

System recruits learners to annotate videos

February 11 2015, by Larry Hardesty



A new system for crowdsourced video annotation could increase the educational value of instructional videos. Jose-Luis Olivares/MIT

Educational researchers have long held that presenting students with clear outlines of the material covered in lectures improves their retention.

Recent studies indicate that the same is true of online how-to videos, and in a paper being presented at the Association for Computing Machinery's Conference on Computer-Supported Cooperative Work and Social

Computing in March, researchers at MIT and Harvard University describe a new system that recruits viewers to create high-level conceptual outlines.

Blind reviews by experts in the topics covered by the videos indicated that the outlines produced by the new system were as good as, or better than, those produced by other experts.

The outlines also serve as navigation tools, so viewers already familiar with some of a video's content can skip ahead, while others can backtrack to review content they missed the first time around.

"That addresses one of the fundamental problems with videos," says Juho Kim, an MIT graduate student in electrical engineering and [computer science](#) and one of the paper's co-authors. "It's really hard to find the exact spots that you want to watch. You end up scrubbing on the timeline carefully and looking at thumbnails. And with educational videos, especially, it's really hard, because it's not that visually dynamic. So we thought that having this semantic information about the video really helps."

Kim is a member of the User Interface Design Group at MIT's Computer Science and Artificial Intelligence Laboratory, which is led by Rob Miller, a professor of computer science and engineering and another of the paper's co-authors. A major topic of research in Miller's group is the clever design of computer interfaces to harness the power of crowdsourcing, or distributing simple but time-consuming tasks among large numbers of paid or unpaid online volunteers.

Joining Kim and Miller on the paper are first author Sarah Weir, an undergraduate who worked on the project through the MIT Undergraduate Research Opportunities Program, and Krzysztof Gajos, an associate professor of computer science at Harvard University.

High-concept video

Several studies in the past five years, particularly those by Richard Catrambone, a psychologist at Georgia Tech, have demonstrated that accompanying how-to videos with step-by-step instructions improves learners' mastery of the concepts presented. But before beginning work on their crowdsourced video annotation systems, the MIT and Harvard researchers conducted their own user study.

They hand-annotated several video tutorials on the use of the graphics program Photoshop and presented the videos, either with or without the annotations, to study subjects. The subjects were then assigned a task that drew on their new skills, and the results were evaluated by Photoshop experts. The work of the subjects who'd watched the annotated videos scored higher with the experts, and the subjects themselves reported greater confidence in their abilities and satisfaction with the tutorials.

Last year, at the Association for Computing Machinery's Conference on Human Factors in Computing Systems, the researchers presented a system for distributing the video-annotation task among paid workers recruited through Amazon's Mechanical Turk crowdsourcing service. Their clever allocation and proofreading scheme got the cost of high-quality video annotation down to \$1 a minute.

That system produced low-level step-by-step instructions. But work by Catrambone and others had indicated that learners profited more from outlines that featured something called "subgoal labeling."

"Subgoal labeling is an educational theory that says that people think in terms of hierarchical solution structures," Kim explains. "Say there are 20 different steps to make a cake, such as adding sugar, salt, baking soda, egg, butter, and things like that. This could be just a random series

of steps, if you're a novice. But what if the instruction instead said, 'First, deal with all the dry ingredients,' and then it talked about the specific steps. Then it moved onto the wet ingredients and talked about eggs and butter and milk. That way, your mental model of the solution is much better organized."

Division of labor

The system reported in the new paper, dubbed "Crowdy," produces subgoal labels—and does so essentially for free. Each of a video's first viewers will find it randomly paused at some point, whereupon the viewer will be asked to characterize the previous minute of instruction. After enough candidate descriptions have been amassed, each subsequent viewer will, at one of the same points, be offered three alternative characterizations of the preceding minute. Once a consensus emerges, Crowdy identifies successive minutes of video with similar characterizations and merges their labels. Finally, another group of viewers is asked whether the resulting labels are accurate and, if not, to provide alternatives.

The researchers tested Crowdy with a group of 15 videos about three common Web programming languages, which were culled from YouTube. The videos were posted on the Crowdy website for a month, during which they attracted about 1,000 viewers. Roughly one-fifth of those viewers participated in the experiment, producing an average of eight subgoal labels per [video](#).

In ongoing work, the researchers are expanding the range of topics covered by the videos on the Crowdy website. They're also investigating whether occasionally pausing the videos and asking viewers to reflect on recently presented content actually improves retention. There's some evidence in the educational literature that it should, and if it does, it could provide a strong incentive for viewers to contribute to the

annotation process.

"We did a bunch of experiments showing that subgoal-labeled videos really dramatically improve learning and retention, and even transfer to new tasks for people studying computer science," says Mark Guzdial, a professor of interactive computing at Georgia Tech who has worked with Catrambone. "Immediately afterward, we asked people to attempt another problem, and we found that the people who got the subgoal labels attempted more steps and got them right more often, and they also took less time. And then a week later, we had them come back. When we asked them to try a new problem that they'd never seen before, 50 percent of the subgoal people did it correctly, and less than 10 percent of the people who didn't get subgoals did that correctly."

"Rob and Juho came up with the idea of doing crowdsourcing to generate the labels on videos," Guzdial adds, "which I think is supercool."

More information: "Learnersourcing Subgoal Labels for How-to Videos": juhokim.com/files/CSCW2015-Learnersourcing.pdf

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: System recruits learners to annotate videos (2015, February 11) retrieved 26 April 2024 from <https://phys.org/news/2015-02-learners-annotate-videos.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is

provided for information purposes only.