# Live Video Analytics

**Ganesh Ananthanarayanan**
Microsoft Research

**Weisong Shi**
Wayne State University

**CAMERAS ARE EVERYWHERE!** According to a January 2013 report published by the National Public Radio, the Chinese government has installed more than 20 million cameras across the country. Similarly, a recent BBC report stated, there is one camera for every 14 people in London. Other large cities including New York, Paris, and Tokyo are also deploying cameras in large numbers, so it is reasonable to claim that networked cameras are everywhere. These cameras are deployed for a wide variety of commercial and security reasons. Further, consumer devices themselves have cameras with users interested in streaming videos live from these devices.

We are all living in the golden era for computer vision and AI, with many recent advancements combined with copious training data and systems infrastructure, largely improving their range of capabilities. Analyzing live videos from these cameras has great potential to impact science and society. Live video analytics has the potential to impact a wide range of verticals ranging from public safety, traffic efficiency, infrastructure planning, entertainment, and home safety.

Analyzing live video streams is arguably the most challenging of domains for systems-for-AI. Unlike text or numeric processing, video analytics require higher bandwidth, consume considerable CPU and GPU resources for processing, necessitate richer query semantics, and demand tighter security and privacy guarantees. The widespread need for video analytics and abundance of video data presents the unique and timely opportunity to design solutions for live video analytics. All aspects of video analytics call to be designed "green-field," from vision algorithms, to the systems processing stack and networking infrastructure, including hybrid cloud solutions. Such a holistic design will enable the democratization of live video analytics such that any organization with cameras can obtain value from video analytics.

Democratizing video analytics has many key challenges including low-cost video analytics, storage of video data and metadata, network design for video streams including hybrid cloud designs, scheduling for multitenant video processing in large distributed systems, training and inference of vision neural networks, microprocessor architectures for machine learning and video processing, secure processing of video analytics, privacy-preserving techniques for video processing, and interactive querying of video streams

In this special issue, many of these problems are under active investigations. The paper "*Effectively Linking Persons on Cameras and Mobile Devices on Networks*" targets the problem of video analytics on multiple cameras networked together by tying together the mobile devices of users with their appearances on the cameras. This is a considerable advancement

*Digital Object Identifier 10.1109/MIC.2019.2930947*
*Date of current version 7 October 2019.*

over the state-of-the-art approaches that only use visual localization or wireless localization alone.

Drones with cameras are increasingly touted to become a key source for videos for a variety of scenarios including disaster relief and ecological surveillance. The paper "*Edge-Based Live Video Analytics for Drones*" is one of the earliest works on real-time video analytics on small autonomous drones that tackles the difficult challenges at the intersection of wireless bandwidth, processing capacity, energy consumption, and accuracy of the results. We believe it will pave the way for many more works in this very important area.

Driven by concerns on performance and privacy, video analytics is predicted to largely be driven by "smart" cameras with compute on them. Thus, virtualizing the resources on the camera is a key requirement. "*Towards Dynamically Reconfigurable IoT Camera Virtualization for Video Analytics Edge Cloud Services*" takes an important stab at this problem across a cluster of cameras and multiple tenants (or video analytics applications).

Finally, key to effective video analytics is not just inference of the neural network models but also training of the models with effective training data. We have work on "*Edge-Based Discovery of Training Data for Machine Learning*" that shows how edge-based discovery of the training data allows for early discard of data at the edge itself, thus when the data reaches the human labeling experts in the cloud, it is already of higher quality. This is valuable to ensuring effective labeling, which in turn will lead to accurate models over time.

We hope that readers will find these articles interesting and informative. We also would like to thank all authors for their submissions and all reviewers for their timely and high-quality reviews.

**Ganesh Ananthanarayanan** is a Researcher at Microsoft Research. His research interests are broadly in systems and networking, with recent focus on live video analytics, cloud computing and large-scale data analytics systems, and Internet performance. He has authored or coauthored more than 30 papers in systems and networking conferences such as USENIX OSDI, ACM SIGCOMM, and USENIX NSDI, with the highest rated paper at ACM Symposium on Edge Computing (SEC) 2018 and Best Demo Award (runner-up) at ACM MobiSys 2019. He received the Ph.D. degree in December 2013 from the University of California, Berkeley, CA, USA, where he was also a recipient of the UC Berkeley Regents Fellowship. His work on "Video Analytics for Vision Zero" on analyzing traffic camera feeds won the Institute of Transportation Engineers 2017 Achievement Award as well as the "Safer Cities, Safer People" U.S. Department of Transportation Award. He has collaborated with and shipped technology to Microsoft's cloud and online products like the Azure Cloud, Cosmos (Microsoft's big data system) and Skype. He is a member of the ACM Future of Computing Academy. Contact him at ga@microsoft.com.

**Weisong Shi** is a Professor of Computer Science with Wayne State University, Detroit, MI, USA. His research interests mainly focus on big data systems, edge computing, energy-efficient computer systems and mobile health. He received the B.S. degree from Xidian University, Xi'an, China, in 1995 and the Ph.D. degree from the Chinese Academy of Sciences, Beijing, China, in 2000, both in computer engineering. He served as the founding steering committee chair of IEEE/ACM Symposium on Edge Computing (SEC). He is a recipient of the National Outstanding Ph.D. Dissertation Award of China and the NSF CAREER award. He is an IEEE Fellow and an ACM Distinguished Scientist. Contact him at weisong@wayne.edu.