

Crowd-Training Machine Learning Systems for Human Rights Abuse Documentation

Jay D. Aronson

Center for Human Rights Science
Carnegie Mellon University
aronson@andrew.cmu.edu

Abstract

Social media and mobile phones with good cameras and Internet access are dramatically changing the nature of human rights documentation, reporting and advocacy. Hundreds of hours of relevant video are added to sites like YouTube, Live Leak, Vimeo, and Facebook every week. In Syria, more than 650,000 videos have been uploaded to social media sites since the conflict started three years ago. Violence and human rights abuse footage from regions experiencing mass protest and conflict (including Turkey, Venezuela, and the Ukraine) are heavily represented in social media as well. This trove of potentially useful information includes documentation of political protests, armed conflict, human casualties, mob violence, police brutality, destroyed infrastructure, pleas for help from civilians, humanitarian relief efforts, political statements, interviews with participants in violence and protest, as well as government information and propaganda.

A significant challenge with this video deluge is determining what is immediately relevant, what may be relevant in the long-term, what is irrelevant to the situation or repetitive, and what is patently false or misleading. These determinations must often be made quickly with limited computing and human resources coupled with the need for a deep understanding of local social and political conditions that may change from day-to-day or even hour- to-hour, local landmarks, the types of weaponry available to local forces, etc. Thus, machine learning and computer vision systems are not adequately equipped to assist human analysts in sifting through the massive amount of visual data available without significant training.

In this talk, I will describe efforts being undertaken in a collaboration between human rights advocates and computer scientists at Carnegie Mellon University to develop tools, methods and algorithms that will make it easier to semi-automate the process of video forensics for human rights purposes. Key to this process, and apropos of this session, is the development of mechanisms to enable “the crowd” (i.e., those individuals around the world who care about human rights and have relevant knowledge) to participate in the training process in real time as new situations emerge so that human rights fact-finders can continue to monitor such situations as public interest dies down or moves on to new issues or places. While such systems will never replace analysis by knowledgeable human analysts, they will hopefully augment their ability to monitor situations in near real time.

A key difference between many crowdsourcing projects and the one proposed here is the use of the crowd to train machine learning systems that can aid the professional community in semi-automating the sorting and analysis of large volumes of video data. In presenting this project, I hope to get feedback from other participants in the workshop on how to achieve this goal, particularly by leveraging knowledge gained from work in other domains. I also would like to develop partnerships with other organizations and institutions engaged in similar work on behalf of the public good.