Helen Chen

Newton, MA

Su20-021

Junior
Computer Science
University of Maryland, College Park

Truth, Trust, and Transparency in Synthetic Media

Helen Chen, John Jack Lewis, Imad Eddine Toubal, Vishal Sandesera, Michael Lomnitz, Zigfried Hampel-Arias, Kannappan Palaniappan, and Prasad Calyam

Authenticity of digital media has become an everpressing necessity for modern society. Since the introduction of Generative Adversarial Networks (GANs), synthetic media has become increasingly difficult to identify. Synthetic videos that contain altered faces and/or voices of a person are known as deepfakes and threaten trust and privacy in digital media. Deepfakes can be weaponized for political advantage, slander, and to undermine the reputation of public figures. Despite imperfections of deepfakes, people struggle to distinguish between authentic and manipulated images and videos. Consequently, it is important to have automated systems that accurately and efficiently classify the validity of digital content. Many recent deepfake detection methods use single frames of video and focus on the spatial information in the image to infer the authenticity of the video. Some promising approaches exploit the temporal inconsistencies of manipulated videos; however, research primarily focuses on spatial features. We propose a hybrid deep learning approach that uses spatial, spectral, and temporal content that is naturally coupled in a consistent way to differentiate real and fake videos. In this work, we build a computationally efficient cloud-ready multimodal system to detect deepfake videos. We evaluate the performance of our proposed system compared to recent approaches, in terms of accuracy and speed, on the Facebook Deepfake Detection Challenge and FaceForensics++ video datasets.