

Real-time vehicle navigation using a holographic memory

Allen Pu, MEMBER SPIE

Robert Denkwalter

Demetri Psaltis, MEMBER SPIE

California Institute of Technology

Department of Electrical Engineering

136-93 Caltech

Pasadena, California 91125

E-mail: allenpu@sunoptics.caltech.edu

Abstract. We describe an optoelectronic information processing system that is capable of real-time vehicle navigation and target acquisition. The system uses a holographic database, based on the DuPont HRF-150 photopolymer, to perform the desired tasks. The architecture and the performance of the system are discussed in detail. © 1997 Society of Photo-Optical Instrumentation Engineers. [S0091-3286(97)00610-7]

Subject terms: correlation pattern recognition; optical information processing; optical correlator; holographic memory; vehicle navigation; photopolymer.

Paper CPR-06 received Feb. 20, 1997; revised manuscript received Mar. 24, 1997; accepted for publication May 23, 1997.

1 Introduction

A machine that uses visual sensing has access to large amounts of data about the surrounding environment. However, extracting useful information from this large data set is a difficult task given that serial computers can be easily overwhelmed by the sheer volume of data to be examined and by the amount of memory required. There has been recent advances in digital-signal-processing (DSP) hardware and control algorithms that have enabled autonomous vehicles to recognize lanes, forks in the road, and landmarks in real time.^{1,2} In this paper, we present a memory-based rather than a model-based approach to navigation using holographic data storage.

Image processing is well suited for optical implementation because of the inherent 2-D parallelism optics can provide. Furthermore, optical data storage offers high storage capacity and fast parallel access to the stored information. If we combine optical image processing and optical data storage, we can have a system that is well suited to machine vision applications. Specifically, we have constructed a system using an optical correlator as the computational element and a holographic database as the storage device. With our system, an input image is simultaneously correlated with a large number of templates in the holographic database. Systems that combine holographic storage with correlation can produce up to 30,000 correlations/s (an input image updated at 30 Hz correlated against 1000 stored templates). By measuring the position and intensity of these correlation peaks, a control algorithm can be devised to perform complex tasks.

In this paper, we describe an experimental vehicle navigation system that uses optical correlation to direct a vehicle along a predefined trajectory. In this experiment, we wanted the vehicle to be able to autonomously retrace a path that it had previously traveled under manual control. The vehicle uses only visual input to perform this task. The only sensor mounted on the vehicle is a CCD camera. Our approach to this problem is simple. A camera on the vehicle records the observed visual scene as the vehicle is manually driven along the desired trajectory. Sequential still frames from this video are then selected to be “way-

points” and they are stored in the holographic database. The vehicle retraces the desired path by simply driving from waypoint to waypoint until the final destination. The vehicle recognizes the waypoints through the correlation process. Furthermore, the correlation is done against many waypoints simultaneously so the vehicle can look ahead to upcoming waypoints to decide its future actions. The vehicle can also be trained to recognize and stalk targets by storing images of the target in the holographic database.

2 Experimental Setup

Figure 1(a) shows a photograph and Fig. 1(b) the schematic diagram of the vehicle used in our experiment. The vehicle, which we constructed using off-the-shelf components, has three wheels and carries a miniature CCD camera, a video transmitter, a remote-control receiver, two drive motors, and two lead-acid batteries. It is 38 cm long, 31 cm wide, and stands about 20 cm tall (not including the antenna mast). At full speed, the vehicle can travel ~10 cm/s forward or backward. Furthermore, since there are two independent drive wheels, the vehicle can turn and also spin on its own axis clockwise or counterclockwise.

The vehicle itself has no intelligence. It transmits the video from the on-board camera to the optical correlator and receives commands to drive the wheels. All the computation is done within a modified VanderLugt optical correlator shown in Fig. 2 and a personal computer. Real-time video is presented to the optical correlator through a spatial light modulator (SLM). The SLM used in this experiment was removed from an InFocus TVT-6000 video projection system. It has 480×440 pixels and a 3- × 2.5-cm clear aperture. The SLM is illuminated with collimated light from a 200-mW doubled Nd:YAG laser ($\lambda = 532$ nm). In Fig. 2, lens C2 is a 100-cm-focal-length lens and it is placed a few centimeters after the SLM. It produces the Fourier transform of the input image at plane C3. Such a long focal length lens was used to ensure good shift invariance of the correlation and to produce a large Fourier transform pattern at plane C3. A cross-hair spatial filter is placed at C3 to accomplish three goals. It blocks the higher

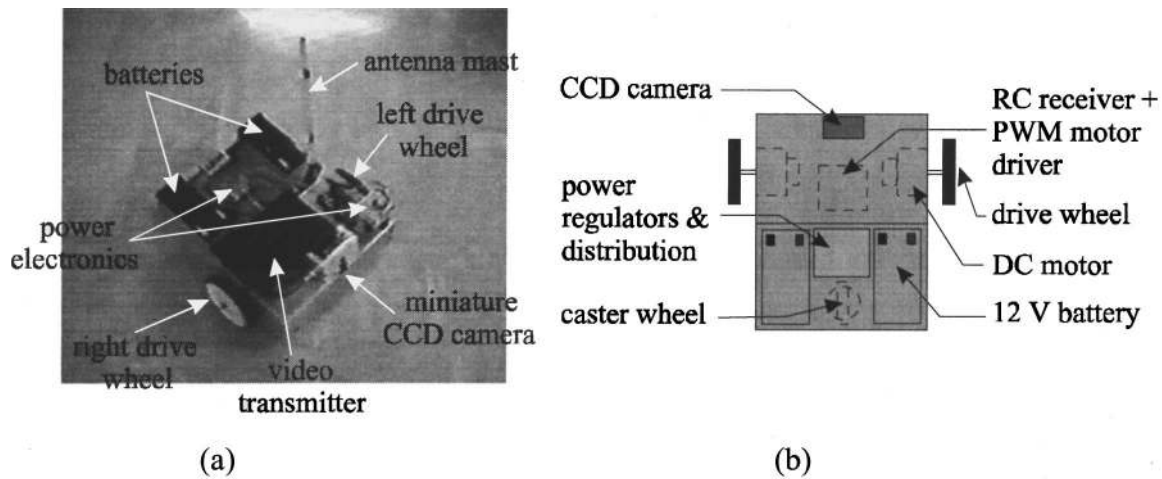


Fig. 1 Photograph (a) and schematic diagram (b) of our vehicle.

diffracted orders that result from the pixelation of the SLM. Spherical aberration in lenses $C2$ and $C4$ limits the precision in which the higher diffracted orders can be collected. The removal of the higher orders gives a clearer image but it reduces the light efficiency of the SLM. The second function of the spatial filter is to block the low-frequency components of the input image. This edge-enhances the input image and greatly reduces the cross-correlation between different images. Finally, a thin cross-hair removes some of the horizontal and vertical lines from the input image to sharpen the correlation peaks. The vehicle navigates in the hallways of a building, where there are many vertical lines representing door ways and horizontal lines representing floor and ceiling tiles. If these horizontal and vertical lines are not attenuated from the input images, then the correlation peaks will also be in the shape of a cross and it will be difficult to determine the exact location of the peaks.

Lenses $C4$ and $C6$ image with magnification one plane $C3$ onto plane $C7$, the plane of the hologram. A second spatial filter is placed between the two lenses at the image plane of the SLM (considering lenses $C2$ and $C4$ as an imaging system). The purpose of this spatial filter is to remove the sharp rectangular edges of the SLM itself. These edges are common to all the images and they are removed to minimize cross-correlation. The aperture of the filter is approximately 3×2.5 mm, since lenses $C2$ and $C4$ form a 10:1 reduction system (the focal length of lenses $C4$ and $C6$ are 10 cm).

The filtered Fourier transform of the input video (at plane $C3$) is imaged onto the recording material at plane $C7$. The recording material we use for this experiment is a $100\text{-}\mu\text{m}$ -thick photopolymer film (DuPont's HRF-150-100). The DuPont photopolymer is well suited for this ex-




Fig. 2 Schematic diagram of the modified VanderLugt optical correlator used in this experiment.

periment because the film develops itself without having to be removed from the optical system for wet processing. This is nice since alignment is fairly critical in an optical correlator. Furthermore, the photopolymer is a write-once, read-many material. Once the recording process is complete, subsequent exposure to light does not decay the previously stored images. Therefore, the correlation process can be continued indefinitely without the need to refresh the stored holograms. Finally, the DuPont photopolymer has very high diffraction efficiency for a thickness of only $100\ \mu\text{m}$. We have found from previous experiments that the M -number of this material is approximately $^3 M/6.5$. A high M -number enables many holograms to be multiplexed on top of each other and still be recalled with sufficient diffraction efficiency. Furthermore, since the recording material is only $100\ \mu\text{m}$ thick, we were able to make the volume hologram effect negligible by using a 100-cm-focal-length Fourier transform lens to obtain full 2-D shift invariance from the optical correlator. By contrast, to get comparable M -number from a LiNbO_3 crystal, we would have to use a thickness of approximately 1 cm, which would result in very limited shift invariance.

The photopolymer laminated to a 1-mm-thick glass substrate is mounted on a rotational stage at plane $C7$. The photopolymer is placed normal to the signal beam and the center of rotation of the rotational stage is aligned with the optical axis of the signal beam. A plane wave reference beam that forms an angle of 30 deg (outside) with the signal beam is introduced to record holograms. Multiple holograms are stored through peristrophic multiplexing.⁴ After a hologram is recorded, the photopolymer is rotated in-plane by 1 deg before the next hologram is recorded. For a correlator, peristrophic rotation moves the correlation peaks in the trajectory of a circle around the signal beam (at correlation plane). The 1-deg clockwise peristrophic rotation moves the correlation peak of the previously stored hologram slightly above the focal point of the reference beam, as shown in Fig. 3 (blowup view of the correlation plane in Fig. 2). This enables another filter to be stored at the same location on the photopolymer and still have spatially distinct correlation peaks. This process continues until all the filters have been stored (up to 180 filters using peristrophic multiplexing alone). Since the holograms are all stored in the same area, all the correlation peaks are produced simultaneously when an input image is presented. Fourier transform lens $C8$ produces at the correlation plane the correlation peaks. A few correlations peaks are shown at the correlation plane of Fig. 2. Three of these peaks are monitored by a CCD camera, $C9$, placed at the correlation plane. Other correlation peaks can be brought into the field of view (FOV) of the CCD camera by rotating the photopolymer film.

To determine the proper action for the vehicle to take, we decipher the video from the CCD camera placed at the correlation plane. The information of interest is the intensity and the position of the peaks. We ignore and suppress other information such as the structure of the correlation peaks, which contains clues about the environment the vehicle is in. We used a personal computer to extract the intensity and the position of the peaks. The video from camera $C9$ is fed into the computer's memory through a framegrabber board. A search algorithm locates and deter-




Fig. 3 Peristrophic multiplexing of filters by rotating the recording material in-plane.

mines the intensity of the three correlation peaks. To follow the waypoints stored in the holographic database, the computer issues commands according to the flowchart shown in Fig. 4. First, the computer locates the brightest peak (the CCD pixel array is turned into a Cartesian coordinate system with the origin at the center of the CCD array). If the brightest peak is centered in x' within a preassigned margin ϵ , then the vehicle is assumed to be on path and both wheels are set to their nominal velocity to drive straight ahead. A centered peak indicates that the vehicle's viewpoint is the same as when the vehicle was driven manually down the path the first time (see Fig. 5). However, if the peak position is shifted to the left or right, the current viewpoint of the vehicle is shifted with respect to the stored waypoint. This happens when the vehicle is heading in a direction different than the original direction. The computer then issues commands to drive one of the wheels faster to turn the car back to the desired heading.

We monitor three peaks with the CCD array, $C9$, placed at the correlation plane. The brightest peak is supposed to be in the middle of the CCD array and it represents our current position on the path (see Fig. 6). The top peak represents the next waypoint the vehicle is looking for and the bottom peak represents the last waypoint the vehicle was at. In this way, we can keep an eye on what is coming up next and where we were before. As the vehicle progresses down the path, eventually the top correlation peak will become brighter than the middle peak. This means that the vehicle has reached the next waypoint and the computer issues a command to advance the holographic memory by rotating the photopolymer film to bring the top peak to the middle. Now a new waypoint appears on top and the old last waypoint drops out from the FOV of the CCD array. The preceding process is repeated continuously until the final way-

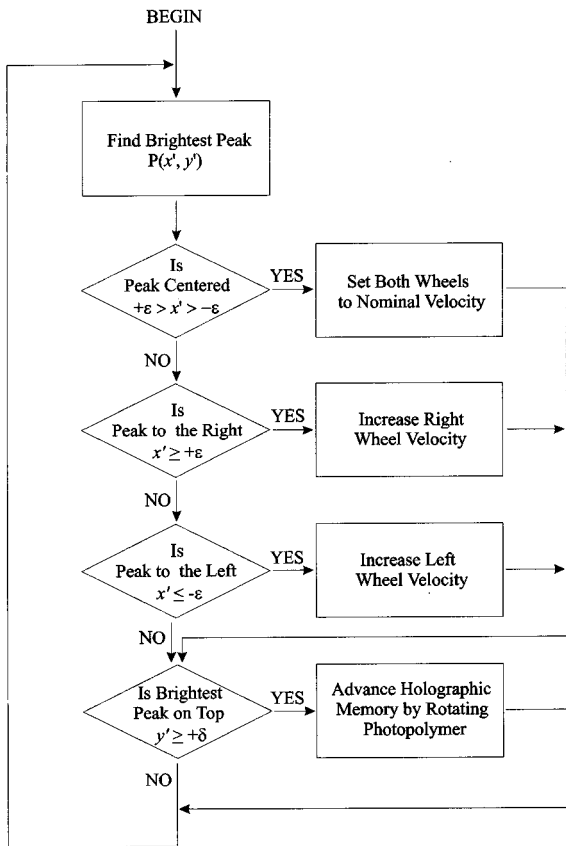


Fig. 4 Flowchart of the basic vehicle control algorithm in the computer.

point has been reached and thus the path is complete.

Figure 7 shows an overview of the entire system. The vehicle sends the live video back to the optical correlator. The system computes the current position and the trajectory of the vehicle from the correlations with stored waypoints, and it determines the location of the brightest peak. It sends control signals to the vehicle, and advances the holographic memory. Notice that once the control algorithm has been




Fig. 5 Position of the brightest peak in x' determines if the vehicle is on path.

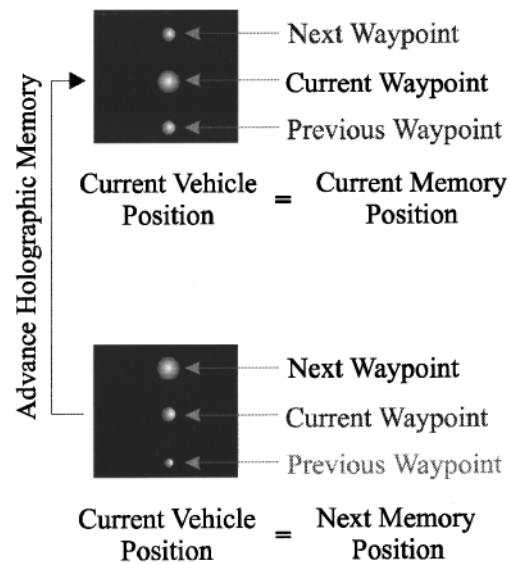


Fig. 6 Position of the brightest peak in y' determines if the vehicle has reached the next waypoint.

finalized, the computer and the CCD array at the correlation plane can both be replaced by a single very large scale integration (VLSI) chip. Since the control algorithm is just a simple set of hardwired decisions, a VLSI chip with CMOS detectors and simple logic would make the optical correlator a self-contained system with input (video from vehicle), computation (correlation), memory (holographic database), and output (commands to drive the vehicle).

3 Modes of Navigation

The flowchart shown in Fig. 4 describes how the vehicle follows the stored waypoints to the final destination. We call this the “follow mode.” If the vehicle is pointed in a direction so that it sees an unfamiliar scene, however, then no correlation peak is produced and the vehicle would be lost. In this case, the vehicle switches to the “acquisition mode,” where it rotates in place. More specifically, if the top correlation peak is undetectable, then the vehicle is instructed to rotate in place until the top peak becomes visible again. This is a logical choice since if the vehicle cannot see the next waypoint, there is no point in going ahead and getting further lost. The acquisition mode is also useful when the vehicle’s FOV is blocked temporarily by an object. When that occurs, the top correlation peak goes blank and the vehicle would simply rotate in place until the interfering object is removed.

With follow and acquisition modes, we can navigate the vehicle through the desired path even when the vehicle experiences unexpected perturbations. To further demonstrate the utility of our system, we decided to have the vehicle find and then stalk a target once the vehicle has reached a destination. The “stalking mode” is just a simple addition to the control algorithm and it works as follows. Far and near images of a target are stored after the last waypoint in the holographic memory (far and near correspond to the distance away from the miniature CCD camera on the vehicle, see Fig. 8). After the vehicle has reached the final waypoint, the computer advances the holographic memory




Fig. 7 Overview of the entire autonomous vehicle navigation system.

to bring the far and near correlation peaks onto the CCD C9. If the target is present and is at the desired distance away from the vehicle, then the two correlation peaks will have equal intensity. Furthermore, if the correlation peaks are also centered with respect to x' , then the vehicle is correctly positioned and will stop both wheels and fixate on the target. However, if one of the peaks is brighter than the other, then the vehicle will either backup or drive forward until both correlation peaks are of equal intensity and the vehicle is at the desired distance away from the target. If the peaks are not centered, then the vehicle will turn in place to bring the target back to the center of its FOV. If there are no detectable correlation peaks, then the target is

not present in the current FOV and the vehicle is instructed to rotate in place until the target is found.

Using these navigational modes, we can drive a vehicle from place to place and then find/stalk a target of interest. Other tasks can be accomplished by using a combination of these basic modes.

4 Holographic Database

Figure 9(a) shows a photograph and Fig. 9(b) the layout of the hallway the vehicle navigates. The width of the hallway is approximately 2.4 m and the floor is covered with glossy square tiles (23×23 cm per tile). The video output from the




Fig. 8 Staling mode operation.

onboard miniature CCD camera was recorded on a laser disk as the vehicle was driven manually through path 1 from room A to room B. We take still frames from this video to store in the holographic database as waypoints for the vehicle to follow. A video sequence was also recorded for a second path. Path 2 starts from room C and goes up to the intersection point with path 1. With waypoints extracted from these videos, the vehicle can start in anyone of the rooms and navigate to any other.

How far apart should the waypoints be separated (in hallway distance) for our system to function properly? If the waypoints are separated too far, then the vehicle will be lost since the top correlation peak corresponding to the next waypoint will be dim and undetectable by the time the middle correlation peak fades away. On the other hand, if the waypoints are too close together, then we will be over-sampling the path and wasting memory on redundant information. Initially we selected the waypoint separation so that the top correlation peak is slightly above the detectable threshold when the middle correlation peak reaches maximum intensity. In this way, the vehicle will always be able to see the next waypoint as it navigates the path. Since the hallway is sometimes noisy (pedestrians or foreign objects blocking parts of the hallway), however, we found that it is

better to ease off from this separation criterion and have some redundancy by selecting waypoints that were closer in distance.

Through trial and error, we found that waypoints separated by one floor tile (23 cm) gives the desired performance in our environment. Of course the required separation would depend on the specific environment. For example, if the hallway was 1 m in width instead of 2.4 m, then the waypoints might have to be closer together. This is due to the fact that the rate of change of the surrounding as seen by the moving vehicle is greater in the 1-m hallway than with the 2.4-m hallway.

Figure 10 shows the 46 waypoints needed to describe path 1 from room A to room B (order of the sequence is from left to right, in a raster scan format). The first 9 frames show the vehicle driving out of room A and approaching a parked bicycle. The bicycle was placed in front of room A to add structure to the wall. Without the bicycle, the vehicle would be looking at a nearly blank wall and the correlation function will fail. Frames 10 through 38 define the straight path down the middle of the hallway. Fortunately, the hallway itself is full of bulletin boards and door frames so that additional structure is not required. The sharp transition from frame 9 to 10 is intentional so that the vehicle would

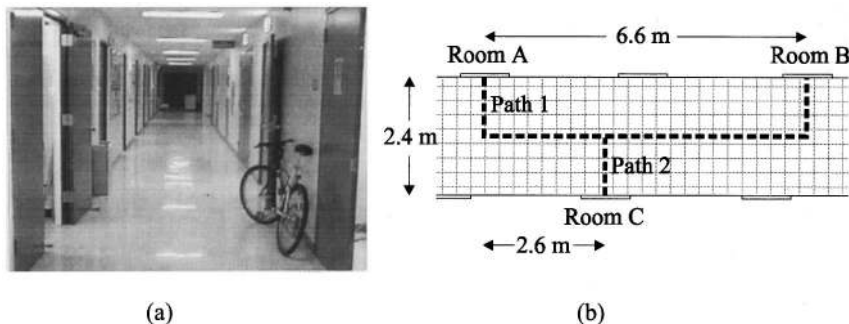


Fig. 9 Photograph (a) and the layout (b) of the hallway the vehicle navigated.



Fig. 10 Forty-six waypoints stored to describe path 1.



Fig. 11 Five waypoints stored to describe path 2.

get lost temporarily. When the vehicle reaches the location of frame 9, the next waypoint the vehicle is looking for is frame 10 (down the hallway). Since the vehicle is still looking at the bicycle, the top correlation peak corresponding from frame 10 is undetectable. Using the control algorithm already discussed, the vehicle is instructed to rotate in place until the next waypoint is found. After the vehicle has rotated far enough so that it is looking down the hallway, the top correlation peak will become bright enough to be detected and the holographic memory is advanced to bring the top peak to the middle. The vehicle then continues navigating on the follow mode. We use the acquisition mode to make turns to save on memory. The alternative would be to store the turn itself as a sequence of waypoints and use the follow mode to complete the turn. This would add an additional four to six waypoints for each 90-deg turn, or a 20% increase in memory requirement for path 1. The final eight frames of Fig. 10 shows the vehicle driving toward room B. Once the vehicle reaches the last waypoint, it is instructed to stop. Path 2 is described by the five waypoints shown in Fig. 11. When the vehicle reaches the last waypoint on path 2, it turns (using acquisition mode) and follows frame 20 of path 1 to reach room B.

The combined total of 51 waypoints (46 from path 1 and 5 from path 2) are stored in the DuPont HRF-150, 100- μm -thick photopolymer using peristrophic multiplexing. Three correlation peaks fit on the active area of the CCD array C9 given the 1-deg peristrophic rotation and the focal length of lens C8. The holograms were recorded with a signal and reference beam intensities of approximately $500 \mu\text{W}/\text{cm}^2$ each. It is desirable for the maximum correlation peak intensity for all the waypoints to be the same or the control algorithm might not advance the holographic memory properly during the experiment. The peaks can be equalized by simply varying the recording time of the waypoints according to the number of edges it contains. For waypoints with many edges, a shorter recording time would produce a correlation peak that is of the same intensity as a waypoint with fewer edges recorded for a longer duration. For path 1, the recording time varied from 1 to 3 s per waypoint and the recording schedule was obtained through trial and error. At first, all the waypoints were stored with the same exposure time. The waypoints were then played back from the laser disk into the optical correlator for autocorrelation. For noticeably stronger or weaker autocorrelation peaks, their recording time was decreased or increased accordingly. A recording schedule that produced visually equalized correlation peaks was obtained after four trials. For the 51 waypoints that were recorded, we stayed well within the linear region of the photopolymer and did not approach dynamic range saturation.




Fig. 12 Sequence taken from an experiment where the vehicle navigated path 1.

The 1-deg peristrophic rotation between holograms has another interesting and desirable effect. The hologram representing the top correlation peak is rotated by $+1$ deg (in orientation) with respect to the hologram representing the middle peak (the middle hologram is at 0-deg rotation with respect to the input SLM). Similarly, the hologram representing the bottom correlation peak is rotated by -1 deg with respect to the middle hologram. During correlation, this $+1$ deg of rotation diminishes the peak intensity of the top and bottom correlation peaks while leaving the middle peak unaffected. This is due to the fact that the correlation function is not rotationally invariant. What this means is that the vehicle must drive close to the next waypoint before the top correlation peak becomes brighter than the middle peak. This inhibition of the top (and bottom) correlation peak is desirable because it prevents the computer from advancing the holographic memory too far ahead of the vehicle's current position, which could occur if not all the waypoints had similar maximum peak intensity.

5 Experimental Results

After the waypoints have been properly selected and recorded, the vehicle was able to consistently and accurately retrace the programmed path using a combination of follow and acquisition modes. Figure 12 shows a sequence of the vehicle navigating path 1. Each frame consists of the video

from the miniature CCD camera (main view), the correlation peaks at the correlation plane (upper left-hand corner), and the image of the waypoint corresponding to the top correlation peak (upper right-hand corner). When the vehicle first starts from room A (frame 1 of Fig. 12), it only has two correlation peaks corresponding to the current position and the next waypoint. It does not have a previous position stored in its memory so the bottom correlation peak is blank. As the vehicle moves forward toward the bicycle, the holographic database advances and all three peaks become detectable. Frame 4 of Fig. 12 shows that the vehicle has reached the midpoint of the hallway and it should execute a turn to the port side at this point. Notice that the top correlation peak in frame 4 is undetectable since the next waypoint the vehicle is looking for is completely different from its current view. This activates the acquisition mode and the vehicle begins to rotate in place. Frame 5 shows the view from the vehicle in the middle of the rotation and all the correlation peaks have gone blank. After the vehicle has rotated far enough, its viewpoint is the same as the next waypoint so the top correlation peak becomes detectable again (frame 6). The control algorithm switches from acquisition mode back to follow mode and advances the holographic memory to bring the top peak to the middle (since the top peak is brighter, frame 7). Frame 8 shows the vehicle navigating using follow mode as it had before in frames 1, 2, and 3. In follow mode, the velocity of the vehicle was 7 cm/ps for the experiment shown in Fig. 12. Later, with an improved control algorithm, the velocity approached the limit of vehicle at 10 cm/s.

The computer monitoring the correlation peaks is aware of the position of the vehicle with respect to the hallway. The computer knows which correlation peak is currently the brightest and by reading the angle on the rotational stage, it can pinpoint the current waypoint the vehicle is at. This information is useful for more complex instructions such as navigating the vehicle to a particular waypoint and then branch off to join another path. For the experiment where the vehicle was navigated from room C to room B, the computer kept track of the vehicle's position on path 2. After the vehicle reached the fifth and last waypoint of path 2, the computer instructed the holographic memory to scan quickly to place frame 20 of path 1 as the next waypoint. The vehicle then used acquisition mode to find path 1 and followed the waypoints to room B. We can also program the vehicle to drive backward from room B to rooms C or A by using the bottom correlation peak as the next waypoint and the top peak as the previous waypoint.

The sequence of frames in Fig. 13 show the types of noise we introduced into the hallway to try to fool the vehicle. In frame 1, a table and a man doing jumping-jacks at the middle and the end of the hallway did not hamper the normal operation of the vehicle. A remote-controlled truck ramming the vehicle, as shown in frame 2, was also unable to make the vehicle lose its way. The vehicle was deviated temporarily from its course but always recovered and was able to drive back onto the path. In frame 3, the vehicle was rotated manually while it was in the follow mode. This caused the correlation peaks to disappear and when the vehicle was replaced on the floor, it rotated in place until it was once again aligned with the path and then continued to its final destination. Frame 4 shows that with up to 40 to




Fig. 13 Types of noise we introduced to test the performance of the system.

50% of its FOV blocked, the vehicle was able to continue driving past the obstacle. However, had the person stepped directly in front of the path of the vehicle, its FOV would be blocked sufficiently to cause the vehicle to rotate in place until the obstacle is removed.

One way we could confuse the vehicle was to move it laterally away from the middle of the hallway by a few feet (toward one of the walls). This changes the perspective of the hallway as seen by the vehicle and correlation with the stored waypoints will fail. The vehicle becomes lost and rotates in place without ever recovering the path.

To demonstrate the stalking mode, we programmed the vehicle for a different path, from room A into another room with a cleared center. At the end of the sequence of waypoints, far and near images of a red monster truck were recorded and designated as the target. During the experiment, the vehicle navigated from room A into the cleared room using the stored waypoints. Once the vehicle has reached the last waypoint, the computer switches into the stalking mode and rotates the vehicle in place until the target is in view. The sequence in Fig. 14 shows the stalking mode in action. In frame 1, the vehicle has acquired the




Fig. 14 Sequence showing the vehicle stalking a monster truck.

target located at the desired distance from the vehicle. Frame 2 shows the target moving toward the starboard side of the vehicle and the vehicle rotates in place to track the target. If the target moves away from the vehicle, the vehicle drives forward until the desired distance has been reached (frame 3). Finally, if the target gets too close to the vehicle, the vehicle backs up and redistances itself from the target (frame 4).

6 Discussions and Conclusions

The optical correlator system shown in this paper was able to guide a vehicle through a preprogrammed path and automatically recognize and track targets in real time. In each experiment, less than 100 holograms were required, which is only a small fraction of the holographic database's total capacity. Up to 1000 holograms have been stored at a single location using the DuPont 100- μm -thick photopolymer³ and up to 10,000 holograms have been stored in photorefractive crystals.⁵ Even larger capacity can be achieved by multiplexing holograms spatially on a disk.⁶ Such a large capacity can enable an entire building or a campus to be mapped for autonomous vehicle navigation, a number of targets stored from all possible perspectives, scales, and elevations for automatic target recognition and tracking.

There are many ways in which the performance of this system can be improved. For navigation, instead of recording only a single path from room A to room B, down the middle of the hallway, several paths can be recorded. This will enable the vehicle to navigate around obstacles blocking a particular path. For stalking, images of the target at many different distances from the vehicle and from different angles can be stored to extend the range of target acquisition. More clever use of the correlation peaks, improvements in the control algorithm, and optimized correlation templates can further increase the capability of our system beyond its current simplistic form.

Acknowledgments

This work was supported by the Defense Advanced Research Projects Agency and the ERC center. Additional support from Rockwell Science Center is gratefully acknowledged.

References

1. E. D. Dickmanns and N. Muller, "Scene recognition and navigation capabilities for lane changes and turns in vision-based vehicle guidance," *Control Eng. Pract.* **4**(5), 589–599 (1996).
2. N. P. Papanikolopoulos, P. K. Khosla, and T. Kanade, "Visual tracking of a moving target by a camera mounted on a robot—a combination of control and vision," *IEEE Trans. Robot. Automat.* **9**(1), 14–35 (1993).
3. A. Pu, K. Curtis, and D. Psaltis, "Exposure schedule for multiplexing holograms in photopolymer films," *Opt. Eng.* **35**(10), 2824–2829 (1996).
4. K. Curtis, A. Pu, and D. Psaltis, "Method for holographic storage using peristrophic multiplexing," *Opt. Lett.* **19**(13), 993–994 (1994).
5. G. W. Burr, X. An, F. H. Mok, and D. Psaltis, "Large-scale rapid access holographic memory," in *Optical Data Storage Meeting 1995*, *Proc. SPIE* **2514**, 363–371 (1995).
6. K. Curtis and D. Psaltis, "3-dimensional disk-based optical correlator," *Opt. Eng.* **33**(12), 4051–4054 (1994).



Allen Pu received his BS from The Cooper Union in 1992, MS and PhD from the California Institute of Technology, Pasadena, in 1993, and 1997, respectively, all in electrical engineering. He is currently employed by Holoplex, a startup company located in Pasadena, conducting research in optical computing with special interests in holographic data storage and optical correlators.

and a consultant to industry. His research interests are in optical memories, optical information processing, holography, pattern recognition, neural networks, and optical devices. Dr. Psaltis has authored or co-authored more than 200 publications. He is a fellow of the Optical Society of America and received the International Commission on Optics Prize in 1989.

Robert Denkwalter: Biography and photograph not available.



Demetri Psaltis received a BSc in electrical engineering and economics in 1974 and MSc and PhD degrees in electrical engineering in 1975 and 1977, respectively, all from Carnegie Mellon University, Pittsburgh. He remained at Carnegie Mellon, as a research associate and later as a visiting assistant professor, for 3 years. In 1980, he joined the faculty at the California Institute of Technology, where he is now executive officer of the Computation and

Neural Systems Department, professor of electrical engineering,