

Autonom-ish Cars

Mobileye's plan to leapfrog Google and Uber to self-driving ubiquity

BY MAX CHAFKIN

ILLUSTRATIONS BY YANN KEBBI

On Sunday, March 18, a 49-year-old woman named Elaine Herzberg walked her bicycle across Mill Avenue, a divided highway just north of downtown Tempe, Ariz., and was hit head-on by a Volvo SUV traveling at about 40 miles per hour.

The death would have been an ordinary tragedy—6,000 pedestrians were killed by cars last year in the U.S. alone—but for one thing: The driver was a piece of software created by Uber Technologies Inc. This was the first known fatality caused by a driverless car. The ride-hailing app company and its main chip supplier, Nvidia Corp., suspended testing indefinitely, and the U.S. National Transportation Safety Board announced it would conduct an investigation. Uber is cooperating.

In Jerusalem, Amnon Shashua began his own tests. Shashua, a 57-year-old Israeli, is a professor at Hebrew University of Jerusalem and the chief executive officer of Mobileye. The company, which he co-founded in 1999, makes a “driver-assist” system roughly the size of a computer mouse. Stuck behind a car’s rearview mirror, the device’s camera and custom chip allow it to predict when a driver is about to hit something and, if necessary, the system slams on the car’s brakes.

Mobileye sells its chips to auto suppliers for about \$55 each—you’ll pay around \$1,000 at the dealership for the full system—and they’re already on the road in 27 million cars worldwide. This uniquely broad distribution was one advantage Intel Corp. gained when it bought Mobileye last year for about \$15 billion. The deal was Israel’s biggest acquisition ever, and a significant departure for the American chipmaker, which folded its old driverless car division into Mobileye. “We’re going to do what it takes to win,” says Intel CEO Brian Krzanich.

The other feature that distinguishes the company is its outspoken approach to safety. Most pitchmen for self-driving cars tend to shy away from discussing possible failings of the technology for fear of spooking customers and regulators. Shashua has made addressing these concerns central to his marketing.

At the end of March, after Tempe police released a dashcam video showing what the car had seen just before it killed Herzberg, Shashua fed the video into the Mobileye system to see how his computer vision system would react. As he wrote in an article on Intel’s website under the headline “Experience Counts,” his software correctly identified Herzberg as a pedestrian. In an interview at Mobileye’s headquarters a month later, Shashua, wearing a gray T-shirt and jeans, said Uber and its main competitors—most notably Alphabet Inc.’s Waymo—were “making something that kind of works.” He argued that Herzberg’s death was “avoidable.”

If there’s something gross about using a tragedy as an opportunity for self-promotion, there are reasons to take Shashua seriously. Mobileye has roughly a 70 percent share of the driver-assist market, and in 2016 the National Highway Traffic Safety Administration reached an agreement with industry groups that will make the systems a standard feature for all new vehicles sold in the U.S. by 2022. It says this will prevent a total of 28,000 crashes and 12,000 injuries by 2025.

Shashua wants to go much further. In 2015, Mobileye began selling more sophisticated systems that can take over highway driving duties. Tesla Inc.’s Autopilot was the first such system, and Mobileye now supplies hardware and software that powers Cadillac’s Super Cruise, Nissan’s ProPilot, and Audi’s Traffic Jam Assist. Shashua is also developing a fully self-driving package that Mobileye will begin offering automakers as early as next year. BMW, Fiat Chrysler, and China’s NIO all plan to use it.

In late April, Shashua took a *Bloomberg Businessweek* reporter for a ride in the fully autonomous prototype—the first time the company demonstrated it to a journalist. This, he argued, was what a road-ready autonomous car might actually look like. He dismissed Waymo’s and Uber’s efforts as science projects crammed with expensive radars and other sensors, plus the tens of thousands of dollars’ worth of servers and chips stuffed in their trunks.

Those companies say their costs are falling, and also that more expensive models are feasible because consumer ownership is an outmoded concept—most cars will eventually be part of a shared fleet of robo taxis rented by the mile. Shashua is determined to deliver a driverless car on the relative cheap. His prototypes, a half-dozen Ford Fusions now testing in Jerusalem, part of a planned fleet of 30 cars that will arrive in California later this year, have no radar or lidar (basically radar, but with lasers). The only sensors are 12 digital cameras hidden around the car’s body. Mobileye eventually plans to add some low-cost radars and lidars, but only as backup if a camera fails. Instead of running on a trunk-size bank of servers and chips, the prototype relies on just four of the same low-cost, custom-made processors in Mobileye’s current driver-assist system. Consumers, Shashua says, will pay less than \$8,000 for the whole thing. “So that means anyone could have a driverless car,” he says, perhaps overestimating the average car buyer’s budget.

Mobileye’s competitors have generally been testing their machines in suburbs with wide, well-marked roads and orderly traffic patterns. Shashua has been testing in Jerusalem, a city with narrow, medieval roads and a driving culture charitably ▶



◀ described as assertive. “There’s a big difference between something that kind of works and something that works,” he says. “I don’t think the public understands this.”

Driverless cars have to do two basic things: see what’s on the road and then react to that information. Shashua has worked on the first problem, computer vision, since the 1980s, long before it found its way into cars. Born in Tel Aviv, he studied computer science at Israel’s elite Weizmann Institute of Science before enrolling in an artificial intelligence Ph.D. program at MIT.

In 1998, five years after earning his Ph.D., Shashua was in Nagoya, Japan, where a company he’d started, CogniTens, had been working for Toyota Motor Corp. using digital photography to measure bumpers as part of Toyota’s quality-control process. During a meeting, he suggested as an aside that a single digital camera inside the cars could be used to warn people if they were about to hit something or drift out of their lane. The automaker was skeptical but agreed to give him a few hundred thousand dollars to do research.

No field is more influenced by state-of-the-art AI than computer vision. Today, massive data sets can teach software to distinguish obstacles in a few weeks. Back then, building a computer that could see was tougher. Mobileye’s engineers, mostly students of Shashua’s from Hebrew University and MIT, started by assembling lists of road features the cars needed to identify. “You start with edge detection,” recalls Gideon Stein, Mobileye’s chief research scientist and an early employee. “And then, well, the back of a vehicle is really a box, isn’t it? So what I’m looking for is a square. Then you say, ‘See that little line between the wheels? Let’s make a detector for that.’” The process took years.

Some of Mobileye’s early code is still running as a backup to newer types of computer vision and to measure distances, but deep learning algorithms have taken over much of the heavy lifting. A team of more than 2,000 workers in Sri Lanka pores over hundreds of thousands of miles’ worth of video collected from open roads. (“Ground truth,” as it’s known in the AI world.) They draw boxes on their computer screens around any item of consequence—a car, truck, bike, pedestrian—and label it. Next, those data are fed into a neural network, which is then trained so it can reliably identify similar items it hasn’t seen before, even if they’re coming at you fast on a freeway.

The team’s big challenge is to improve its software’s ability to read road signs and other cues to identify “free space”—parts of the road that are open to the car. Ground-truth

workers annotate signs, sidewalks, bike lanes, and other areas off-limits to vehicles, then draw on the road to show where the car can and can’t go. Each scene has so many elements that annotating a minutes-long clip can take all day. “It’s like writing a big vocabulary book,” says Gaby Hayon, a top Mobileye research executive.

Drivers, Hayon says, unconsciously “compose a story” that includes not just data (car, right lane, 45 mph) but a sense of past and future. You can guess the car is in the right lane and slowing down because it’s going to make a right turn, meaning it’s likely to slow down a lot more and may require you to hit the brakes or try to swerve around it. The challenge is to feed in enough data to make sure the computer can make the same guess. “What we’re doing with the sensing models,” Hayon says, “is building words for your stories.”

On a warm spring day in Jerusalem, a Mobileye prototype car called Clara flicked on her left turn signal and prepared to merge onto Jerusalem’s main north-south freeway, Begin Boulevard. Our backup driver for the trip was Shai Shalev-Shwartz, the vice president for technology who wrote the software that takes what the computer sees and decides what to do with it. Shashua rode in the back seat.

“Merges and lane changes are very, very complicated,” Shalev-Shwartz warned as cars and trucks sailed past, packed close together at 50 mph. It would have been a difficult situation for a human, and it seemed hard to see how Clara would find her way in. “We are waiting for the right moment,” Shalev-Shwartz said, then paused. “Hopefully.”

Paradoxically, driverless cars present a threat on top of the prospect that they’ll transform into robotic death machines. They may simply bore us to death. Carmakers, anxious to avoid collisions, have designed their algorithms to be extra careful, which means AI drivers tend to take a lot longer than humans, and trips may include lots of sudden stops. In late April, a video posted to Twitter showed a Waymo car struggling to enter a freeway. The modified Chrysler minivan signaled, slowed to a crawl, then, looking as if a flustered teenager were driving it, ambled straight through to the exit ramp and off the highway.

“One of these vehicles driving in Mountain View is nice,” Shashua said, referring to Google’s hometown and the site of Waymo’s early testing. “Ten thousand will likely obstruct traffic. Society won’t want these vehicles, not because they’re not safe, but because they are not useful.” A Waymo spokesman notes that it’s the only company testing cars without safety drivers.

Human driving instructors and our parents try to convey the importance of decisiveness in negotiating the highway, but there are situations in which no amount of coaching can help. Consider, for instance, the double merge, which happens when two highways converge. Since cars are merging from opposite sides, there’s no clear right of way or prescribed set of rules to dictate what you should do other than survive. Drivers survey the road and use their intuition to figure things out.

“There’s a big difference between something that kind of works and something that works”



This makes it a good application for deep learning. Engineers teaching AI software can train it to run through many possible paths in a scenario to figure out what's most effective. The approach is gamelike—developers speak about giving “rewards” to the computer as its algorithms improve—and has been most successful navigating systems with well-defined rules and a limited universe of possibilities. In training its bot to beat the world's top player in the Chinese board game Go, Alphabet's DeepMind developed strategies no human players ever tried.

Games, however, are more forgiving than the road. In 2016, Shalev-Shwartz fed a computer 100,000 or so double-merge scenarios, and out came a function to handle double merges. Great, he thought, and set up a simulation to test his function overnight. The first morning, there was only one accident after about 100,000 more simulations. He tweaked the rewards so the accident would no longer occur and went home. The next morning, he found another virtual crash. He fixed that, and there was a third accident, then a fourth and a fifth.

This was especially disturbing because simulations are inherently more predictable than real-world driving. What would happen if this algorithm were let loose on a real highway? “Deep learning generalizes very nicely on typical data that is similar to what it has seen,” Shalev-Shwartz says. “It has a really hard time generalizing rules based on things that it has never seen.”

In all likelihood, it will be the weird, one-in-a-million corner cases that cause accidents in a world with lots of autonomous cars. A driver on LSD, for example, or a kangaroo on the road, or a woman walking a bike where a computer doesn't expect her. According to a report by the tech news site The Information, Uber's software correctly identified Herzberg as a pedestrian, but failed to stop the car.

Citing anonymous sources, the article said executives believed they'd improperly “tuned” their algorithms, a fancy way to say the software didn't work and no one is sure why. (A spokeswoman says the company is conducting a “top-to-bottom safety review.”) That's because of what's known as the black box problem. An AI algorithm can't tell us why it picked a particular approach to a given problem, nor can it tell the people who programmed it why a particular corner case caused a model to fail. “The guarantees of machine learning are always statistical guarantees,” Shalev-Shwartz explains. “They tell you that with a probability of 99 percent, the function that you learned will be close to the optimal function.” Inevitably, in the fraction of a percent, and even with the best systems, there will be fatalities.

Mobileye knows that better than almost anyone. In 2013, Shashua struck a deal with Tesla CEO Elon Musk to install an advanced version of its driver-assist system in every new Tesla, enabling the cars to steer themselves on highways. Musk wasn't the first carmaker Shashua pitched, but he was the first yes. “Elon accelerated things,” Shashua said after the Clara demo. “He pushed everything fast.”

To make sure Tesla didn't oversell the system's capabilities,

Shashua visited Musk at his factory in Fremont, Calif., and urged him to make sure the system would force customers to keep their hands on the wheel. “His whole engineering team was waiting for me because they couldn't talk to him,” Shashua says, suggesting that Musk had been unpersuadable on the issue. “So I went there, I had the meeting, and I convinced him. He promised it wouldn't be hands-free.”

Two months later, in late 2015, Musk unveiled Tesla Autopilot. He told customers to keep their hands on the wheel and that they would be responsible in the event of a crash. But nobody listened, and Tesla's cars didn't always stop them. Within months, YouTube was full of videos of Tesla drivers falling asleep, playing Jenga, or riding in the back seat.

Then in May 2016, a Tesla driver switched on Autopilot while on a highway in Florida. A truck coming from the opposite direction turned in front of him and its white body, blending with the sunlight, confused the Mobileye camera. It failed to stop the car. The driver, who apparently had his eyes off the road, hit the truck at almost 75 mph, and died instantly.

Tesla has since tweaked its Autopilot system so drivers who ignore warnings to keep their hands on the wheel can no longer use the feature for the remainder of the drive. At the time, Tesla mostly blamed the Mobileye camera. That angered Shashua because the system had never been designed to detect crossing traffic. Soon after, he says, he canceled the partnership. For Tesla, he says, automation is “just a story. But for us, it's all we can do.” Tesla declined to comment for this story but has denied that it marketed a hands-free system and said it, not Mobileye, initiated the breakup.

The incident helped prompt Mobileye to create a fundamental set of do's and don'ts that supersedes its software's moment-to-moment decisions, an effort to make sure cars don't do anything reckless even if they haven't encountered a situation before. The company's Responsibility-Sensitive Safety, which it first proposed in a research paper last year, tries to formalize the basic rules of the road that human drivers internalize with experience, such as what constitutes a safe following distance.

Mobileye calls this a mathematical model to guarantee safety, but, as Shalev-Shwartz says, no system is guaranteed to be crash-free. Really, the Responsibility-Sensitive Safety system is meant to allow society to adjudicate blame when a crash happens. Did the car's sensors screw up, or was the fault the driver's? The answer is critical for the driverless industry's survival.

Of course, the other half of the equation is making the cars actually drive themselves. After we attempted to enter Begin Boulevard, Clara seemed ready to move, but she jerked us back into our lane as a car to the left accelerated and cut us off. She quickly recovered, edging almost onto the lane line—a maneuver designed to mimic the behavior of a human driver—then made the change with a forceful acceleration. A few seconds later, the traffic slowed and Clara hit her brakes, hard. “It's not Mountain View,” Shashua cracked from the back seat. “We are in survival mode here.” **E**