F28

NOT Too Far FROM THE TREE

Farmers are redesigning their orchards to make them easier for workers to pick. That makes them perfect for the first wave of robot harvesters.

BY JOHN H. TIBBETTS

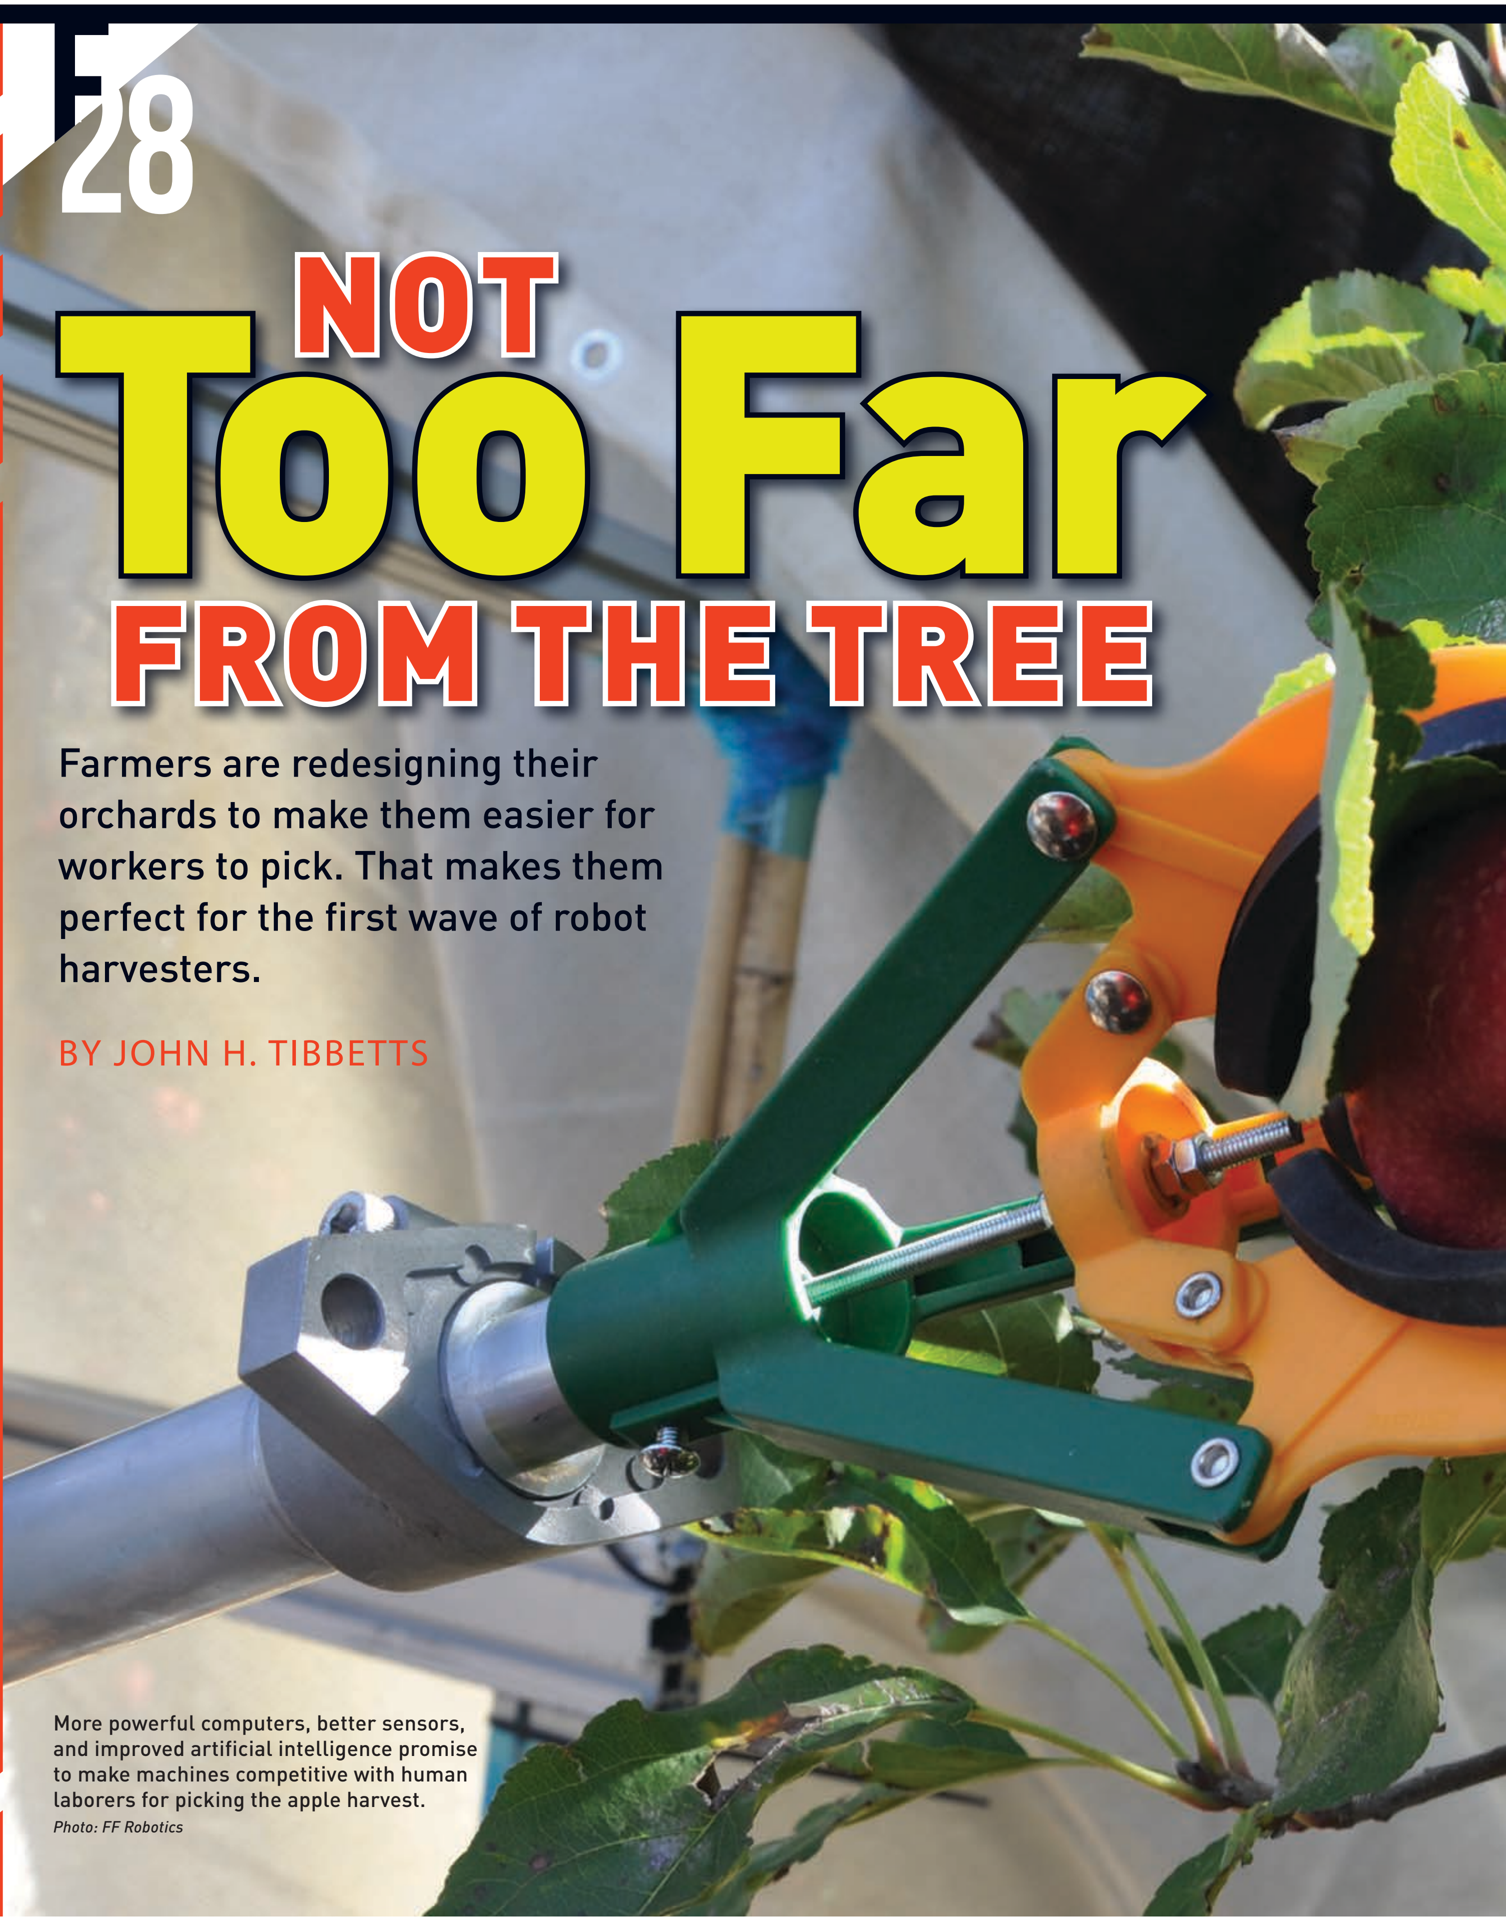

More powerful computers, better sensors, and improved artificial intelligence promise to make machines competitive with human laborers for picking the apple harvest.
Photo: FF Robotics

1/17/2018 2:15:28 PM

Exhausted by his long days picking apples, Avi Kahani looked forward to his life as an outdated technology.

This was in the early 1980s, when Kahani managed his communal farm's orchards on Israeli hills overlooking Lebanon. Every year, kibbutz members pitched in to pick fruit. And each year, they raced to harvest the apples while they were ripe enough to command the highest prices.

Kahani, who would soon leave to study mechanical engineering, could already see the first wave of industrial robots sweeping through industry. If robots could replace workers in factories, he reasoned, why not laborers in orchards?

"We thought robots would be coming soon," Kahani said, laughing at the memory. "Unfortunately, apple-picking robots have taken a lot longer than we expected."

Today, Kahani is making that vision a reality. In 2013, Kahani and Gad Kober cofounded Israel-based FFRobotics, one of two companies racing to commercialize the world's first mechanical apple picker. Abundant Robotics, based in Hayward, Calif., is the other.

Attached to a slowly moving tractor, FFRobotics' arm needs only two degrees of freedom to pick most apples.
Photo: FF Robotics

FFRobotics' robot sits in the back of a human-driven truck. As it moves through the orchard, its four arms reach out to grasp apples in their three-pronged hands. Once they have the apple, the effectors rotate to snap it off the branch and then drop it in a container. The system could replace an army of pickers, and that is the point.

"What's driving robotics in fruit picking is the same in Israel, Europe, and the United States," Kahani, the company's CEO, said. "Seasonal workers are hard to find, and their wages are rising."

Each apple for the world's lucrative fresh market is still picked by hand. In the United States alone, 40,000 people labor in apple orchards and packing sheds each harvest season. Temporary migrant laborers make up about one-third of the peak-season labor force, arriving on H-2A guest visas. An unknown number of seasonal workers are undocumented.

This past year, as immigration—both legal and illegal—became a highly charged political issue, many farms and orchards had trouble finding enough seasonal workers. Without those workers, unpicked apples will rot on the branch.

"Growers feel powerless about the current labor market and guest worker system," said Karen Marie Lewis, a Washington State University tree fruit horticulturist. "The state of Washington is highly dependent on visa and immigration law, which can determine whether we get the full complement of our workforce."

Even before worker shortages, wages for seasonal workers were rising. The guest worker program requires growers to pay migrants the same prevailing wage they pay domestic laborers, which has been increasing, Lewis said. Between labor's rising cost and its uncertainty, growers are eager for feasible alternatives.

FFRobotics is developing that alternative. After three years of field-testing in Israeli orchards, Kahani and his colleagues plan to test their apple-picking robot on Washington's 2018 harvest, which runs from mid-August through mid-November. The state accounts for nearly three out of every five apples grown in the United States.

Fruit picking is hard, repetitive work, just the type of semi-skilled labor typically replaced by automation in manufacturing. But to perform their tasks reliably, factory robots need a highly structured environment. They require

1/17/2018 2:15:32 PM

precisely oriented parts delivered to specific locations, and consistent lighting, colors, and background conditions to avoid confusing their vision systems.

Traditional orchards are anything but that. Their canopy of leaves and branches presents a shifting collage of confounding light and shadow as the sun moves across the sky, altering the appearance of each apple. A robotic picker must distinguish a ripe fruit from less mature apples, and from tree limbs, branches, twigs, and leaves.

It must also do this under floodlights, since the only way for a grower to afford a robotic picker would be to run it 24/7.

There is little room for error. Apples are among the most easily bruised fruits. If a robotic picker misses and an apple falls, it may be ruined for the lucrative fresh market.

Fortunately, farmers have been making orchards robot-friendly for years—and they never even knew it.

MODERN ORCHARDS

A traditional apple tree has a thick trunk and branches that stick out randomly in every direction. The trunk and branches soak up a lot of water and energy before they get to the fruit.

So, geneticists developed dwarf varieties with narrower trunks and branches that deliver a greater percentage of water and nutrients to the apples. This produced more apples per acre.

Orchard managers train these dwarf varieties to make them more accessible. They prune and thin them to establish wide, shallow canopies in nearly vertical planes or V-shapes, so each apple grows within the reach of a human arm. They support weaker trunks with wood or concrete, and keep branches from sagging with wires and trellises.

They plant rows up to 1,800 feet long, so human pickers can move efficiently along their "fruit walls." These apple trees stick out about a foot from the trunk on each side of the row, with fewer limbs and shadows to hide or obscure fruit. This makes it easier for a laborer to identify, reach, and pick each apple.

"In Washington, we are building repeatable, predictable, and accessible canopies for human pickers and now we are primed for engineering solutions, including robotic harvesting," said Lewis. "A human-friendly orchard is a robot-friendly one."

Modern orchard designs also allow engineers to build simpler apple-picking systems, added Amir Degani, founder of the Civil, Environmental, and Agricultural Robotics (CEAR) Lab at Technion-Israel Institute of Technology in Haifa. Degani advised with FFRobotics on developing its robotic arm.

"To harvest apples on a natural tree, a robot arm would require six to seven degrees of freedom, and that's too expensive," Degani said. "With a simpler tree, you can actually use a 3-degree-of-freedom robot that FFRobotics is using."

FFRobotics' multiple arms, for example, sit on a truck that moves slowly through the orchard, eliminating one degree of freedom. To reach an apple, an arm moves up-and-down and forward-and-back, two degrees of freedom. Once it grasps the fruit, the hand rotates (a third degree of freedom) to detach it from the tree.

FFRobotics' only challenger to date, Abundant Robotics, also deploys multiple robotic arms guided by algorithms and cameras. Each arm has a vacuum tube end-effector that sucks fruit off trees and drops them into bins.

Abundant, which spun out from the independent lab SRI in 2016, field-tested its robot in Washington and Australia in 2017. The company received $10 million in funding led by Google's investment arm in May 2017, bringing its total financial backing to $12 million.

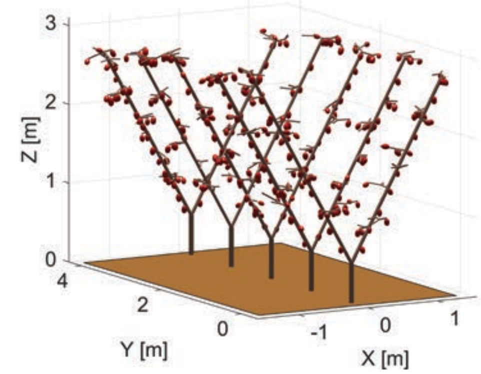

Left, farmers train modern orchards to make apples more accessible to pickers. Right, an illustration of how a robot sees apples on a tree.

Photos: FF Robotics

The Washington Tree Fruit Research Commission was an early supporter of Abundant Robot's picker, providing $500,000 in seed funds. FFRobotics also has a proposal before the commission for funding.

"With or without that support, we will be in Washington in 2018," Kahani vowed.

Washington will be ready. A grower typically sets aside 5 percent of an orchard's acreage annually to plant new trees, which produce apples within three years. At that pace, transitioning an entire orchard to new canopies would require two decades. In Washington, however, some growers are moving faster.

"The industry is spending a boatload of money to plant new acres or transition older acreage to robot-ready canopies," Lewis said.

The combination of labor woes and successful robot tests have raised interest among orchard owners across the country, including Rod Farrow, who grows 28 apple varieties on Lamont Fruit Farm's 520 acres in Waterport, N.Y.

"Even a couple of years ago, we thought we would not see a commercial apple picker in our lifetimes, but things are moving very quickly," he said. "Robot systems can pick an apple and put it in a bin at an equivalent rate as our workforce with very little damage to the apples. I believe we will have a commercial version working in our own orchards in five years."

DEEP LEARNING

Before locating and grasping an apple, a robot picker must recognize it as one—despite variations in light, dust, wind, and a background of leaves and branches.

Most fruit-picking systems depend on deep learning software to recognize apples, said Manoj Karkee, an associate professor at Washington State University's Agricultural Automation and Robotics Lab. A deep learning model consists of interconnected layers of computation. It is a type of neural network, named because it is a simplified imitation of the human brain's system of neurons and synapses.

"Deep learning is the go-to method for working with these kinds of robotic fruit-picking problems," Karkee said. "It gives us so much accuracy and robustness in detecting and localizing objects in complex environments."

To train a deep learning network to recognize apples takes two things. The first is computing power, which arrived on the scene in the form of graphics processing units (GPUs) used to rapidly crunch numbers for realistic, rapidly changing video game displays. The second is examples of apples, which are readily available from Internet's vast trove of images and videos.

To understand how deep learning works, consider how Facebook tags friends in photographs. Its network's first computational layers might detect simple things like image elements, such as edges. The next layers might bring together these edges to form parts of faces—noses and eyes, for instance—and further layers might weave those features together to identify the person. Each successive layer of the model builds on the knowledge of the previous layers. The model learns how to recognize features—color, texture, shape—that the programmers tell it to look for.

FFRobotics trained its deep learning model similarly, feeding it labeled digital images of apples and tree backgrounds. The model learned by trial and error to correlate an apple's "face"—its specific color, shape and texture—with its corresponding image label. It then trained the model on images of tree trunks, branches, twigs, and leaves.

Over time, the model learned to get it right. When it consistently misidentified a blob of leaves as an apple (or vice versa), the programmers tweaked the algorithms until they were more likely to identify an apple accurately next time. Then they asked the model to identify features in unlabeled digital images that it had never seen before, refining the algorithms as they moved forward.

Even with this training, FFRobotics preps the robot before it enters an unfamiliar orchard.

"We always take a lot of pictures in the same orchard where the robot will pick," Kahani said. "When we go to Washington for field tests, we will take additional pictures in their orchards because even small changes—such as the angle of the sun from Israel to the Pacific Northwest—can change the model's results. Training the model is an ongoing process."

Not only do apple tree varieties and cultivation practices vary, but modern orchards often cultivate premium apples in shades of red, pink, green, and gold.

Abundant Robotics uses a vacuum effector to pull apples off the tree and deposit them in a container.
Photo: Abundant Robotics

"To gain maximum efficiency, you must train the system with pictures of the fruit variety in the orchard you want to pick," Kahani said. "Then we classify the apples based on the individual grower's criteria of apple size and color."

A member of the FFRobotics' team also scans each tree's critical points and junctions with an encoder, digitizing the tree. The data goes into Matlab to create a 3-D model of each tree.

The four-armed robot, which rides on a tractor driven by a human, uses a single camera to capture an image of an individual apple. The 3-D model identifies the fruit's location and orientation and differentiates the fruit from the rest of the tree based on the deep learning algorithm. A controller then directs the robotic arm to that location.

OPEN OR CLOSED

FFRobotics is still struggling with whether to go with open- or closed-loop controller. The open-loop system recognizes a specific fruit and sends the gripper to that location. If a strong wind moves the apple left or right, the gripper does not follow.

The closed-loop system tracks the movement of the fruit by distinctive points on the apple's face as guides and adjusts the arm as it moves closer to the apple. While closed-loop systems are more effective, they are also too expensive, Kahani said.

"From our experience, the open loop system with this robot and this gripper is working very well even in windy situations," Kahani said. "But we are thinking and working all the time to improve the system, so we may change it."

The three-fingered gripper at the end of the robotic arm has a single motor that grips the apple and rotates it about 90 degrees to detach it from the tree. The arm retracts several inches and drops the apple into a container.

The robot's vision system functions best when apples are clearly visible or only partially hidden. Apples blocked by branches would be too difficult to pick, so the robot leaves them for a human to pick later. The machine can pick 80 to 90 percent of a crop, depending on the tree design, said Kahani, shaving harvesting costs by about 25 percent.

"Even picking 70 percent of a crop would make it financially viable," said Karkee. "A manual crew may have to follow up by picking the other 30 percent, but that would take care of seasonal labor demands at harvest."

Because farmers would use an apple picker only 80 days per year, Kahani would like to adapt the mechanism to other fruits. He envisions rewriting the algorithm and reengineering the gripper for peaches or citrus fruits. The robot would then become a migrant laborer, working the fall harvest in the north and heading south for the winter harvest in winter.

Some humans are likely to join them, however, Farrow said. Robots, at least the early versions, cannot reach every apple and they are prone to breakdown.

"The federal guest worker program is a nightmare to work with, but if you give up those workers to use these robots, then those people aren't available to you anymore," Farrow said.

Farmers will still need laborers, but Farrow expects most growers to embrace robot pickers as they transition to new canopies. "Large operations that are already transitioning will adopt robots quickly. But you need at least couple hundred acres with the right canopy to justify making it work financially, and many operations aren't there yet."

Maybe not, but they are certainly a lot closer than they were when Kahani was picking apples as a young man. ME

JOHN H. TIBBETTS is a freelance writer based in Charleston, S.C.