

NIER: Practical Neural-enhanced Low-bitrate Video Conferencing

Anlan Zhang¹ Yuming Hu² Chendong Wang³ Yu Liu¹ Zejun Zhang¹ Haoyu Gong²

*Ahmad Hassan¹ Shichang Xu⁴ Zhenhua Li⁵ Bo Han⁶ **Feng Qian¹***

¹University of Southern California ²University of Minnesota, Twin Cities

³University of Wisconsin-Madison ⁴Google ⁵Tsinghua University

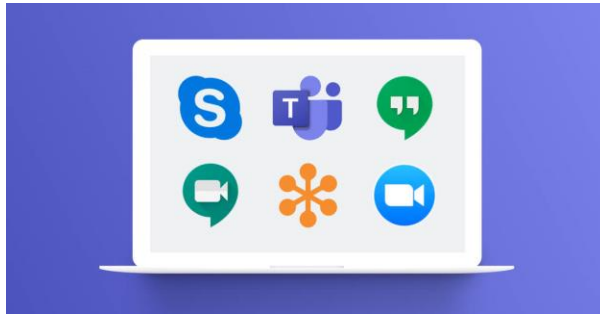
⁶George Mason University

September, 2025

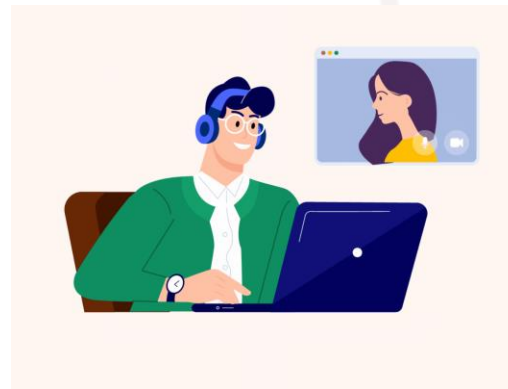


Why Low-bitrate Video Conferencing?

- Today's video conferencing consumes a substantial amount of data
 - Zoom [1] requires 1.2 Mbps UL/DL for 1-on-1 video all at 720p → ~1.05 GB for a 1-hour session
- Benefit multiple stakeholders



Streaming platforms

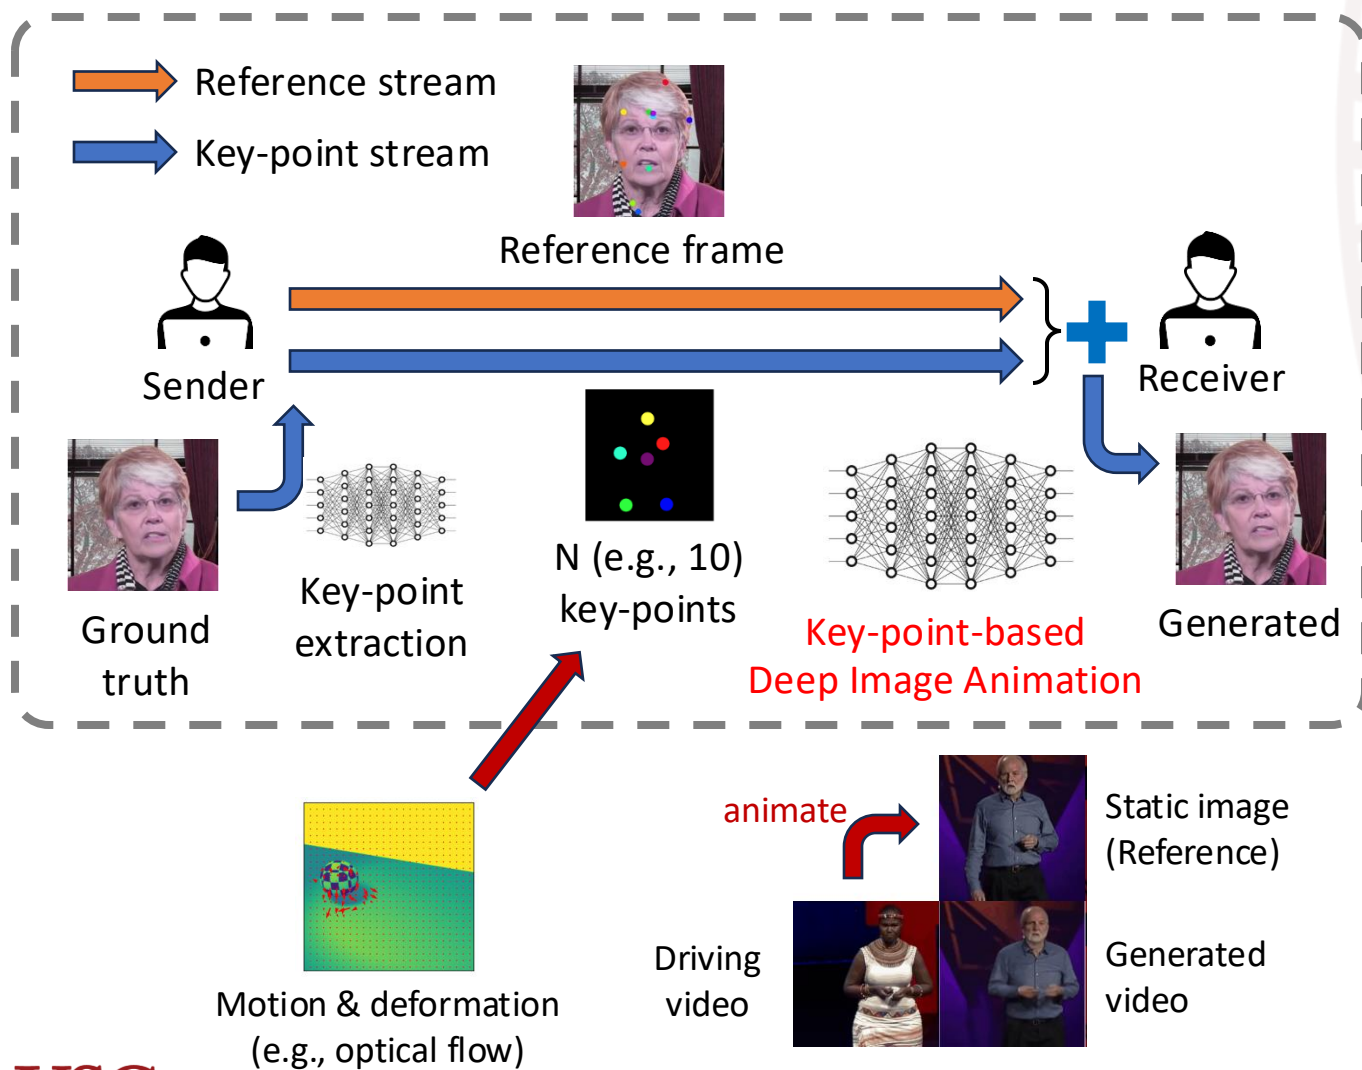


Mobile customers



Cellular providers

NIER: Neural-enhanced Low-bitrate Video Conferencing

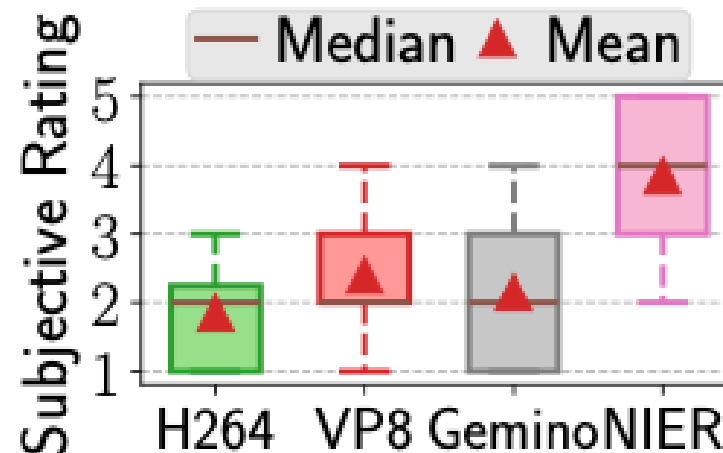


Key design of NIER

- **Consistent good quality**
 - Dynamic reference frame update & reuse
- **Bandwidth adaptation**
 - Layered key-point coding
 - Reference streaming reshaping
 - Priority-based packet scheduling
- **Resilience to packet loss**
 - Inferring missing key-point
 - Loss-resilient reference frame split
- **DIA acceleration**

User Study w/ 20 Participants (10 1-on-1 Calls)

- Candidates: H264, VP8, Gemino [1], NIER
- Bandwidth: {75, 45} Kbps + 128 Kbps (for audio)
- Packet loss rate: {0.05, 0.1, 1, 5, 10}%
- Device: MacBook Air M1 2020



NIER outperforms other low-bitrate schemes in subjective ratings by up to 34%.