

Human Computation for Image and Video Analysis (GroupSight)

Kurt Luther

■ *The second Workshop on Human Computation for Image and Video Analysis (GroupSight) took place in Quebec City, Canada, on October 24, 2017, as part of the fifth AAAI Conference on Human Computation and Crowdsourcing (HCOMP). The aim of this workshop was to promote greater interaction between researchers and practitioners engaged in exploring how to combine human and computer efforts to convert visual data into discoveries and innovations that benefit society.*

What would be possible if we could accelerate the analysis of images and videos, especially at scale? This question is inspiring widespread interest across research communities as diverse as computer vision, human-computer interaction, computer graphics, and multimedia.

The goal of the Workshop on Human Computation for Image and Video Analysis was to promote greater interaction between this diverse set of researchers and practitioners, who explored ways to combine human and computer efforts to convert visual data into discoveries and innovations that benefit society at large. This was the second meeting of the GroupSight workshop to be held at the AAAI Conference on Human Computation and Crowdsourcing (HCOMP). It was also the first time the workshop and conference were colocated with the ACM Conference on User Interface Software and Technology. The workshop was held in Quebec City, Quebec, Canada, on October 24, 2017.

The workshop featured two keynote speakers in human-computer interaction (HCI) doing research on crowdsourced image analysis. Meredith Ringel Morris (Microsoft Research)



The Workshop Was Held in Quebec City, Quebec, Canada.

presented work on combining human and machine intelligence to describe images to people with visual impairments. Walter Lasecki (University of Michigan) discussed projects using real-time crowdsourcing to rapidly and scalably generate training data for computer vision systems.

One major theme of the workshop papers was building systems that leverage the visual capabilities of crowdsourced human workers. Abdullah Alshaibani and colleagues at Purdue University presented InFocus, the workshop's Best Paper award winner. This system enables untrusted workers to redact potentially sensitive content from imagery. Kyung Je Jo and colleagues at KAIST presented Exprgram, the workshop's Best Paper Runner-Up award winner. This paper introduced a crowd workflow that supports language learning while annotating and searching videos. GroundTruth, a system by Rachel Kohler and colleagues at Virginia Tech, combines expert investigators and novice crowds to identify the precise geographic location where images and videos were created.

Another major theme of the papers was creating synergies between crowdsourced human visual analysis and computer vision techniques. Steven Gutstein and colleagues from the US Army Research Laboratory presented a system that integrates a brain-computer interface with computer vision techniques to support rapid triage of images. Divya Ramesh and colleagues from CloudSight presented an approach for real-time captioning of images by

combining crowdsourcing and computer vision.

A third theme of the workshop papers was improving methods for aggregating results from crowdsourced image analysis. Jean Song and colleagues at the University of Michigan presented research showing that tool diversity can improve aggregate crowd performance on image segmentation tasks. Anuparna Banerjee and colleagues at the University of Texas, Austin presented an analysis of ways that crowd workers disagree in visual question-answering tasks.

The workshop also featured a poster session and break-out groups. Participants used a bottom-up approach to identify topical clusters of common research interests and open problems. These clusters included real-time crowdsourcing, worker abilities, applications (both to computer vision and in general), and crowdsourcing ethics.

Danna Gurari (UT Austin), Kurt Luther (Virginia Tech), Genevieve Patterson (Brown University and Microsoft Research New England), and Steve Branson (Caltech) served as cochairs of this workshop. The steering committee comprised James Hays (Georgia Tech), Pietro Perona (Caltech), and Serge Belongie (Cornell Tech). Some of the papers are available on the workshop website, groupsight.github.io.

Kurt Luther is an assistant professor in the Department of Computer Science at Virginia Tech.