AMVP: Adaptive CNN-based Multitask Video Processing on Mobile Stream Processing Platforms

Mengyuan Chao, Radu Stoleru, Liuyi Jin, Shuochao Yao*, Maxwell Maurice†, Roger Blalock†
Department of Computer Science and Engineering, Texas A&M University
*Department of Computer Science, George Mason University
†National Institute of Standards and Technology (NIST)

The Fifth ACM/IEEE Symposium on Edge Computing (SEC'20)
Virtual, November 11-13, 2020

Outline

- Background & Motivation
- Related Work
- System Overview
- Adaptive Multitask Video Processing Framework
- Experimental Results
- Conclusion & Future Work

Background – Effective Disaster Response

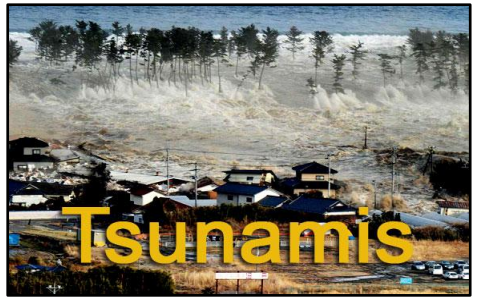

(200,000+, $15B, Indonesia, 2004)

(300,000+, $8.5B, Haiti, 2010)

(2,900+, $90B, Puerto Rico, 2017)

Effective Disaster Response (EDR)

Challenges

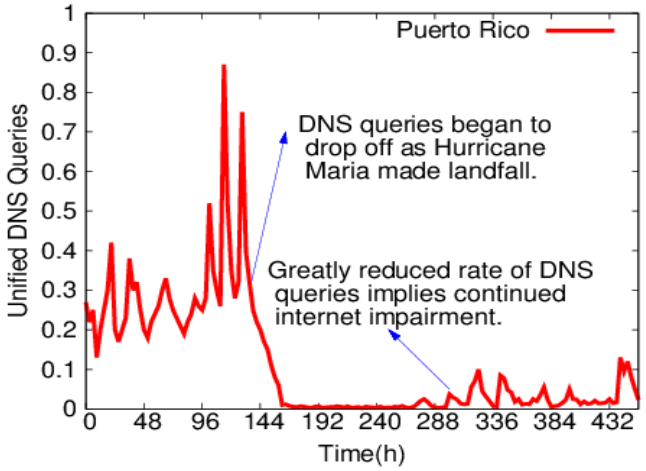

No ISP Network

No 3G, 4G, LTE

No Social Networks

No Cloud Service

Edge Computing Communication (ECC)

Pictures source: http://www.stayinbusiness.com/

Background – Techniques for ECC Enabling EDR

- **Current: System On Wheels (SOW)**

Incident Response Vehicle

AWS Public Safety & Disaster Response

Mobile Communication Vehicle

- **Limitations of SOW**
  - Need long time to arrive
  - Cannot go harshest area
  - Expensive to equip with

Picture source: https://www.fema.gov, Federal Emergency Management Agency

- **Next Generation: Edge Bubble (EB)**

The Next Generation First Responder

First responders work as a team

- **Advantages of EB**
  - Shorter time to arrive
  - Can reach harshest area
  - Cheaper to equip with
  - Closer to data source

Pictures source:https://www.fedscoop.com/dhs-wearables-first-responders/

Motivation – Multitask Video Processing at EB

- **Scenario**: First responders need automatically record multiple information about survivors

Old or young?

Male or female?

Looks OK or not?

Multiple similar CNN-based vision analysis pipelines

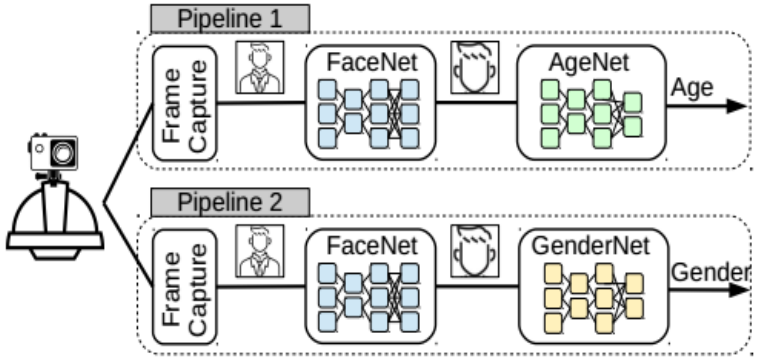

- **Challenges**
  - Limited computing resources on mobile devices
  - Computation intensive video stream analysis
  - Dynamic computing resources and networks
  - Diverse user performance requirements

- **Idea**: Combine multiple pipelines into a single one and adaptively share and offload CNN layers

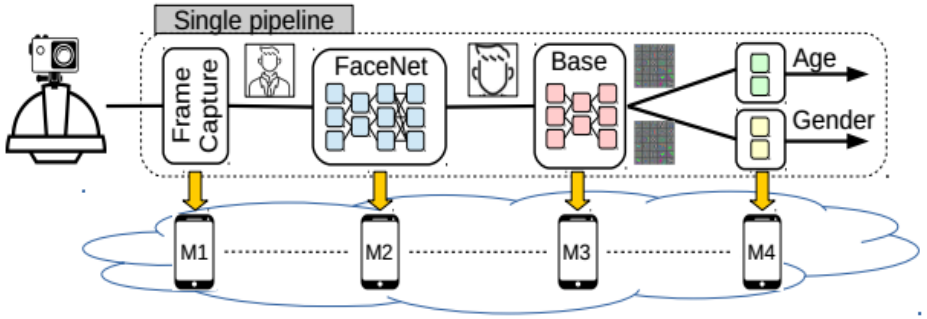

Picture source: https://depositphotos.com/178854658/stock-photo-rescue-team-treating-injuring-field.html

Related Work – CNN-based Mobile Vision Processing

- **CNN Offloading**
  - **Cloud**: DDNN@ICDCS'17, Neurosurgeon@SIGARCH'17, JALAD@ICPADS'18, µLayer@EuroSys'19
  - **Cloudlet**: MODI@HotEdge'18, IONN@SoCC'18, DADS@Infocom'19, Couper@SEC'19
  - **IoT/Mobile**: Modnn@DATE'17, Mednn@ICCAD'17, MusicalChair@arXiv'18, DeepThings@TCAD'18
- **CNN Compression**
  - **One-fit-all**: XNOR-Net@ECCV'16, Thinet@ICCV'17, FactorizedCNN@ICCV'17, ShuffleNet@CVPR'18
  - **Adaptive**: DeepX@IPSN'16, AdaptDNN@LCTES'18, OnDemandDNN@MobiSys'18, ContextDNN@ICDCS'20
  - **Feature**: DeepFCPX@ICIP'18, ContextFCPX@CVPR'18, LosslessFCPX@MMSP'18, LossyFCPX@ICM'19
- **CNN Sharing**
  - **Inside a Model:** NestDNN@MobiCom'18, FoggyCache@MobiCom'18, DeepCache@MobiCom'18
  - **Among Models:** MCDNN@MobiSys'16, Mainstream@ATC'18

We study adaptive CNN-based multitask video processing on mobile devices with dynamic computing resources, network, and user goals.

System Overview – Hardware Architecture

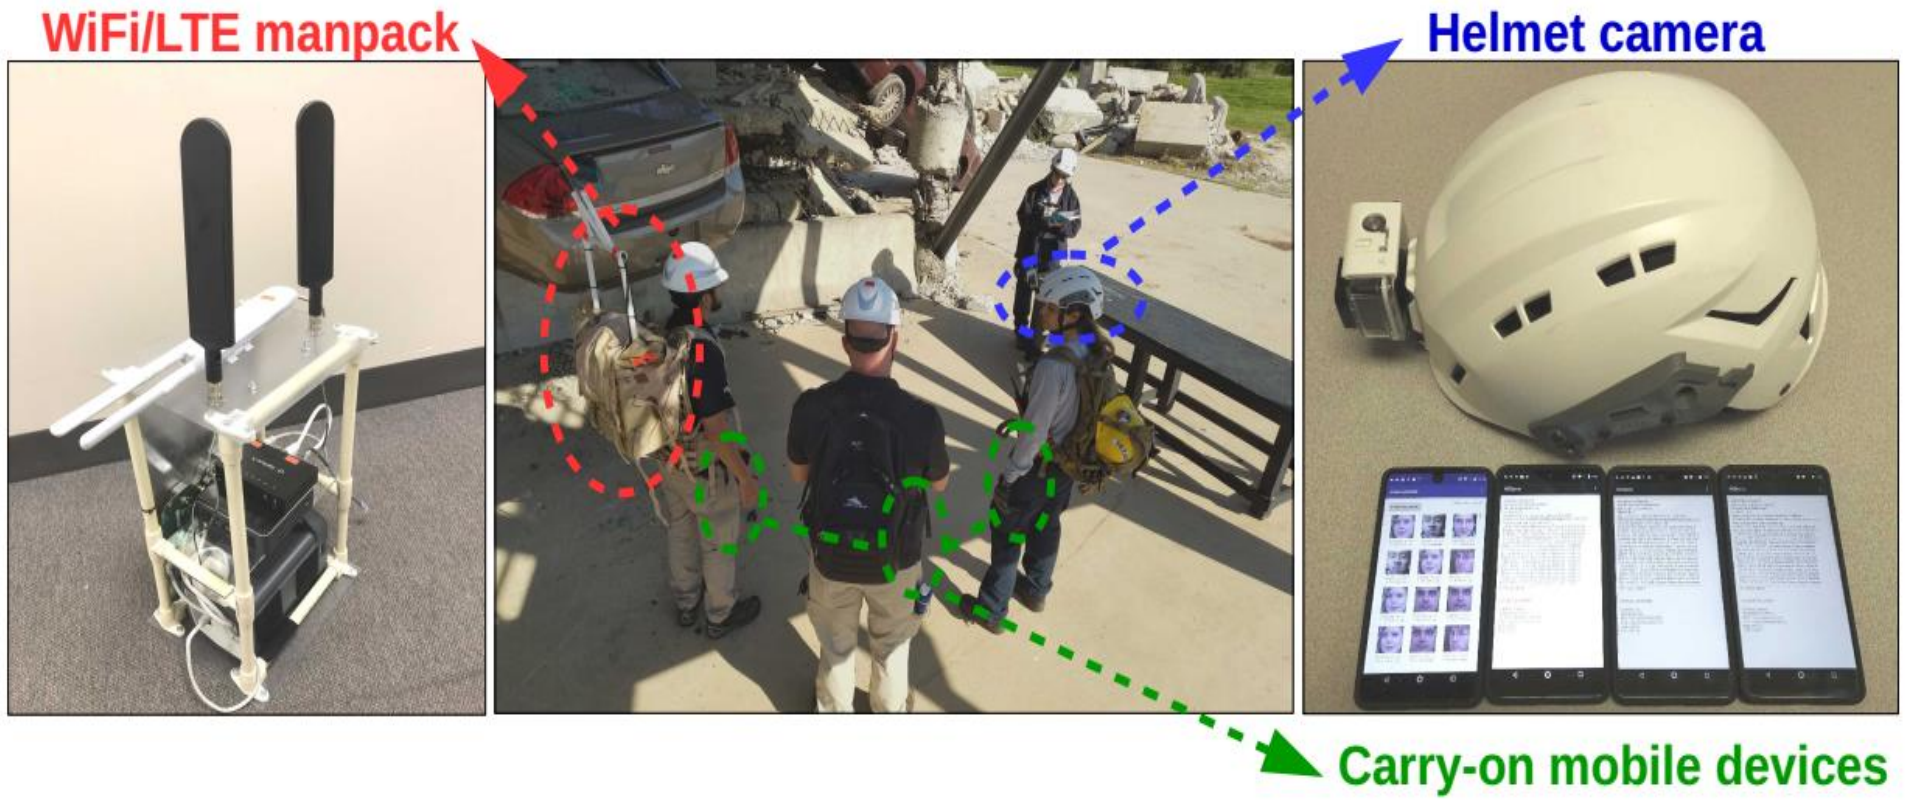

Hardware Implementation of Edge Bubble

System Overview – Software Architecture

- **AMVP**: Adaptive CNN-based Multitask Video Processing Framework

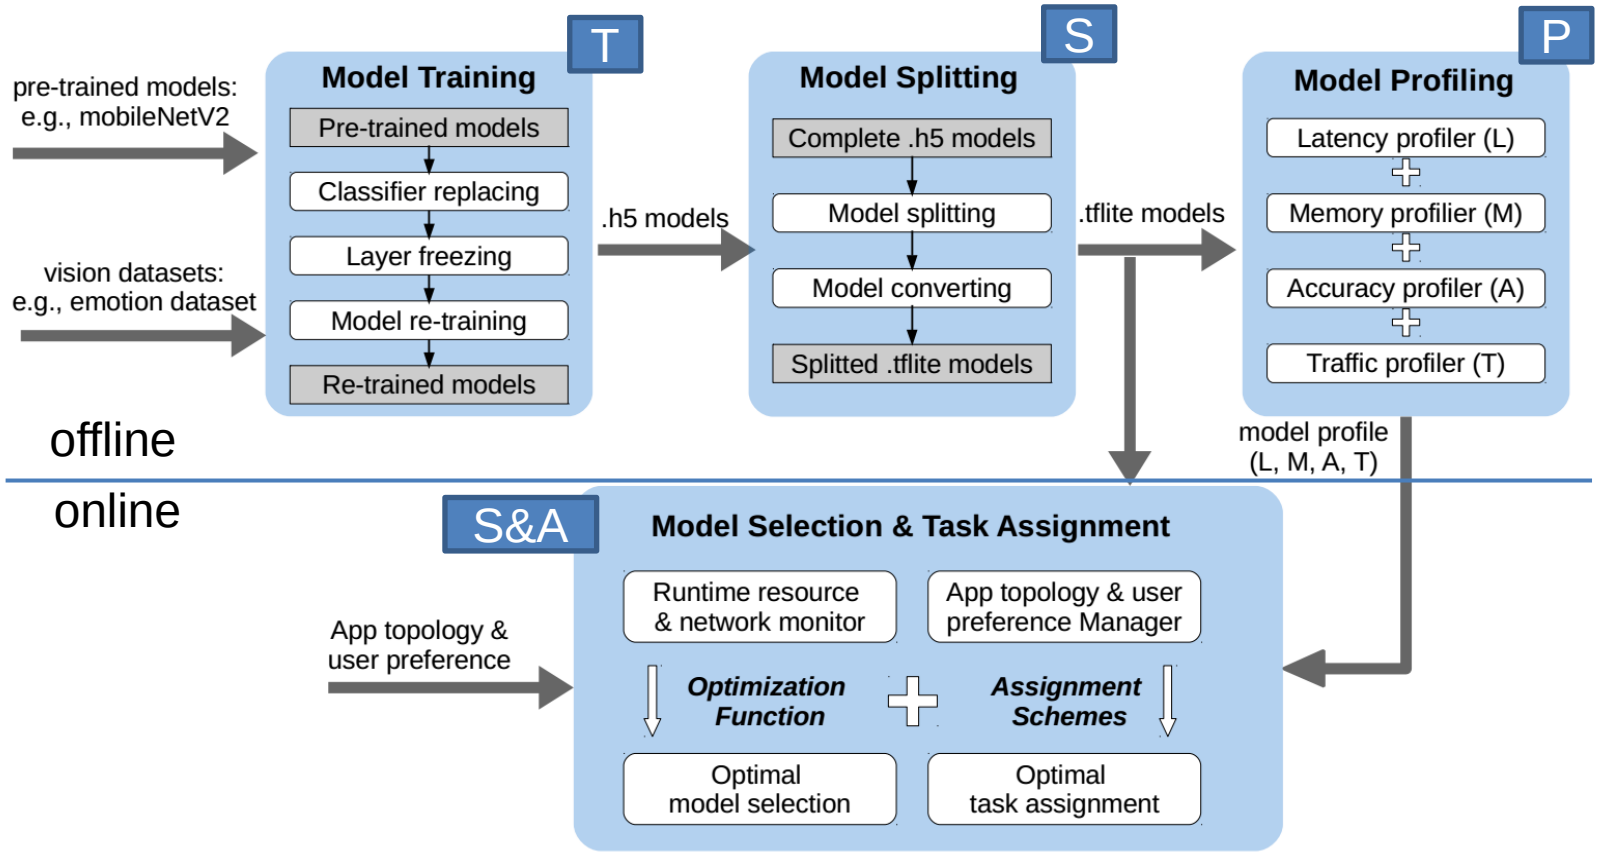

Software Architecture of AMVP

Adaptive Multitask Video Processing Framework

- **Model Training via Transfer learning**

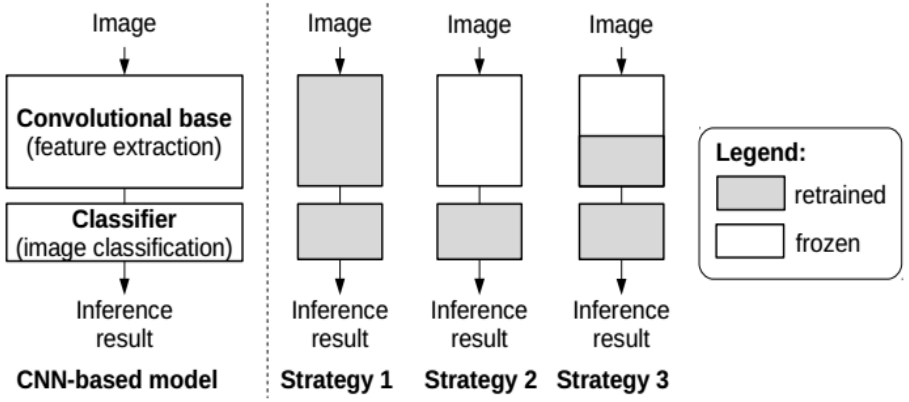

- **Strategy 1**: retrains all weights, requires a large dataset and a lot of computation
- **Strategy 2**: freezes whole convolutional base, only trains classifier; suitable for training tasks similar to original task
- **Strategy 3**: freezes some layers in base, trains the rest; a small dataset, freezes more; a large dataset, retrains more

- **Transfer learning results**

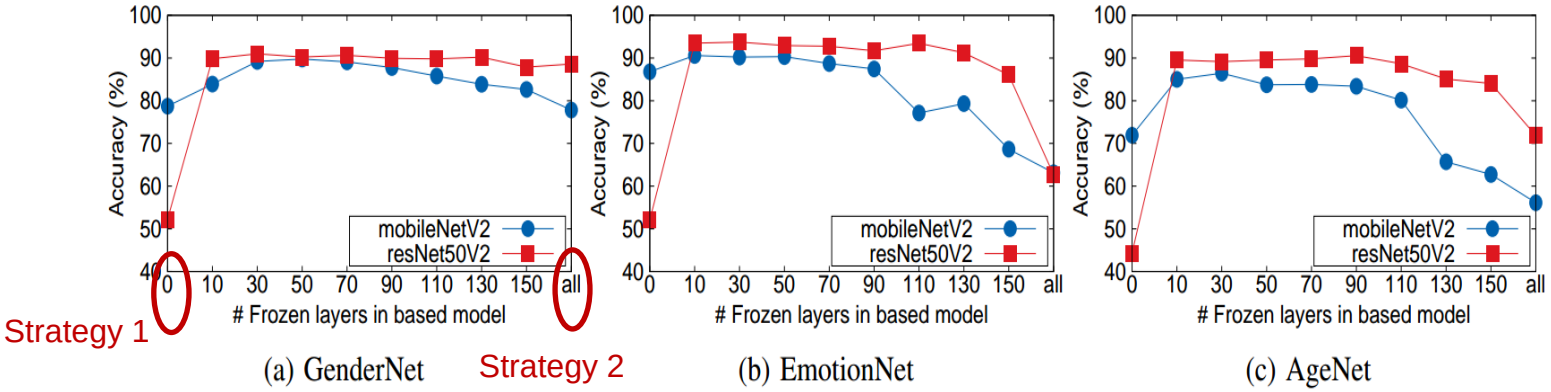

- **Observation 1**: resNet50V2-based are more accurate than mobileNetV2-based
- **Observation 2**: Strategy 3 achieves optimal performance for our tasks and datasets
- **Observation 3**: Simpler task is less affected by # frozen layers than complicated tasks

Adaptive Multitask Video Processing Framework

T P S S&A

- **Model Profiling**

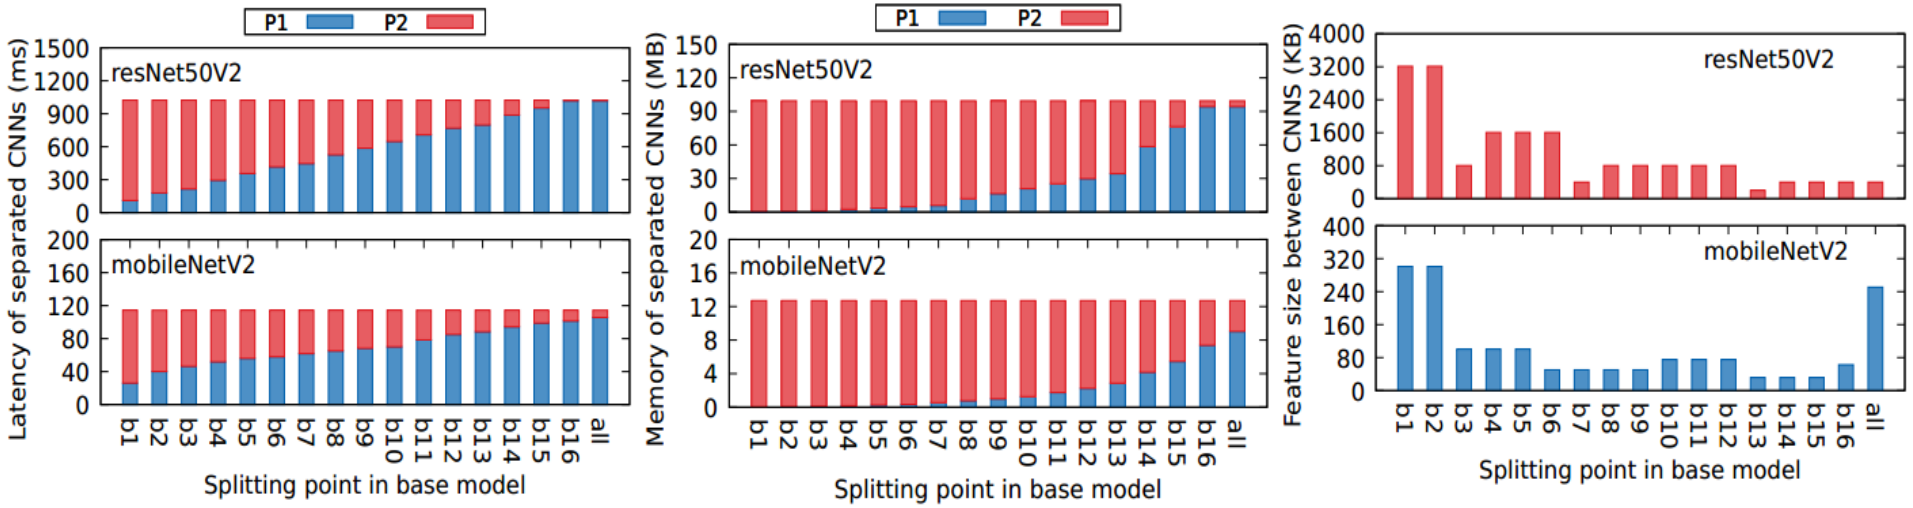

(a) Latency (b) Memory size (c) Feature size

- **Observation 1:** latency of P1 increases with block number
- **Observation 2:** memory is concentrated in the last few layers
- **Observation 3:** feature size gradually gets reduced along with inference process

Adaptive Multitask Video Processing Framework

T P S S&A

- **Model Splitting**

Split at earlier layer	VS.	Split at latter layer
➢ Higher inference accuracy ➢ Better computation balance ➢ Share less among CNNs ➢ Worse memory balance ➢ Larger feature traffic size		➢ Share more among CNNs ➢ Better memory balance ➢ Smaller feature traffic size ➢ Lower inference accuracy ➢ Worse computation balance

- **Quantization-based Feature Compression**

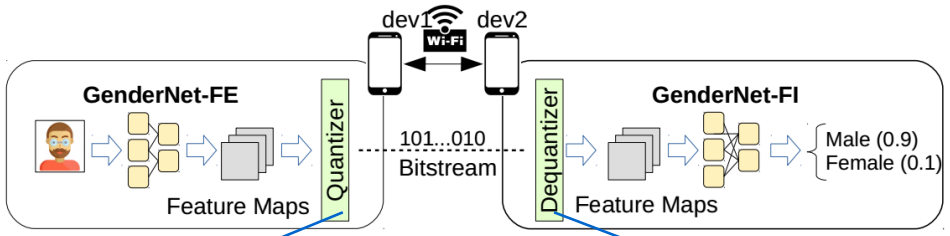

32 bits ==> n bits

$$\overline{\mathbf{F}} = \lfloor \frac{\mathbf{F} - \min(\mathbf{F})}{\max(\mathbf{F}) - \min(\mathbf{F})} \cdot (2^n - 1) \rceil$$

$$\widehat{\mathbf{F}} = \frac{\max(\mathbf{F}) - \min(\mathbf{F})}{2^n - 1} \cdot \overline{\mathbf{F}} + \min(\mathbf{F})$$

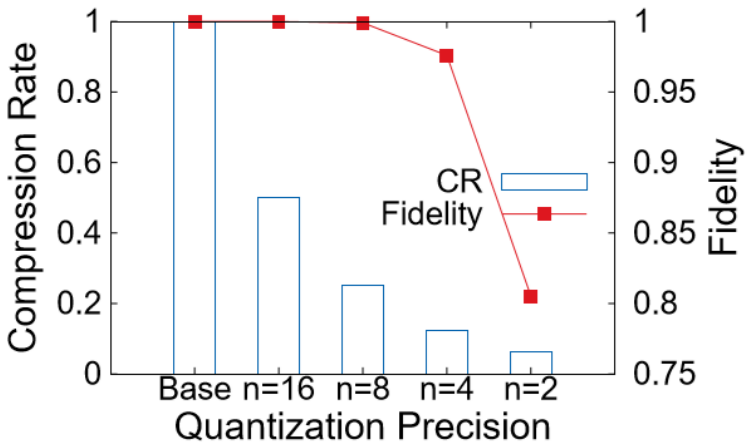

Adaptive Multitask Video Processing Framework

T P S S&A

- **Model Selection and Task Assignment (MSTA)**
  - Cost function for each task

accuracy goal

accuracy

$$C(\boldsymbol{m_s}, \boldsymbol{u_s^k}, s) = \alpha_s \cdot \frac{A_s - A(\boldsymbol{m_s})}{A_s} + \beta_s \cdot \frac{\max(0, L(\boldsymbol{m_s}, \boldsymbol{u_s^k}) - L_s)}{L(\boldsymbol{m_s}, \boldsymbol{u_s^k})} + \gamma_s \cdot \frac{\max(0, T_s - T(\boldsymbol{m_s}, \boldsymbol{u_s^k}))}{T_s}$$

$$\boldsymbol{m_s} = (m_s^1, m_s^2)$$

$$\boldsymbol{u_s^k} = (u_{m_s^1}^{k_1}, u_{m_s^2}^{k_2})$$

P1's delay

P2's delay

communication delay

$$L(\boldsymbol{m_s}, \boldsymbol{u_s^k}) = \frac{l_{m_s^1}^{k_1}}{u_{m_s^1}^{k_1}} + \frac{l_{m_s^2}^{k_2}}{u_{m_s^2}^{k_2}} + D(f_{m_s^1 m_s^2}, b_{k_1 k_2})$$

delay goal

$$T(\boldsymbol{m_s}, \boldsymbol{u_s^k}) = \min\{u_{m_s^1}^{k_1} t_{m_s^1}^{k_1}, u_{m_s^2}^{k_2} t_{m_s^2}^{k_2}, \frac{1}{D(f_{m_s^1 m_s^2}, b_{k_1 k_2})}\}$$

throughput goal

P1's throughput

P2's throughput

communication throughput

  - Problem formulation

$$\begin{aligned} \underset{\boldsymbol{m_s}, \boldsymbol{u_s^k}}{\text{minimize}} \quad & C \\ \text{subject to:} \quad & \forall s : C(\boldsymbol{m_s}, \boldsymbol{u_s^k}, s) \leq C, \\ & \forall k : \sum_{\{m_s^i\}} u_{m_s^i}^k \leq U_k, \\ & \forall k : \sum_{\{m_s^i\}} r_{m_s^i}^k \leq R_k \end{aligned}$$

computing resource constraints

memory constraints

Adaptive Multitask Video Processing Framework

- Greedy algorithm for MSTA

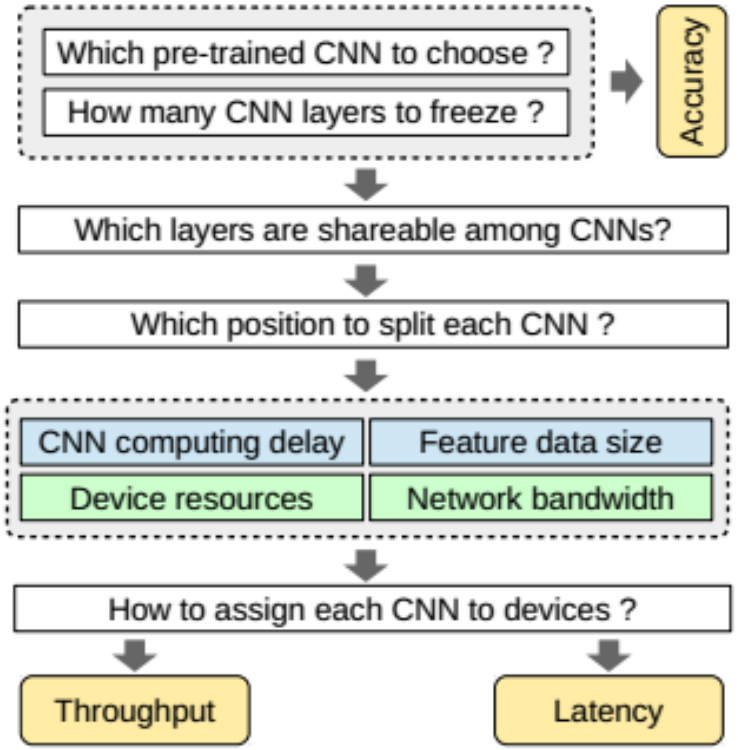

(a) Basic workflow of MSTA

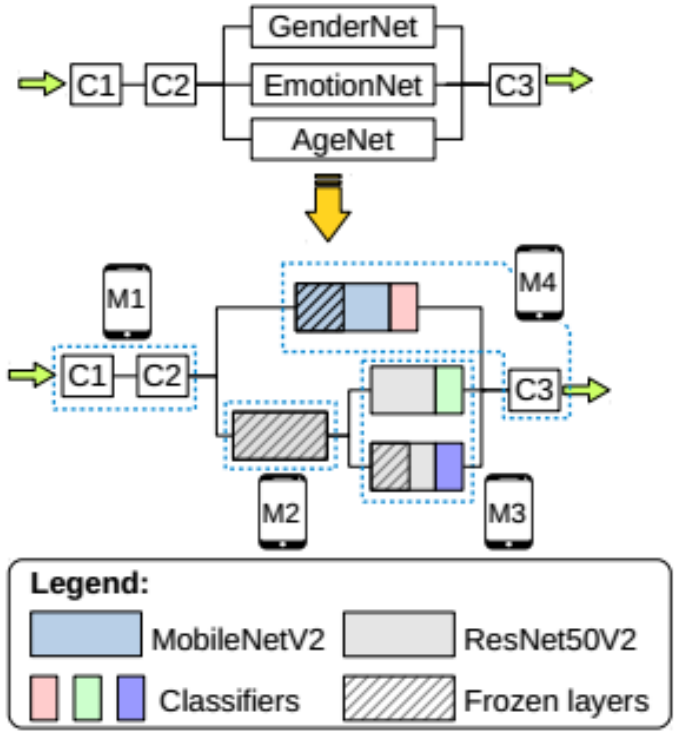

(b) An example of MSTA

Experimental Results

- **Adapt to different accuracy requirements**

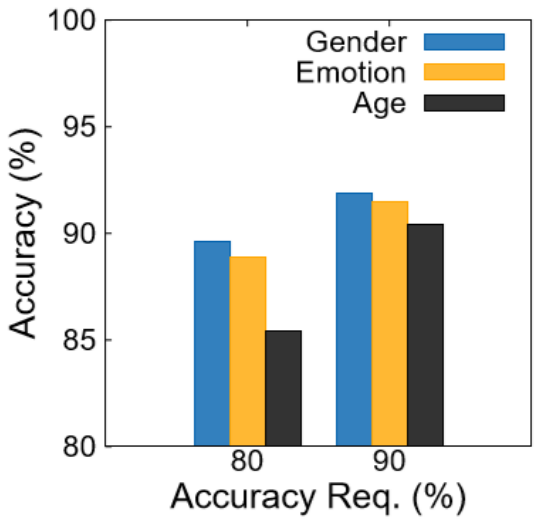

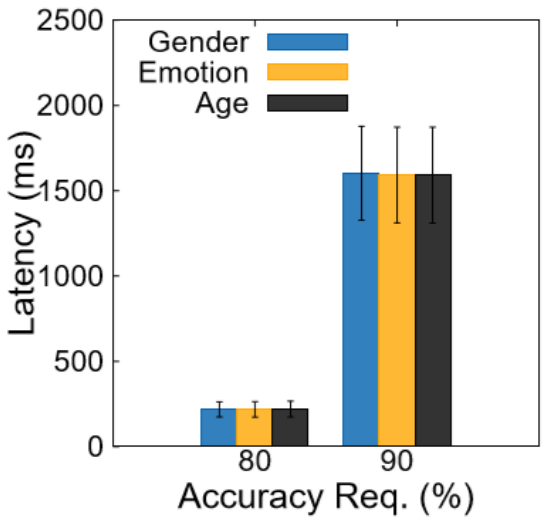

- **Adapt to different latency requirements**

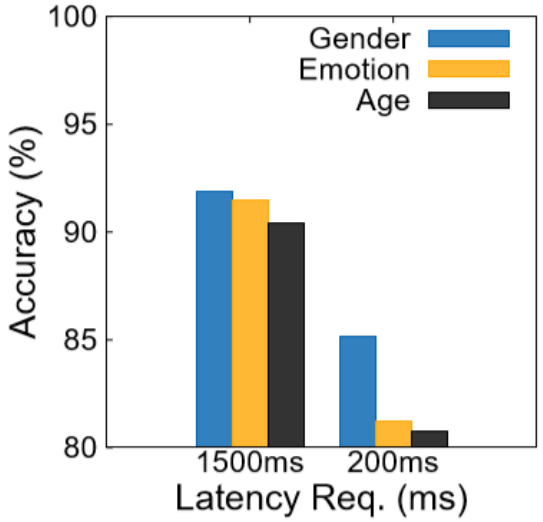

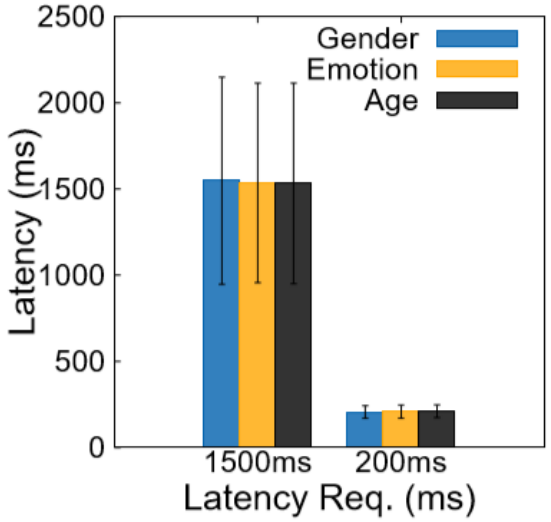

Experimental Results

- **Adapt to different computing resources**
- **Adapt to different network conditions**

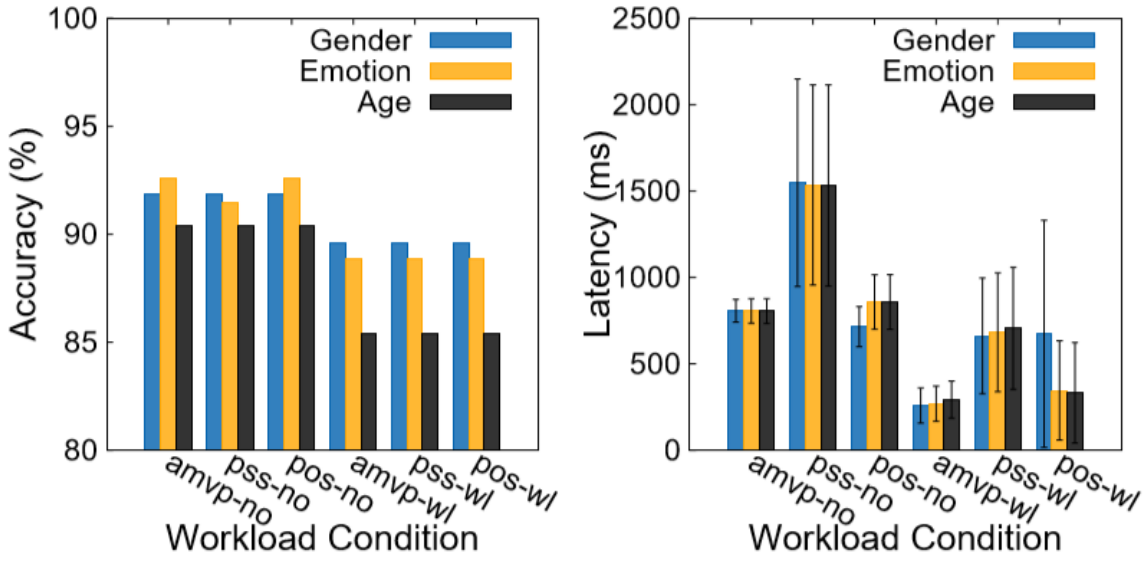

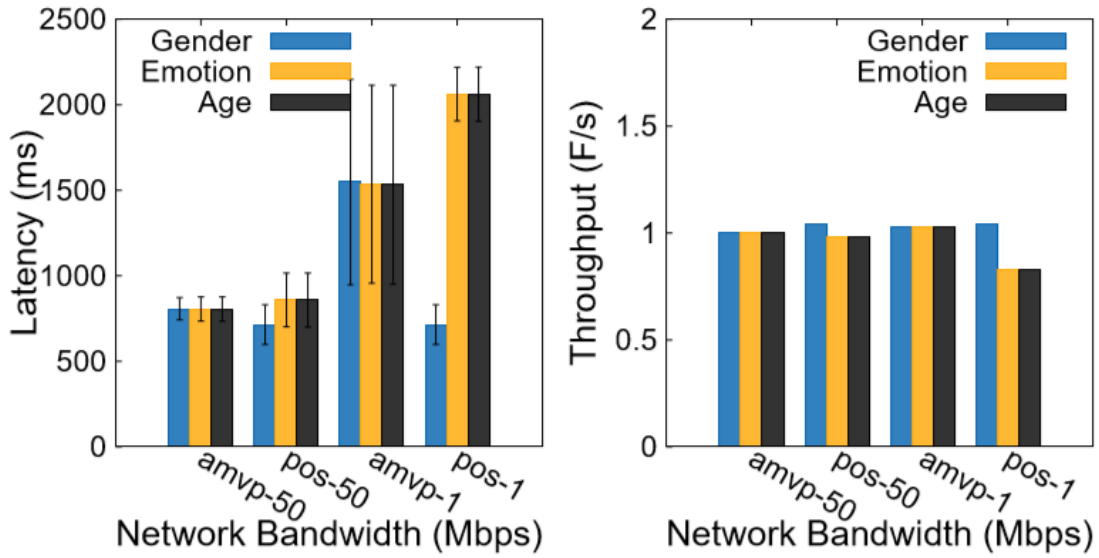

Conclusion and Future Work

- **Conclusion**: We propose AMVP, an Adaptive Multitask Video Processing Framework which supports dynamic CNN layer sharing among multiple CNN-based vision analysis tasks and adaptive CNN layer offloading from one mobile device to other devices at an Edge Bubble.

- **Future Work**

  - Apply accuracy metrics as in Chameleon [Sigcom'18], which includes precision, recall and F1 score.
  - Generalize AMVP to deploy on a heterogeneous stream processing platform including edge server
  - Support simultaneous processing of multiple videos
  - Support other AI applications such as voice recognition, speech recognition, NLP, etc.

Thank you!
Q&A

Contact: stoleru@cse.tamu.edu