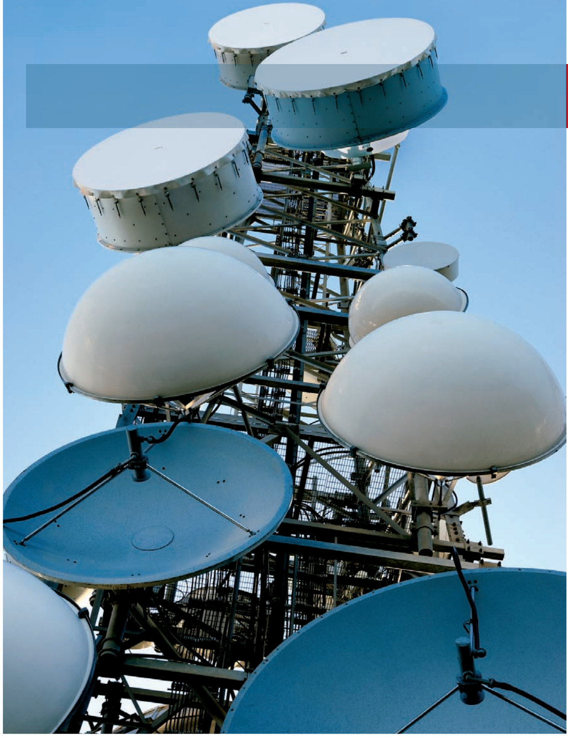

IN THE NEWS

Senior News Editor: **Lee Garber**, lgarber@computer.org

Drones Learn to Navigate, Avoid Obstacles

George Lawton

Drone technology has garnered attention because of its military use for surveillance or attack.

However, considerable research on drones—also called unmanned aerial vehicles (UAVs)—is taking place in universities and other research settings, as people hope to use flying robots for civilian purposes such as surveillance, search and rescue, infrastructure monitoring, and aerial photography.

There is great interest in this, despite the privacy concerns.

Stanford University doctoral student Hmong Huang said, "For the most part, a lot of [drones'] control problems have been solved. The hard part is sensing information about space and figuring out where the obstacles are."

"This," he said, "is why most UAVs are remotely piloted. It is not feasible to let them loose in an unstructured environment where there are lots of obstacles around."

To overcome this, considerable drone research is taking place on environmental sensing, as well as automated navigation and obstacle avoidance.

Controlling Drones

UAVs use a variety of sensors—including cameras, lasers, and gyroscopes—to identify their location and flight orientation.

Humans frequently remotely fly drones. In some cases, though, controllers—either onboard or on a central computer—operate the UAVs based on flight plans or destination information.

Modern drones contain two levels of flight controllers, noted Cornell University assistant professor Ashutosh Saxena.

A high-level controller specifies the general direction in which a person wants to send the vehicle. A low-level controller translates the high-level goal into a series of instructions for each onboard motor.

There is considerable room for improvement in onboard controllers and sensors, according to Maxim Likhachev, an assistant professor at Carnegie Mellon University (CMU).

Better Image Recognition

Cornell University researchers are working on better image-recognition systems capable of running on small UAVs.

Saxena explained, "Most of the past work has been on controllers that can do amazing things in free environments. However, when there are obstacles in the environment, such as when the UAV has to fly into a building or tree canopy, the vision-based controller becomes important."

He is researching robotic perception for inferring 3D depth from onboard cameras' 2D images so that a controller could keep a UAV from crashing into objects. The algorithms his team developed are compact enough to run natively on the resource-constrained UAVs.

The team is also working on cameras, other types of sensors, and controllers that can run completely onboard with no need to communicate with external elements.

1541-1672/13/$31.00 © 2013 IEEE
Published by the IEEE Computer Society

This is important because if UAVs must communicate with external components, they can't travel out of transmission range, which limits their usefulness. And if they don't have their own cameras, they could work only in areas where there are already cameras to observe their movements.

The Cornell researchers are now designing fast crash-avoidance algorithms that run directly on UAVs. They designed their algorithms to be compact enough to run on the type of little processor with minimal RAM that a small drone with limited battery power could support. Adding bigger processors and more RAM would either reduce existing UAVs' flight time or require bigger, more expensive drones.

To generate a three-dimensional model of an area using 2D data from a camera, the scientists utilized Saxena's Make3D algorithm in conjunction with Markov random fields, often used to model various low- to mid-level image-processing tasks.

The researchers have demonstrated their system on a simulator and are now working on real-world testing.

The Cornell techniques can work with any camera that supports greater than 640 × 480 resolution, which is necessary to provide sufficiently precise depth information.

Better Robotic Planning

CMU researchers have focused on developing more efficient planning algorithms for drones using data from laser-based range sensors. The algorithms plan a drone's movements by translating high-level goals into a series of rotor commands.

Carnegie Mellon's Likhachev said processing laser data is more efficient than using image-recognition techniques on camera information. The laser data includes depth information, while camera-based techniques must take time and effort to compute depth.

Greater efficiency means that the system needs less processing power to handle the data. This lets the system run on smaller drones with less expensive CPUs.

The CMU team's algorithms plot a drone's flight path and movements and then regenerate new plans on the fly in response to environmental changes such as obstacles encountered.

The researchers have also been figuring out how to land UAVs in environments with trees or other obstacles and on different kinds of objects. They discovered that various types of objects, depending on their texture and orientation, require different landing patterns and algorithms.

A harder, as-yet-unsolved problem is having drones respond quickly to changes in environments that include other moving objects that might also have to change their trajectories.

To help with this, the CMU team is combining data from multiple sensor types—including a laser range finder—that detect other objects. They also use an inertial measurement unit and compass, which incorporate data about the drones' current movements into the rotor-control algorithm.

The researchers are also developing graph-based algorithms to calculate and update plans in real time. This technique starts by showing elements of a drone's flight plan as a graph with axes representing space, time, and rotor power.

The algorithms search through possible trajectories and identify the one that avoids obstacles while using the least energy.

The efficient graph-based approach can deliver robust results on a smaller processor.

The CMU scientists want to develop algorithms that run in real time on UAVs, which must make corrections and generate new flight plans in hundreds of milliseconds, Likhachev explained. Algorithms that must run on an external computer would require a network connection that could slow these operations.

The researchers have been working on small systems like microquadrocopters and larger ones such as helicopters.

Using UAVs Strategically

Much of the UAV-controller work has focused on the tactical aspects of moving a drone to a target location, which could be necessary for tasks such as search-and-rescue missions.

Stanford researchers are studying the strategic aspects of positioning UAVs to support larger goals.

They looked at ways to guide flying robots to the best vantage point to help people playing capture the flag. In this game, two teams have a flag or some other marker at their headquarters. Each squad tries to capture the other's flag and take the marker to its own headquarters.

The Stanford team developed algorithms that let UAVs automatically identify the places that humans have trouble seeing from the ground and the overhead locations where they need to be to observe these places. The algorithms work with a digital map that represents obstacles as ones and empty spaces as zeros.

A probabilistic interference algorithm identifies points behind buildings and in other locations that one team could use to launch attacks but that the other team might not be able to see, Stanford's Huang explained. The algorithm cross-references data about the strategic importance of various locations on the 2D map with data about whether spots are visible from one another.

A centralized computer aggregates players' location data, collected by UAVs via GPS information from participants' smartphones, and then runs the algorithms, as Figure 1 shows.

A similar approach might one day enable UAVs to find the best vantage points during search-and-rescue operations.

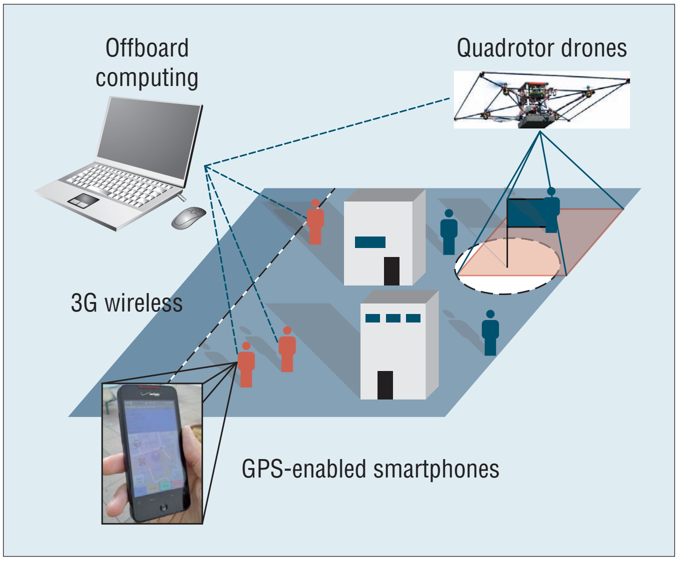

Figure 1. Stanford University researchers are developing technology that could let drones help capture-the-flag game players. The flying robots could identify important places in the playing area that humans can't see from the ground and the overhead locations from which the aircraft could observe them. An offboard computer collects players' GPS-based location data from their smartphones and then runs the algorithms.

Paving the Way for Civilian UAVs

Considerable technical and publicity work remains before UAVs overcome obstacles to widespread adoption.

"For most people," Stanford's Huang said, "thinking about surveillance drones flying overhead and taking pictures is creepy [because of] security and privacy concerns. How we resolve our attitudes towards them will play a role in [UAVs'] deployment."

In domains in which privacy is less of a concern, such as search-and-rescue operations, the technology must improve to allow better control of groups of drones.

Saxena said, "If you want a robot that can go into a disaster zone like a flood, you cannot control each UAV manually. You want the UAV to be able to find missing people or broken infrastructure [on its own]. The other thing we need is machine-learning algorithms to train the vision systems to recognize and improve their interactions with new objects."

Saxena predicted that future research will help improve UAV commercialization via the development of better obstacle-avoidance systems.

He said he is also looking forward to the development of standardized platforms that, like operating systems, could run on different types of robots.

Saxena explained, "This would make it easier to develop applications on one platform that could run on multiple UAVs."

Application Uses AI to Identify Rough Drawings

George Lawton

An international team has developed a computer program that recognizes rough sketches, even as they're being drawn.

Brown University and Technical University of Berlin researchers designed the WhatsMySketch application for the iPhone.

This is one of the first research projects to develop algorithms for recognizing sketches of objects, said Technical University of Berlin doctoral student Mathias Eitz.

It's also the first computer application that enables semantic understanding of abstract sketches by automatically classifying them, the scientists say.

The advance could clear the way for vastly improved sketch-based application interfaces and search engines.

Sketch Recognition

Some applications can identify drawings if they are realistic, such as programs that can match police sketches with mug shots.

Others, such as those using sketch-based retrieval algorithms, could find photos with shapes similar to those in an input sketch but couldn't identify what type of object it represents.

The Brown University and Technical University of Berlin researchers' work is the first to identify highly informal or abstract sketches.

Recognizing sketches is different from identifying photographs because there hasn't been a repository of labeled drawings for applications to learn from, said the University of Berlin's Eitz.

Also, while photos typically look largely the same, no matter who takes them, people draw the same object using different degrees of realism and styles. For example, a person might sketch a rabbit with just big ears, buck teeth, and a cotton-ball tail, which wouldn't look like an actual bunny.

Thus, computers have had trouble identifying informal sketches, noted Brown University assistant professor James Hays.

WhatsMySketch

The researchers started by identifying popular categories of objects. To do this, they used a dataset of labeled photographs called LabelMe, which includes metadata on many potential sketch subjects. The scientists selected the 250 most common items and hired humans—using Amazon's Mechanical Turk crowdsourcing service—to sketch them.

They created about 20,000 sketches of the various items and added metadata about them. They then fed the sketches and metadata into a system that used machine-learning and image-recognition technologies to learn the appropriate object category for each sketch.

WhatsMySketch needs at least 100 different sketches of an object to accurately identify new drawings of the item.

Brown's Hays said, "When giving training data to a computer, there are two big decisions to make: How do I computationally represent the data, and what learning algorithm do I use to leverage the data?"

The scientists decided to represent a sketch as a collection of statistics, correlating shapes in a sketch with object categories.

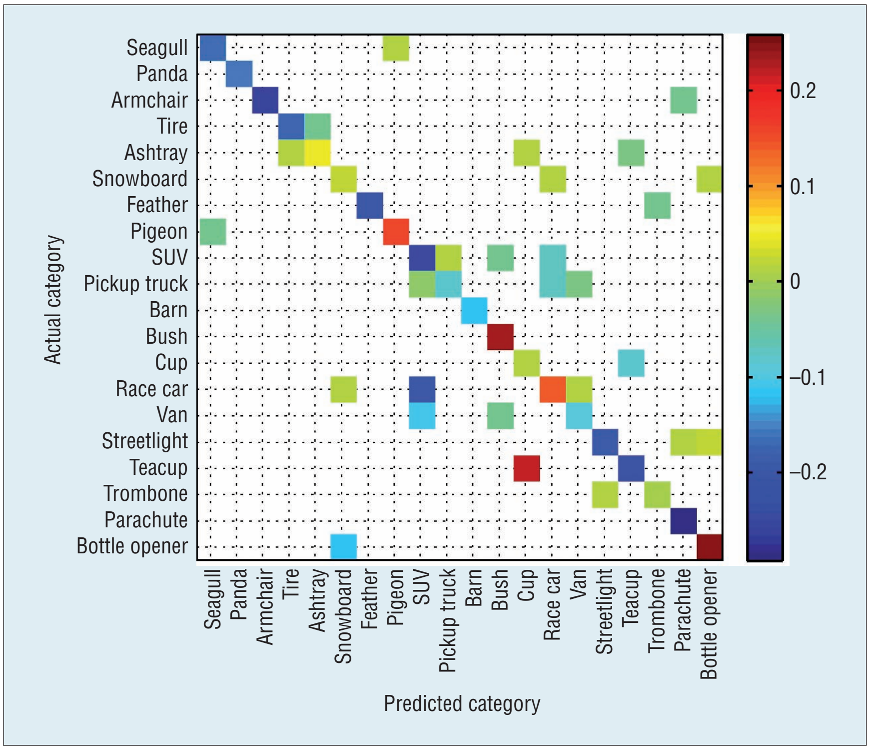

Figure 2. This matrix shows the difference between how well humans and the WhatsMySketch application recognized selected images as they were being drawn. The red, orange, and yellow squares indicate where humans did better. The other colored squares show where the application performed better.

They also used off-the-shelf machine-learning algorithms—including one-versus-all support vector machines, which try to predict whether an image is in a specific grouping—with radial basis function kernels to help the machine-learning algorithm evaluate potential object categories for sketches.

This enabled the team to develop an application that could identify drawings of the 250 objects even as they were being sketched.

After each pencil stroke, the system extracts features from the drawing. It then uses support vector-machine model classifier algorithms to classify the sketch as to the type of object it appears to represent, and displays the 20 highest-scoring possible object categories, all within 100 milliseconds.

The initial WhatsMySketch version identified objects with 56 percent accuracy, compared to 73 percent for humans. Figure 2 compares the application's and humans' performance for selected images.

In a second experiment, the researchers looked at 3D shape retrieval. In this approach, a user inputs a sketch, and a system retrieves an image of the object it represents in a 3D file format suitable for uses such as CAD drawings. For this, the scientists trained their system using 1,814 sketches in 250 categories from an existing dataset of 3D shapes.

This work led to development of a compact targeted-feature-transform algorithm that recognizes sketches made by people with different drawing styles.

Despite all this progress, Hays noted, sketch recognition is still difficult when there are many ways to draw the same object or when different objects in multiple categories look very similar.

Uses

Sketch-based applications haven't been widely used because they haven't been able to identify drawings well, said Hays.

If sufficiently accurate, WhatsMySketch could enable sketch-based interfaces that could improve computer accessibility for various people, including those who can't read.

It could also make searching for information based on a sketch easier and more intuitive. And the system could identify the category of a sketched object and then help users search for and retrieve photographs of similar items based on their drawings.

The software could also make applications more useful across multiple languages because it isn't language based.

Future Work

The researchers, who have no immediate plans to commercialize their technology, are planning additional work on expanding their sketch database to include more categories of drawn objects.

They also want to look at various other areas related to sketching.

For example, they want to explore sketch synthesis, which lets a computer automatically generate a drawing of an object that could be immediately recognizable to a human.

They also want to look at sketch beautification, which lets an application work with a human to simplify and improve a sketch.

The University of Berlin's Eitz said, "We hope that better computational understanding of sketches will lead to better computer accessibility. Virtually everybody is able to sketch a face or recognize a sketched face. The ability to write and read, which today are still the standard ways of communicating with computers, is much less widespread."

"If computers were to understand sketches as we do," he added, "sketching would give a much larger audience access to the data that has been gathered digitally over the last decades."

Intelligent Program Wins at Fantasy Soccer

George Lawton

An academic research team has developed an intelligent system that chose a successful fantasy soccer team better than 99 percent of the 2 million people whose results were used for comparison.

In fantasy soccer, participants create imaginary teams from players in real leagues. At the end of each round of games in the real leagues, each fantasy team gets points based on the aggregate performance of its own players.

The University of Southampton and Technical University of Crete scientists say their research could also help in organizing work, sports, or other types of teams based on participants' individual and team-oriented skills.

How It Works

University of Southampton postgraduate researcher Tim Mathews said the new research predicted how well an individual soccer player would perform on his own merits and how well a team would do with the person playing for it.

According to Matthews, he and his colleagues encountered three main challenges.

First, they required a statistical model able to produce reasonable predictions of a player's point-scoring potential and the team's performance with him on it. For this, they used a Bayesian belief model, which organized game statistics to let existing player-selection algorithms evaluate the performance of the player, his real team, and the fantasy squad that might select him.

This increased the odds that the algorithms would pick the best player combination for a fantasy soccer team.

Next, the researchers created squads of footballers, maximizing their total point-scoring ability subject to a number of challenging constraints, including the rules governing player choices and trades.

For this, they used a multidimensional knapsack-packing approach, designed to make the best player selections without breaking any constraints.

The scientists also wanted to select teams that would perform well as far into the future as possible. They used a belief-state Markov Decision process, a framework that assesses the long-term value of player selections by solving a dynamic programming equation.

In addition, they employed Bayesian Q-learning techniques to take into account the uncertainty about various selections' quality and to hone in on the best choices. As multiple player choices are explored, the system becomes more certain of their expected long-term reward, explained Technical University of Crete assistant professor Georgios Chalkiadakis.

"This is important in a domain as unpredictable and dynamic as football," he noted.

The researchers' approach also utilized the *value of perfect information* concept, a mathematical framework that quickly explored various combinations of players to determine the likelihood that a particular combination would score well.

Looking Ahead

The scientists' approach has some weaknesses, the University of Southampton's Matthews explained. For example, he said, the statistical model makes numerous simplifications to make the algorithms easier to implement.

Nonetheless, the scientists say their technology is effective, and they plan to develop it into a Web service for fantasy-soccer participants, to further their research.

They also intend to combine their work with multiagent collaboration and coordination techniques. This would, for instance, help the gaming industry by improving the opponents that computers generate for humans to play against, said Matthews.

And, he said, their techniques would work well with other team-based sports for which there is a statistical model for players' point-scoring abilities.

The new algorithms, he added, could also help solve long-term planning problems commonly found in other fields such as crowdsourcing, as well as in organizing and scheduling emergency responders.

Similar to the process of selecting players for fantasy-soccer teams, people who operate crowdsourcing applications typically want to pick which individuals—from among many participants—will provide the best services that are being sought.

The same applies to supervisors who select the emergency personnel who have the capabilities and experience that will let them best respond to fires, disasters, and other emergencies, noted University of Southampton lecturer Sarvapali Ramchurn. ■