

AI-Driven Video Compression and Enhancement: From Core Technologies to Real-World Applications



Yuxin "Zoe" Liu
CEO and Co-Founder
Visionular

Thursday, October 16
11:00 A.M. • MSEE 112

Abstract

We are experiencing explosive growth of video applications across on-demand, live streaming, and real-time communications. The demand for intelligent and efficient video compression, understanding, processing, and enhancement has reached unprecedented levels. In this talk, I will present recent advances in AI-driven video technologies we have developed, focusing on how these innovations are grounded in academic research and actively deployed in industrial applications at scale.

We begin by exploring the entire video pipeline, from video capture on acquisition and processing devices, to handling on the server and in the cloud, and finally to presentation on the player side. We will then highlight the following six subtopics:

Intelligent CPU encoding: We will detail how to achieve per-CPU encoding efficiency optimization to deliver optimal compression efficiency while reducing CPU usage.

1. Intelligent GPU encoding: We will dive into an NVIDIA GPU NVENC-based encoding approach that utilizes three core models: no-reference VQA, adaptive enhancement, and ROI-guided rate control, all implemented with GPU-based machine learning. This approach enables high-density, low-latency encoding while significantly improving perceived video quality and reducing bitrate.
2. Film-grain handling: We will describe a joint grain synthesis and grain preservation approach to effectively encode film grain visual content. We also propose a novel full-reference video quality metric specifically designed to evaluate the compressed visual quality of grain content.
3. Machine-oriented compression: We will highlight why machine-oriented compression is needed and how it differs from human-oriented compression.
4. AI for live video streaming: We will present a series of technologies beyond video compression, including intelligent video editing based on ML-driven video understanding and processing, targeted at live-streaming use cases.
5. AOM AV1 & AV2: AOM, the open-source, royalty-free media industry consortium that introduced AV1 in 2018, has just announced that AV2 will be finalized by the end of 2025. It highlights four key AV2 features: support for XR/VR, split-screen, tools for screen content coding (SCC), and low-complexity coding. It is worth digging into why these features matter to the industry.

This talk is designed for researchers and students interested in video and image processing, machine learning, and multimedia systems, and it aims to highlight the synergy between academic innovation and industry-scale deployment in AI-driven video compression, enhancement, and streaming.

Bio

Zoe (Yuxin) Liu has 25+ years of experience in video compression and streaming technologies. She holds BE/ME degrees from the Department of Automation at Tsinghua University, and earned her PhD from School of ECE at Purdue University, West Lafayette, IN. Throughout her career, she has worked at Bell Labs, Nokia Research Center, HP Labs, Apple, and Google, contributing to the development of core video technologies such as Apple FaceTime and Google Glass Video Call. As a key contributor to the AOM/AV1 video coding standard, she introduced AV1 to the global developers at Google I/O 2018.

In 2018, Zoe co-founded Visionular and currently serves as its CEO. She leads the development of AI-driven video compression, processing, understanding, and streaming software and cloud services. Visionular has served nearly 200 enterprise customers and employs 100+ FTEs globally. Headquartered in Silicon Valley, the company also has offices in Denver, London, Singapore, Bangalore, Hangzhou, and Beijing. Visionular is profitable and continues to see steady growth in both revenue and profitability.

Zoe has published nearly 50 peer-reviewed academic papers and delivered over 30 talks at major industry conferences. Zoe has been recently selected as one of six recipients of the 2025 Purdue Outstanding Electrical and Computer Engineer (OECE) award, the achievements of alumni.

Zoe can be reached via email at zoeliu@visionular.com. Visionular's headquarters is located at 280 2nd Street, Suite 218, Los Altos, California, 94022, USA.

Host: Maggie Zhu, zhu0@purdue.edu

Zoom link: <https://purdue-edu.zoom.us/j/94633494251> ~ Meeting ID: 946 3349 4251



Elmore Family School of Electrical
and Computer Engineering