$50,000 to solve the most complicated puzzle ever attempted

November 17 2011

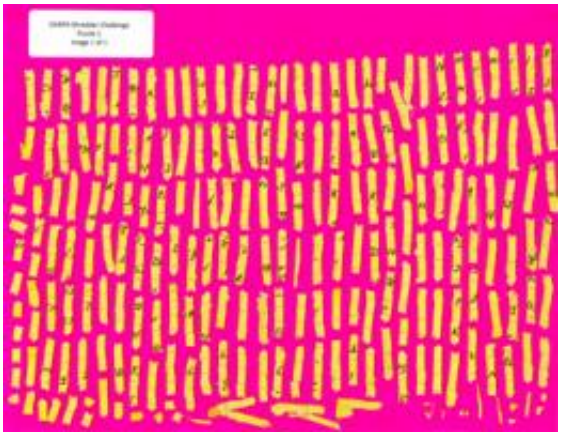

(PhysOrg.com) -- Every few years the Pentagon's Defense Advanced Research Projects Agency (DARPA) holds a public competition to stretch the outer limits of what technology can do. Two years ago they dispersed 10 large, red weather balloons at undisclosed locations across the U.S. The celebrated 2009 DARPA Network Challenge to find the balloons was solved in just nine hours by a team from MIT. Now, Manuel Cebrian, a member of that winning team, is aiming for a repeat win – only this time, the challenge is exponentially harder.

According to Dr. Cebrian – who is now a research scientist at the University of California, San Diego – instead of just looking for 10 balloons, the new DARPA Shredder Challenge's ultimate puzzle involves

piecing together roughly 10,000 pieces of different documents that have been shredded. “This is almost certainly the most challenging puzzle ever created,” said Cebrian. “A combinatorial number of possibilities makes the problem intractable by computer algorithms alone. A combination of crowd-sourcing and advanced computer-vision algorithms is necessary. This is exactly our approach.”

Crowdsourcing was also Cebrian’s approach with the winning MIT team, but he says that approach won’t be sufficient to win the challenge this time. He admits that it will take many more people in the “crowd” than was necessary to solve the 10-red-balloons hunt last time. “For that, we estimated that we needed a few thousand participants on our side to take up the manual challenge. To complete this new challenge, it could take as many as 100,000 people,” said Cebrian, who is in the Computer Science and Engineering department at UCSD. “We need a collaborative platform to move the puzzle pieces around, and we need computer vision to improve the partial solutions that humans achieve. It has to be a continuous human-computer-human feedback loop.”

Given the multidisciplinary nature of the Shredder Challenge, the core UCSD team reflects that variety: members include game theorist Andrea Vattani, Internet expert Karyn Benson, crypto-analyst Wilson Lian, network scientist Dan Ricketts, and computer vision researcher Sam Kwak. The latter will be running crowd-sourced partial solutions through advanced algorithms for clustering pieces and finding which pieces are likely to go together. This will make subsequent puzzle-solving easier for participants.

The UCSD team spent a lot of time thinking about how to create an incentive structure that will encourage participation. The resulting ‘recursive’ incentives are similar to those used by Cebrian’s team in 2009, particularly the referral-based crowd-sourcing. The insiders on the UCSD team will not get any of the prize money, even if the team wins: if

there is money left over after paying out all amounts to participants, it will go to charity.

If UCSD wins, the amount of money an individual will win is based on how many edges they connected in the puzzles. For every edge connected by the participant, he or she would receive $1. The person who recruited that participant would receive 50 cents, and the person above that recruiter would receive 25 cents. Depending on how many people are in the participant's 'family tree', a further fractional amount will accrue all the way up the tree.

"The rationale behind this payment scheme is that the main challenge involves fitting puzzle pieces together," said Cebrian. "However, it is unreasonable that one person would be able to complete all the puzzles by themselves, so recruiting players is essential to solving all the puzzles, especially as they get bigger. So we created a scheme that rewards the solvers, but the successful recruiters as well."

The core team, however, would get bragging rights if they win (and Cebrian knows all about the public attention lavished on winners of past DARPA challenges). All other referral-based participants have a financial stake in the outcome, especially if they have large social networks to tap. Even if they don't solve two edges themselves, they get credit for any of the people below them who do help solve parts of the puzzle. (For details, visit the UCSD Shredder Challenge website at shredder-challenge.ucsd.edu.)

If the UCSD-led team wins, the money will be split among the participants who helped reconstruct the shredded documents and the participants who recruited them. If you refer someone who then refers someone else in their social network, that third person gets the biggest prize per puzzle solved (since they solved that pair). Each of the people above them gets prize money too. In all cases, the prize money is

rewarded only if the team wins the competition. So there is a strong incentive for participants to go with the team that is best-equipped to go all the way.

The competition, announced in late October, involves five separate contests. The first two have been solved by a few teams, as they could be solved purely manually. The UCSD team chose to wait during those rounds, because the complexity – and the prize money – grow many times larger with the larger puzzles, which will certainly require some mix of manual assembly and computer vision. The team plans to solve the first two downloadable puzzles in little time – it can still gain points from this – and build up towards the 5th and conclusive puzzle.

The staggering of the complexity is purposeful. According to the Pentagon in announcing the Shredder Challenge, “DARPA hopes to gain new insight into which of these or other techniques are quicker and more efficient, and whether wide availability of high-resolution photography, communication and crowd-sourcing strategies offer unexpected advances.”

DARPA also settled on the Shredder Challenge because troops often confiscate the remnants of destroyed documents in war zones, and it’s too costly and time-consuming to reconstruct the documents. Said Dan Kaufman, Director of DARPA’s Information Innovation Office: “The ability to reconstruct shredded documents will potentially yield information that may save lives or offer critical information about an adversary’s plans.”

In general, DARPA’s Shredder Challenge goes far beyond battlefield tactics. The agency is not even sure that anyone can solve the final, 6,000-piece puzzle.

“The UCSD team is using the same financial recursive structure to

recruit people to collaboratively solve the problem. But we need more than that," said UCSD's Cebrian. "Our MIT approach won't be enough to solve this problem. We need a collaborative platform to move the puzzle pieces around. And we need computer vision to improve the partial solutions that humans achieve. A continuous human-computer-human feedback loop."

"Ultimately we need a lot of people, more than any experiment of this type before, to help us solve this problem," concluded Cebrian.

The crowd-sourcing task is tailor-made for potential participants with large, existing social networks. Even if the first participant doesn't try to solve any of the puzzle, they can still get some of the $50,000 benefit if their friends online do so. The more eyes the team has looking at the puzzle, the better. Even if you don't have time to put the shredded paper back together, perhaps your friends will have time – or your friends' friends.

Provided by University of California - San Diego