same location. Co-location is advantageous because it makes calculating the distance to objects unnecessary. Consequently, there is no need to perform costly computations required for depth estimation and 3D tracking. Also, a single homography will map the camera and projector image planes regardless of the scene. If the image sensor and DMD chip are placed very close to each other, the beam splitter is not required. Reactive visual systems with a similar design have been described by [11], [16], [20], but their systems are too slow for high-speed automotive applications. High latency in conjunction with road effects like wind turbulence and vibration will require complex prediction algorithms that will add latency to the system making it unusable.

3 Design and Implementation of a Prototype System

We designed and implemented a prototype with low latency and high data throughput (Figure 1), and conducted road tests to demonstrate the feasibility of our DMD-based reactive visual system design as a headlight. The camera and SLM must have a very fast frame rate, e.g., kilohertz range, to capture images of fast moving objects and to create illumination patterns that are imperceptible to drivers. Consequently, a lot of data must be transferred to and from the processing unit with minimal latency. To achieve these goals, components with high-speed interfaces were tightly integrated through hardware and software. The prototype measures 45 cm wide, 45 cm long, and 30 cm tall and is currently too large to install in a vehicle as a headlight. The current size is due to using off-the-shelf components. Specialized embedded hardware with an integrated imaging, processing, and SLM unit will be required to create a compact headlight. Road tests were conducted by securing the prototype to the hood of a vehicle with a suction-cup based vehicle mount. A custom acrylic enclosure protects the system from dust, dirt, and moisture. We demonstrate in Sections 5 and 6 that the prototype performs a variety of tasks at typical traffic speeds.

3.1 Sensing the Road Environment

A camera (Basler acA2040) with a CMOS sensor highly sensitive to light with correlated double sampling to significantly reduce noise was used to capture images. The camera is sensitive to visible and near infrared light since most objects of interest are detectable within this spectrum of light. Monochrome imagery is used to avoid the computational overhead associated with demosaicing the Bayer pattern. A global shutter with area scan is used to avoid distortion effects common with the rolling shutter. Latency is reduced via a pipelined pixel architecture that permits exposure during readout. The camera's extended CameraLink configuration has transfer rates of up to 6.8 gigabits per second. The camera is mounted to a set of linear stages for fine control during calibration.

3.2 Image Processing and System Control

A desktop computer provides an interface between the camera and SLM, performs image analysis, and controls the system. The computer was custom built using an Intel Core 3.4 GHz (i7-2600K) CPU with eight cores and hyper-threading technology. A PCI express 2.0 frame grabber (Bitflow Karbon SP) that transfers image data directly into computer memory without any buffering. The main processing tasks were parallelized to reduce latency and increase system responsiveness. The three-stage processing pipeline is shown in a timing diagram (Figure 2) with times measured from the prototype system as described in Section 4. Capture refers to the integration time of the camera. TX denotes the time to transfer image data to the host computer and the time to transfer data from the host computer to the SLM. Process refers to image analysis and system control. Illumination refers to directing light to the scene for a single cycle.

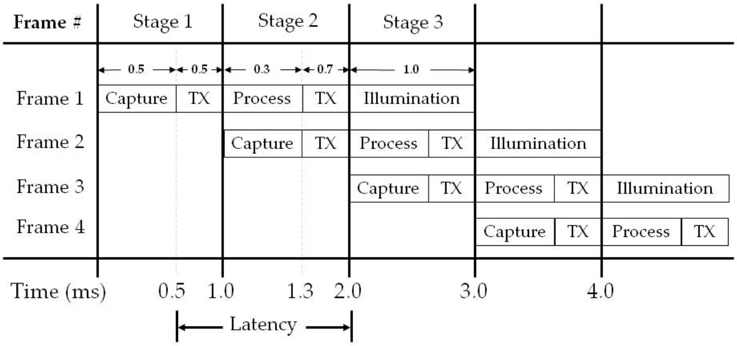

Fig. 2. Timing diagram of the three-stage pipeline with execution times in milliseconds. Capture refers to camera exposure. Process refers to analysis of images and system control. TX denotes data transfer between camera and computer or between computer and SLM. Latency is the time required to illuminate the scene after capturing an image.

Since execution time is critical, the focus of image analysis algorithms is on speed rather than accuracy. Image analyses were performed using OpenCV compiled with Intel Integrated Performance Primitives and Thread Building Blocks to maximize parallelism. Functions that perform per-pixel operations were combined using SSE2 intrinsic functions, when possible, to reduce the computation time associated with multiple iterations over the image. Pre-computable operations such as distortion correction and perspective transformation were initialized and stored in look-up tables. After analyzing images, illumination patterns are encoded and stored in an array then transmitted to the SLM.

3.3 High-Speed Illumination of the Road Environment

A DMD chip is used as a SLM for its spatial and temporal resolution. They are used in consumer DLP projectors, but are driven by video frame rates, which are well below our kilohertz target. A DLP development kit (WinTech W4100) based on the Discovery 4100 (Texas Instruments) was used as the basis of our SLM because the board contains a user programmable FPGA (Xilinx Virtex-5) to achieve fast update rates. The DMD chip is 0.7" with XGA (1024×768) resolution, which, essentially means the headlight beam can be divided into 786,432 smaller beams each of which can be turned on or off. This type of modulation gives unprecedented control over the illumination in space and time. Illumination patterns are received from the host computer by USB 2.0.

The DLP development kit does not include any optics. The optics and light source from a consumer DLP projector (InFocus IN3124) were used instead of designing custom components. We chose this projector because it uses the same DMD chipset as the development kit and uses a lamp (4800 Lumens) brighter than most vehicle high beams. The projector's native DMD chip was removed from the optics module and replaced with that of the development kit via a custom machined mount. A copper heat sink and fan were installed to improve heat dissipation. All of the native DLP electronic boards were left attached to maintain operability even though only the optics and lamp are actively used.

The FPGA was programmed to display patterns faster than 1 kHz. In our design, the FPGA receives data streamed from the host PC and produces the commands for a DMD controller to display the appropriate patterns on the DMD. Each row (1024 pixels) of the DMD is represented as a bit-vector. Transferring a 1024-bit vector for each of the 768 rows was too slow (over 1.5 ms). Instead, the rows are subsampled by a factor of four by representing 1024 pixels by a 256-bit vector. Some resolution is lost, but the visual impact is negligible. Data was further compressed to increase system speed by reading out every other row from the image sensor. The missing rows of the resulting illumination pattern are filled-in by duplicating the previous row on the FPGA. Thus, the image is down-sampled by a factor of 4 horizontally and a factor of 2 vertically.

3.4 System Calibration

Calibrating the system consists of co-locating the camera and SLM, and computing the homography between the camera and SLM image planes. To achieve co-location, a beam splitter with 50% transmission and 50% reflection (Edmund Optics) is used. The projector, rigidly affixed to the optical breadboard, illuminates an object. The camera is translated in all three cardinal directions and rotated until shadows cast by the object are no longer observed by the camera. This recursive co-location procedure takes about 10 minutes to perform.

After positioning the camera and SLM along the same optical line of sight, a perspective transform is calculated for the homography. Radial and tangential distortion by the camera lens is characterized by capturing an image of a checkerboard image and estimating the camera's intrinsic parameters and distortion coefficients. A homography is computed by first projecting a checkerboard pattern and capturing an image. The image is then undistorted and detected corner points are used to compute a perspective transform. After performing these calibration steps, the transformations are stored in look-up tables for later use and the system can be used anywhere without modification. These calibration steps were performed using functionality available in the OpenCV library.

4 Measuring System Latency

As discussed in Section 3.2, the system is pipelined in three stages: image capture and transfer, image processing and transfer, and illumination. Latency of the system is the time between capturing an image and illuminating the scene. There are several factors that contribute to latency. Image size is directly related to camera/computer and computer/SLM transfer time, and image processing time. The size and number of detected objects also has an effect on latency requiring more processing time and thus increases latency. Lastly, the computer's operating system has timing jitter and interrupts that add uncertainty to the latency.

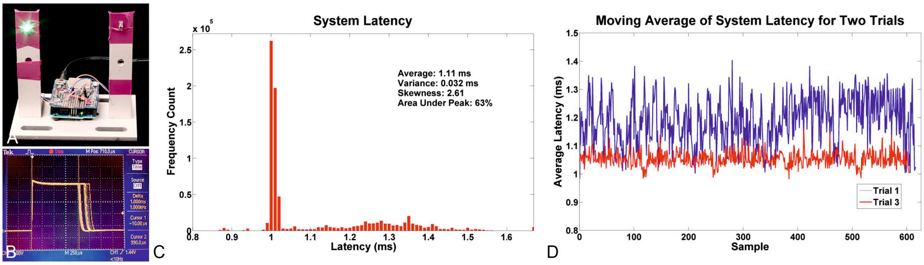

Fig. 3. A: Circuit for measuring system latency consists of an LED and a phototransistor connected to a micro-controller board (Arduino Uno). To measure the system's reaction time, the micro-controller measures the time for the system to detect the illuminated LED then illuminates/dis-illuminates the phototransistor. B: Latency is observed on an oscilloscope (typical readout shown) and measured/recorded by the micro-controller board. C: Histogram (0.01 ms per bin) shows data collected over 30 minutes. The system has some uncertainty, but typically reacts within about 1 ms. D: Moving average (1 second intervals) of latency for a trial with and without uncertainty.

To measure system latency, a circuit was built to measure the system's time to react to an illuminated LED (Figure 3A). The circuit consists of an LED, phototransistor, and micro-controller (Arduino Uno). The high-level idea is to measure the response time of the system by enabling an LED and timing how long it takes for the system to detect the LED and project light onto the phototransistor. To achieve this, the system was programmed to illuminate the phototransistor every other frame. Observing the signal from the phototransistor with an oscilloscope reveals a step response as shown in Figure 3B. The plateau of the signal corresponds to the time that the phototransistor is illuminated. Time was measured with microsecond precision and recorded with the micro-controller.

Data were collected for thirty minutes evenly divided over six separate trials to assess repeatability. During these trials, the image resolution was 800×220, exposure time was 750 μs, and frame rate was 1 kHz. Latency for all the trials is shown in Figure 3C. Across the six trials, the system most often reacts within 1 ms and 63% of the time reacts within two standard deviations from the peak. The average reaction time for all six trials was 1.11 ms with a variance of 0.032 ms. The histogram also reveals uncertainty in the system. This variability was studied by averaging every 1 second worth of data. Shown in Figure 3D are averaged data for two trials: one trial with little variability and one trial with a lot of variability. The plot shows that, in either situation, latency consistently varies by small fluctuations within a narrow band. In the worst case, the fluctuations range from 1 to 1.4 ms.

Several strategies can be utilized to account for latency variability. The uncertainty can be simply included in the illumination pattern by artificially increasing the size of detected objects. Light throughput will decrease, but accuracy will improve. Alternatively, temporal information can be used to predict the location of detected objects. Care must be taken to ensure the prediction model does not add too much time to the system's latency. At high frame rates, a linear model will suffice for most applications. The time to perform processing tasks was measured, in software, with a high resolution timer (Windows API). The average time for processing was 0.3 ms with a standard deviation of 0.04 ms (Figure 2). The time to send data to the DMD board over USB was measured for 5 minutes with an average of 0.76 ms and a standard deviation of 0.07 ms.

5 Anti-Glare High Beams

Glare from the headlights, especially high beams, of oncoming vehicles cause significant stress and distraction at best and temporary blindness at worst. Trucks and other vehicles with headlights at high positions are the worst offenders. Although, glare is not often reported as a cause of accidents, hundreds of fatal night crashes attribute glare as a contributing factor every year [15]. Glare is especially problematic for the elderly whom take eight times longer to recover from glare as compared to a 16-year old [14]. Although high beams are a nuisance to other drivers, they are beneficial on narrow, curvy, and poorly lit roads, especially in rural areas where wildlife routinely jumps onto the road.

Anti-glare headlights are currently being deployed by car companies, e.g., [21], [22], [23], [24]. The details of their systems are publicly unavailable, but it is known that these systems utilize multiple LEDs and sensors placed at different locations in the vehicle, e.g, [10], [9], [2]. Based on this information, it can be inferred that spatial resolution is limited to the number of LEDs. Camera frame rates of these headlight systems are limited to 30 - 60 Hz and thus have high latency [26], [27], [28]. In this section, it will be shown that a high-resolution SLM with low latency produces the best light throughput.

System Requirements and Comparisons. Computer simulations were performed to determine the latency required to maintain high light throughput. Camera parameters and the position of our prototype on a vehicle were used in simulations where two vehicles traveled towards each other at 225 kph on a two-lane, straight road. Detection and prediction were set to be error-free guaranteeing that only system latency contributed to light throughput. Light throughput was calculated for latencies of 2, 16, 30, 50, and 100 ms (Figure 4A). Throughput remains above 90% for all latencies tested when the vehicles are farther than 20 m apart. The reason for this is the oncoming vehicle is moving towards the camera and its position in the image has little variation. However, as the vehicles move closer towards each other, light throughput substantially decreases with higher latency. It is clear that the system needs a latency of at least 2 ms to maintain 90% light throughput when the vehicles are close to each other. The same would be true for vehicles in further lanes or on curved roads.

Computer simulations were conducted to compare the performance LED-based headlights to DMD-based headlights with the same latency. Since specific details of LED-based systems are publicly unavailable, several assumptions were made: (a) LEDs were positioned in a linear array parallel to the road and (b) all LEDs in the array that would illuminate the oncoming driver are disabled.

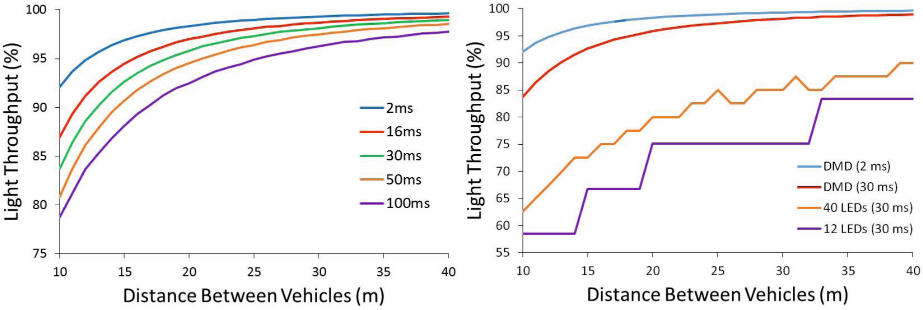

Fig. 4. Results of computer simulations of anti-glare headlights. Detection and prediction are assumed to be perfect and vehicles were traveling towards each other in adjacent lanes at a relative speed of 225 kph. Left: Light throughput as a function of distance between vehicles for different system latencies. Right: Light throughput for DMD- and LED- based anti-glare headlights for different latencies. Simulations show lower latency and higher resolution results in higher light throughput and accuracy, which will be even more relevant for curvy, multilane roads with multiple vehicles.

Simulation results are shown in Figure 4B along with those of the DMD-based system. The low spatial resolution of LED-based systems results in lower light throughput and also creates flicker (abrupt changes in light throughput) for the driver. The flicker can be reduced by turning off more LEDs, but with the trade-off of sacrificing light throughput.

Our Headlight Design as Anti-Glare High Beams. The anti-glare problem and our solution is illustrated in Figure 5A. Headlights from oncoming vehicles are detected in the captured image. Headlights are detected using the assumption that they are the brightest objects in the system's field of view. A very short exposure (100 μs) time is used and the image is thresholded. False detections can be reduced by excluding connected components that are too small to be headlights. Once the locations of the vehicles are known in the camera's reference frame, it is transformed to the headlight reference frame and the spatial light modulator blocks light in that direction. Since the resolution offered by the SLM is very high, only a small region above the detected headlight overlapping the oncoming driver's head is dis-illuminated. This type of beam blocking can be done for any number of oncoming drivers without significant loss of illumination. Compared to the system settings used to evaluate latency in Section 4, the image resolution was increased to 1000×340 to provide the largest field of view possible resulting in a system latency to 2.5 ms.

Demonstration on the Road. The system was tested on the road at night with three oncoming vehicles. Figures 5B-D show video frames captured from inside vehicles driving towards the programmable headlight. In Figure 5B, the blinding glare as the vehicles near each other is shown. Figure 5C and D show the benefit of our anti-glare headlight. Clearly, the difference in visibility is significant allowing drivers to see the road, vehicle, and surroundings. The prototype was able to function for all three drivers at the same time with little light loss.

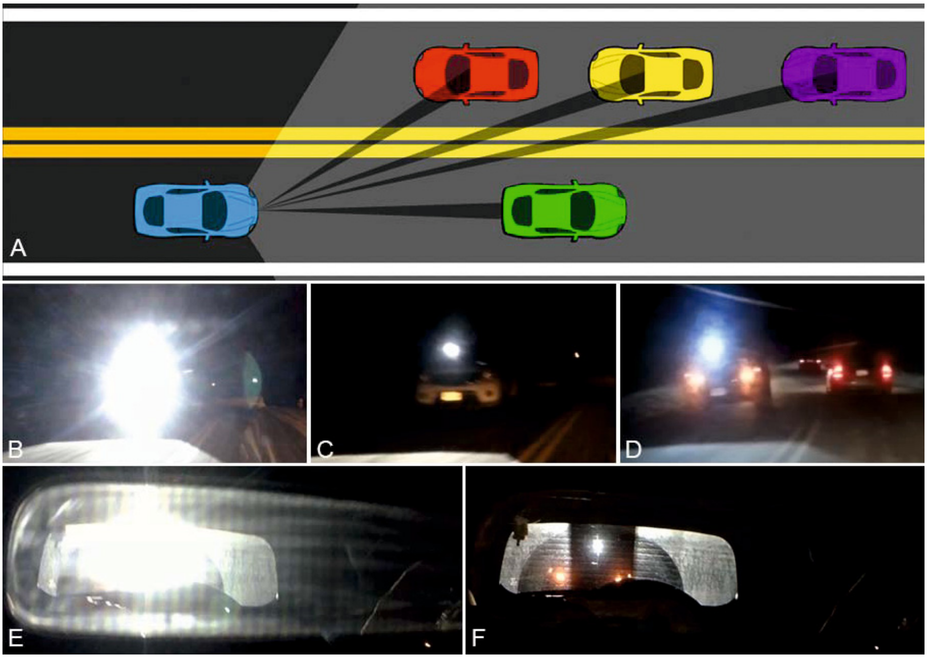

Fig. 5. A: Illustration for eliminating high beam glare. Vehicles are identified and small regions around drivers are dis-illuminated while maintaining illumination elsewhere. Drivers with programmable headlights can then potentially use high beams without worry. Middle row shows view while driving towards our prototype. B: Glare typically seen from high beams (anti-glare feature disabled). C: Reduced glare when the anti-glare feature of our headlight is enabled. D: Anti-glare headlights allow the driver to better see other vehicles on the road. E: Glare in a rear view mirror caused by a following vehicle. F: Tail lights are detected to avoid illuminating the rear-view mirror.

Fig. 6. View shown from the perspective of the vehicle equipped with our prototype. Left: Anti-glare feature is disabled acting as a typical high beam. Middle and Right: Anti-glare feature is enabled detecting multiple oncoming vehicles and reducing light only in the direction of each driver. Notice no discernable difference between images.

The average light throughput was calculated from saved images to be 93.8% with a standard deviation of 3.3%. In Figure 5F, tail lights were detected to avoid illuminating the driver's rear-view mirror and glaring them from behind. As shown in Figure 6, there is no discernible difference to the driver with the programmable headlight when the anti-glare function is enabled. The odd shape of the light beam is due to the system's position and the perspective of the capturing device. Installation in the headlight bay will create a more uniform shape and the spread of the light beam can be increased with a wide angle lens.

6 Demonstrating System Design Versatility

Thus far, computer simulations and demonstrations have shown that the proposed headlight design is advantageous to current anti-glare headlight designs. Our headlight can also be programmed to perform other tasks, whereas, other advanced lighting systems may require additional light sources, sensors, mechanical parts, etc., or are insufficient due to low spatial resolution or high latency. Here we show several tasks, such as visibility improvement in snowstorms (using artificial snow) and illuminating roads with better contrast and lane definition (visual warning of obstacles can be seen at [25]). Also shown is a computational photography application to examine high-speed events.

6.1 Improving Visibility during Snowstorms

Driving in a snowstorm at night is incredibly difficult and stressful. Snowflakes are illuminated brightly and distract the driver from observing the entire road. Researchers in computer vision have proposed methods for removing snow from videos [17], [18], [19]. Processed videos can be displayed for the driver, but current implementations are not intuitive and, at times, distracting for the driver. We can address this problem with a solution similar to that for anti-glare, i.e., reacting to detected bright objects. The main difference, however, is that the density, size, and speed of snowflakes requires high-resolution, low-latency illumination to be effective. Therefore, we exploit the high-resolution and fast illumination beam control of our prototype to distribute light between falling snowflakes to reduce backscatter directly in the driver's visual field (Figure 7). However, this application is significantly more challenging since (a) the size of snowflakes is very small compared to an easily detectable vehicle and (b) the quantity of snowflakes is several orders higher than the number of cars on the road. The goal is to send as much light as possible from the headlight to sufficiently illuminate the road for the driver while dis-illuminating snowflakes.

Computer simulations performed in [11] demonstrate that the idea is feasible. They estimate that, for a vehicle traveling at 30 kph, the system's latency needs to be 1.5 ms or less to have high light throughput and accuracy. We demonstrate improved visibility outside at night with artificial snowflakes. Snowflakes were detected by performing background subtraction and binary thresholding. To compensate for any small detection errors, dilation with a structuring element of a radius equivalent to that of a snowflake was applied. The visibility improvement can be seen by comparing Figures 7B and 7C. Even though the snowflakes fall chaotically, no prediction was required because of the system's fast speed. For comparison, the system by [11] (13 ms latency) was demonstrated for rain drops falling along a straight path and required a linear prediction model.

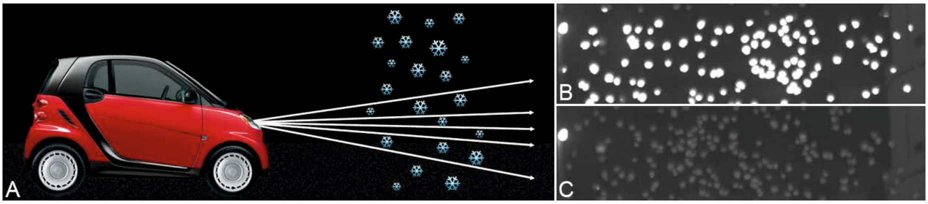

Fig. 7. A: Our headlight has unprecedented resolution over space and time so that beams of light may be sent in between the falling snow. Illustration adapted from [11]. B: Artificial snowflakes brightly illuminated by standard headlight. C: Our system avoids illuminating snowflakes making them much less visible.

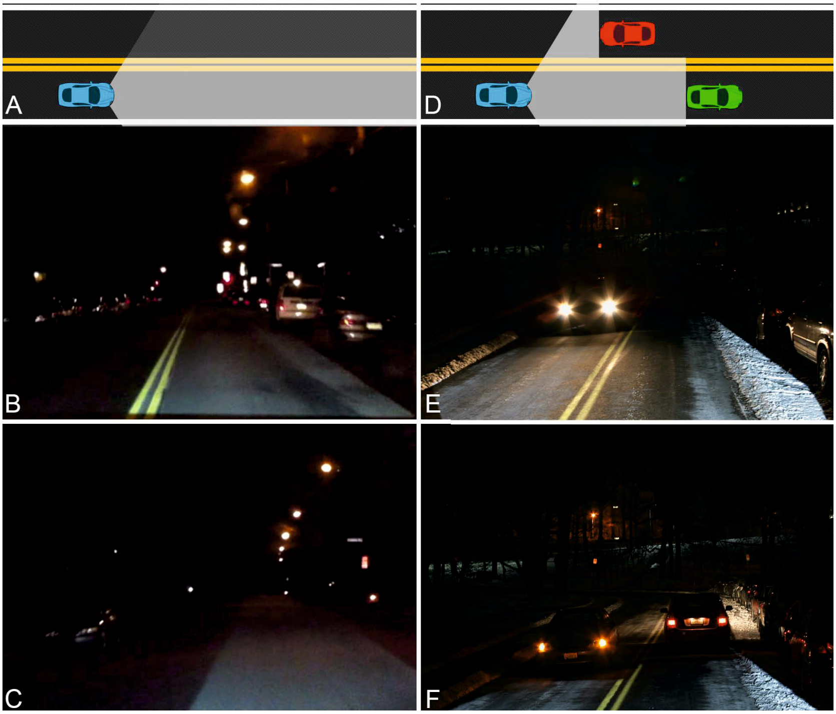

Fig. 8. A: Concept of illuminating the driver's lane with high-intensity light and illuminating the adjacent lane dimly to improve the contrast of the driver's lane. B: Driver's lane more brightly illuminated than the adjacent lane. C: Demonstration while driving on an unmarked road. D: Concept of adjusting lane illumination based on the presence of other vehicles. E: Illumination for the left lane stops at the oncoming driver to avoid projecting lane patterns on the vehicle. F: Lane illumination stops in front of the vehicle in the adjacent lane and behind the vehicle in the driver's lane.

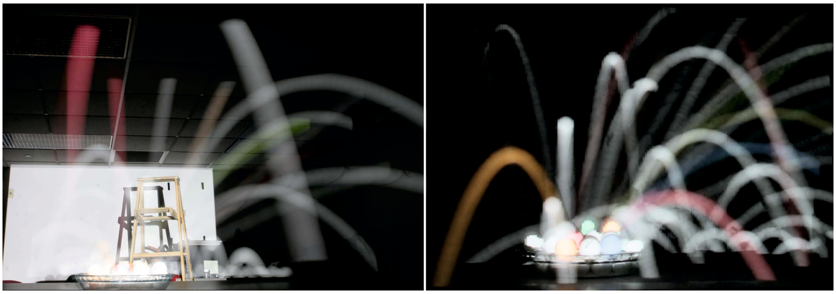

Fig. 9. A tennis ball is thrown into a bowl of ping pong balls causing them to fly through the air. A digital image was captured with a long exposure (3 seconds) to observe the trajectory of ping pong balls. Left: Image captured with scene brightly illuminated. Ball trajectories are not visible against the illuminated wall behind the scene. Right: Image captured with only the ping pong balls being illuminated. The trajectory of the balls is much more visible against the unlit background.

6.2 Improved Lane Illumination

Sometimes the road is not clearly visible and no amount of illumination from a standard headlight can assist the driver. A few examples of such situations are snow covered roads, roads without lane markings or shoulders, and poorly lit roads. Our prototype can be used to brightly illuminate only the driver's lane to provide them with a visual guide. Opposing lanes, curbs, and sidewalks can be dimly illuminated to create a strong contrast with the driver's lane and also provide sufficient illumination to see obstacles (Figure 8A). For this application, images do not need to be captured or analyzed, and objects do not need to be tracked. After computing the homography with the road plane, the headlight acts only as an illumination device. For proof-of-concept, illumination patterns were pre-determined for the stretch of road where experiments were conducted. In practice, the position and speed of the vehicle will be used to dynamically determine the illumination patterns required for the road.

In Figure 8B, the driver's lane and lane markings are fully illuminated, and the adjacent lane is dimly illuminated. The same contrast is used while driving on a dark, unmarked road in Figure 8C. The opposing lane is dimly illuminated while the driver's lane remains fully illuminated creating a demarcation line for the driver to follow. Vehicles driving on the illuminated lane will experience disorienting illumination patterns because the system is calibrated to illuminate the road plane. Therefore, the beam can be adjusted where vehicles are detected in either lane as illustrated in Figure 8D. The adjacent lane can be illuminated up to the location of an oncoming vehicle while maintaining full illumination of the driver's own lane (Figure 8E). Illumination can be controlled in the presence of vehicles in both lanes as well (Figure 8F).

6.3 Observing Events with Computational Photography

Generally speaking, our programmable headlight is a low-latency reactive visual system with many uses outside of the automotive field. It has the flexibility of illuminating or dis-illuminating any fast moving object. An interesting application is studying the trajectory of fast moving objects or fast events. Typically, to capture these types of images, an expensive camera is needed and the room needs to be brightly lit causing a decrease in contrast. Instead, with our system, only the objects of interest need to be illuminated.

For example, a handful of ping pong balls were placed in a bowl. A tennis ball was thrown into the bowl causing the ping pong balls to fly through the air. The ping pong balls were illuminated with infrared LEDs so that they would be detectable in the dark. To observe the trajectory of the ping pong balls, a long exposure (3 seconds) image was captured with a camera. The ping pong balls were detected and immediately illuminated. As shown in Figure 9, the trails are barely visible when the scene is fully illuminated, but are clearly visible when our system is used to illuminate just the ping pong balls.

7 Conclusions and Future Work

The automotive headlight should not be a passive device that can only be completely switched on or off. It should be capable of adapting to the environment to improve safety in poor visibility conditions. Moreover, the design for adaptive headlights should not be limited to a single task. It should be capable of performing many different tasks to help the driver in multiple road environments. Our headlight design provides unprecedented light beam control over space and time. We have demonstrated the flexibility of the headlight for numerous tasks: allowing drivers to use high beams without glaring any other driver on the road, allowing drivers to see better in snow, and allowing better illumination of road lanes, sidewalks and dividers. Our prototype can quickly react to the road environment within 1 to 2.5 milliseconds, and, thus does not create any flicker to be seen by the human eye. Further research and development is needed to make the prototype compact to fit within actual vehicle headlight compartments. Further engineering is required to make the system reliable in the presence of vehicular vibrations and heat. Lastly, more sophisticated algorithms and reliable software need to be developed before deploying our headlight design.

Acknowledgements. This research was funded in parts by a grant from the Intel Science and Technology Center for Embedded Computing, a grant from the U.S. Department of Transportation (Carnegie Mellon University Transportation Center (T-SET)), a gift from Ford Motor Company, a grant from the Office of Naval Research (N00014-11-1-0295), and an NSF CAREER Award (IIS-0643628). The authors also thank the NavLab group at Carnegie Mellon University, Robotics Institute for providing an experimental vehicle platform and Zisimos Economou for helping with the timing circuit.

References

1. Rice, L.: Headlight with Single LED Module. SAE Technical Paper 2010-01-0295 (2010)
2. Söllner, T.: Audi - The Leading Brand in Lighting Technology. Audi Press Release (2013)
3. Plucinsky, T.: BMW Develops Laser Light for the Car. BMW Group Press Release (2011)
4. Bölling, C.: Osram Presents Future Technologies for Car Headlights. Osram Press Release (2013)
5. Schuellerman, D.: GE Lighting Unveils High-Performance Headlamp Lighting Solutions. GE Lighting Press Release (2012)
6. Ford Motor Company: Next Generation of Ford Motor Company's Headlights Make Nighttime Driving Safer. Ford Press Release (2005)
7. Volkswagon Group: To the Point: The New Polo GTI: Extremely Strong and Exceptionally Fuel Efficient. VW Press Release (2010)
8. Boeriu, H.: 2011 BMW 5 Series. BMW Press Release (2009)
9. Wiese, M.: BMW Innovations in Vehicle Lights. "Dynamic Light Spot" for Actively Illuminating Persons, the "Glare-Free High Beam Assistant" and Full-LED Headlights Provide Even More Safety at Night. BMW Press Release (2011)
10. Mercedes Benz: Mercedes-Benz Announces New Active Multibeam LED Headlights. Press Release (2013)
11. de Charette, R., Tamburo, R., Barnum, P.C., Rowe, A., Kanade, T., Narasimhan, S.G.: Fast Reactive Control for Illumination Through Rain and Snow. In: IEEE International Conference on Computational Photography (ICCP), Seattle, Washington (2012)
12. National Highway Traffic Safety Administration: Report on Drivers' Perceptions of Headlight Glare from Oncoming and Following Vehicles (2003)
13. National Highway Traffic Safety Administration: Traffic Safety Facts 2011: A Compilation of Motor Vehicle Crash Data from the Fatality Analysis Reporting System and the General Estimates System (2011)
14. AAA Foundation for Traffic Safety: How To Avoid Headlight Glare (2013)
15. National Highway Traffic Safety Administration: Nighttime Glare and Driving Performance (2007)
16. Wang, O., Fuchs, M., Fuchs, C., Davis, J., Seidel, H.-P., Lensch, H.P.A.: A Context-Aware Light Source. In: IEEE International Conference on Computational Photography (ICCP), Cambridge, MA (2010)
17. Shen, Y., Ma, L., Liu, H., Bao, Y., Chen, Z.: Detecting and Extracting Natural Snow From Videos. Information Processing Letters 110, 1124–1130 (2010)
18. Zhen, C., Jihong, S.: A New Algorithm of Rain (Snow) Removal in Video. Journal of Multimedia 8(2) (2013)
19. Garg, K., Nayar, S.K.: Detection and Removal of Rain from Videos. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR), vol. I, pp. 528–535 (2004)
20. Toshiyuki, A., Osamura, K., Fujisawa, M.: Controlled Illumination for the Object Recognition with Projector Camera Feedback. In: IAPR Conference on Machine Vision Applications, pp. 152–155 (2011)
21. Wiese, M.: BMW Lights the Way into the Future. BMW Press Release (2014)
22. HELLA Inc.: HELLA Develops Unique Matrix LED Headlamp System With Audi. Press Release (2014)

23. Giesen, N.: New Generation CLS with the Future's High-Resolution Precision LED Technology: Leading the way with Better Light. Daimler Group Press Release (2014)
24. Fröberg, P.: Volvo Cars Makes Driving at Night Safer and More Comfortable with Innovative, Permanent High Beam. Volvo Car Group Press Release (2013)
25. Illumination and Imaging Laboratory Project Web Page for Smart Headlights, `http://cs.cmu.edu/~ILIM/SmartHeadlight`
26. MobileEye Camera Matrix, `http://www.mobileye.com/technology/development-evaluation-platforms/cameras`
27. Advanced Driving Assistance and Active Safety Systems, `http://media.opel.com/media/intl/en/opel/vehicles/opel_eye/2009.html`
28. HELLA Group Website, `http://www.hella.com/hella-com`