

Practical image compression with deep neural networks

Luis Herranz

Computer Vision Center
Universitat Autònoma de Barcelona

March 2022

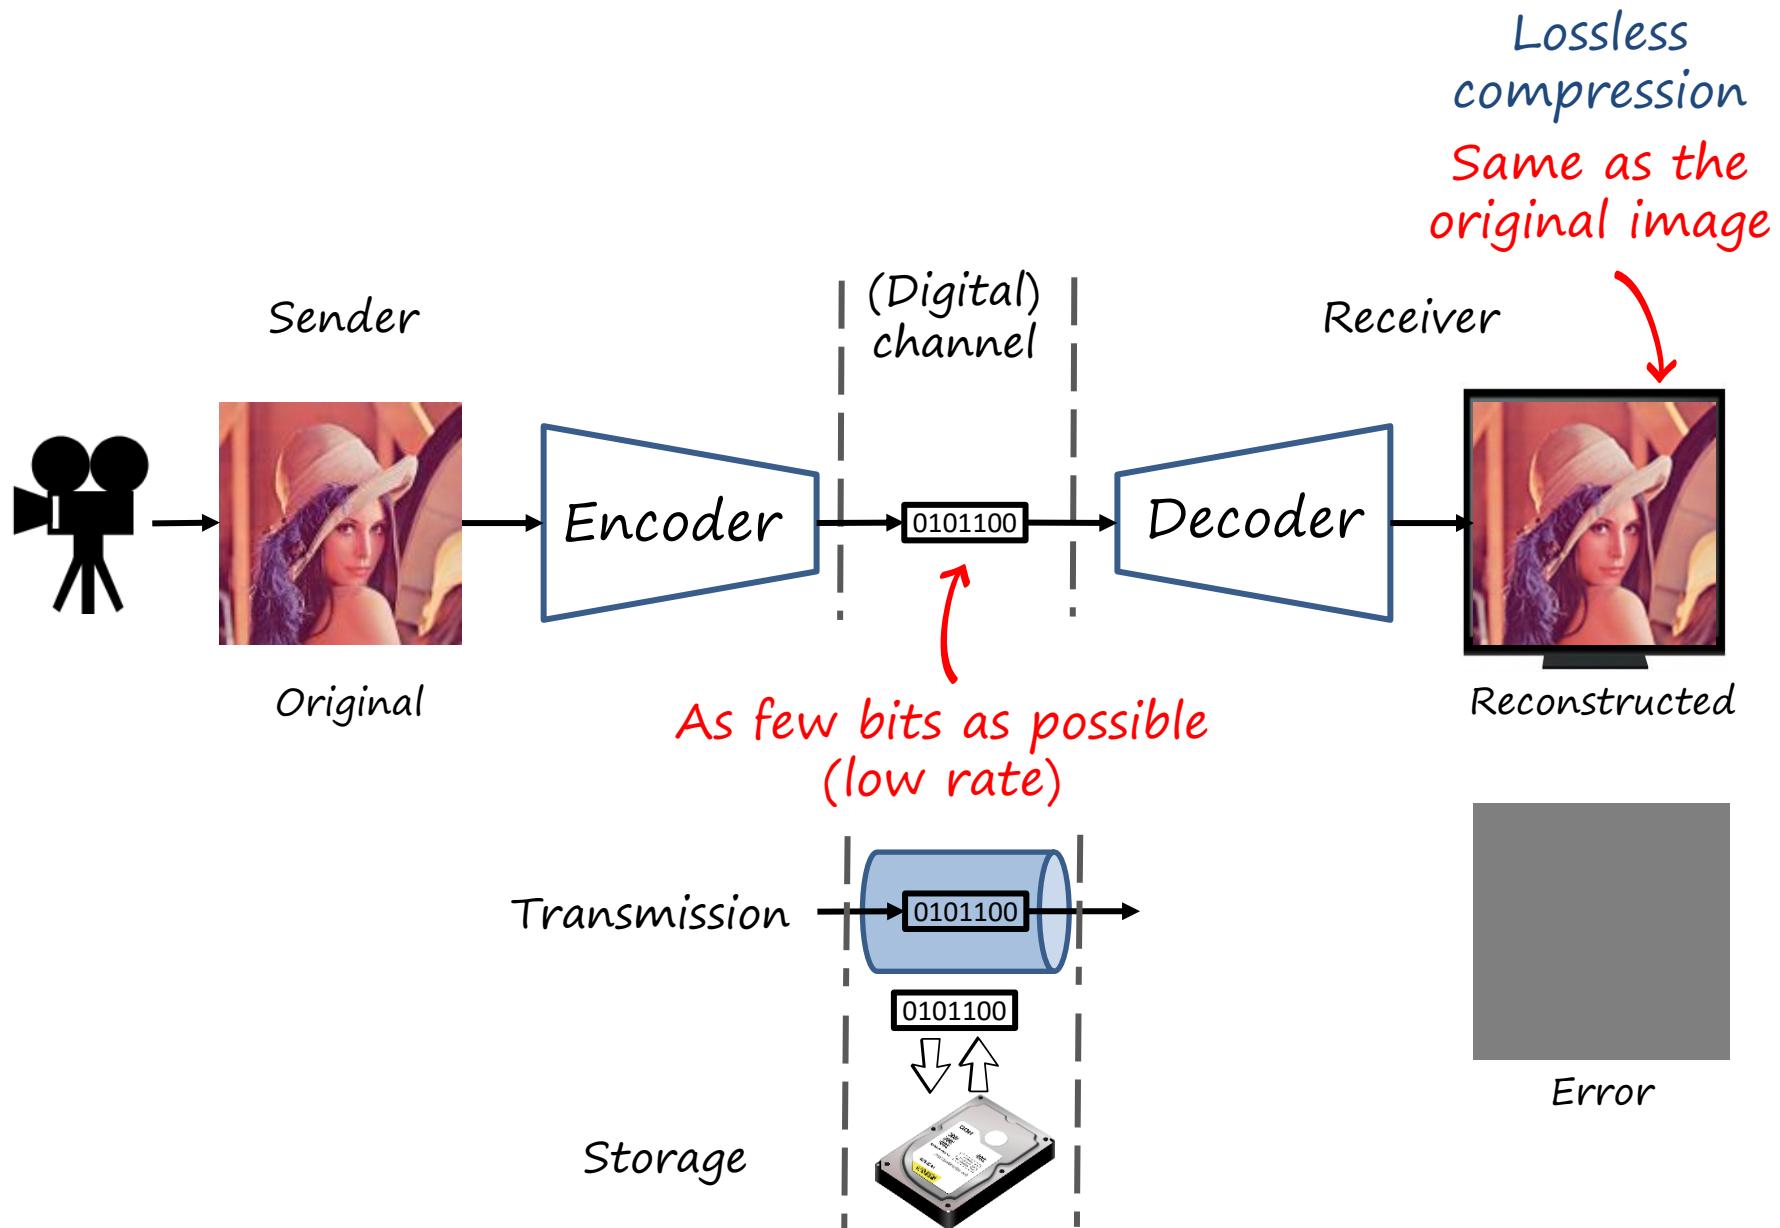
Outline

- Introduction: image coding
- Compression with neural networks
- Towards practical image compression

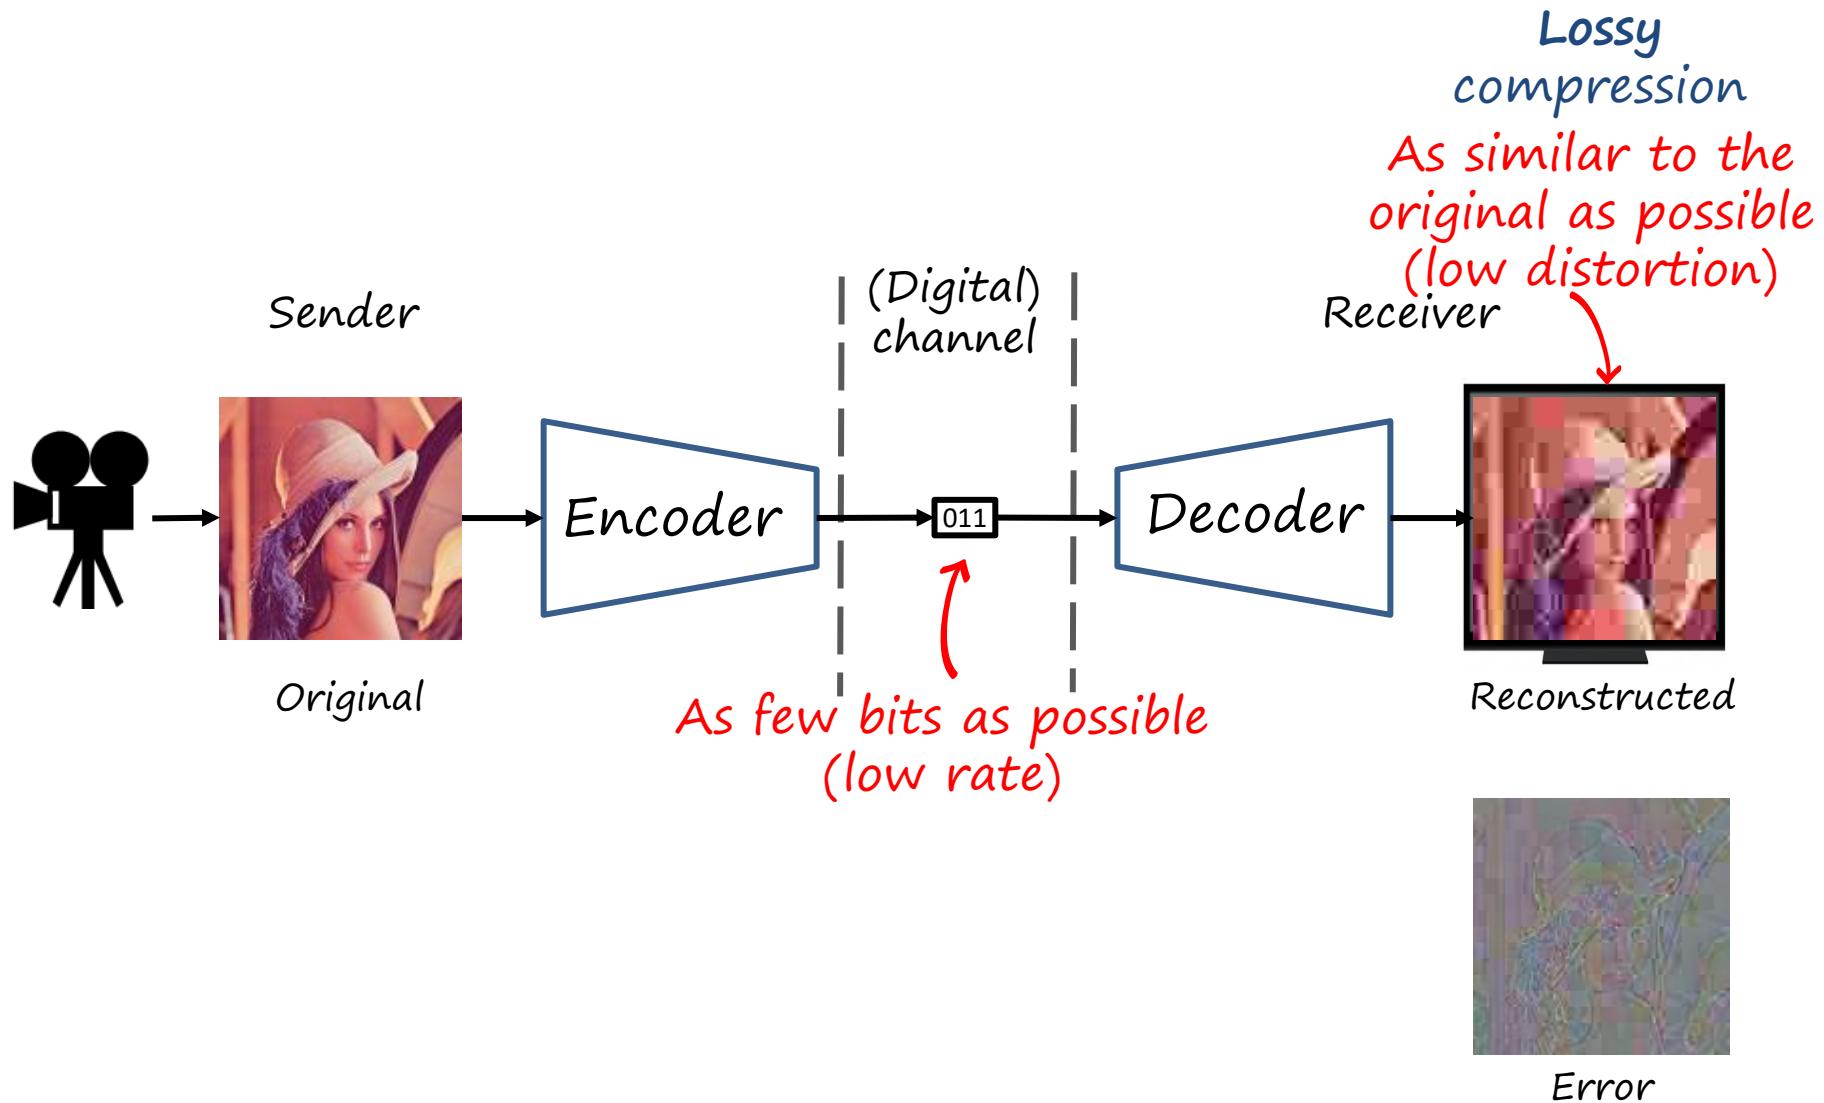
Outline

- Introduction: image coding
- Compression with neural networks
- Towards practical image compression

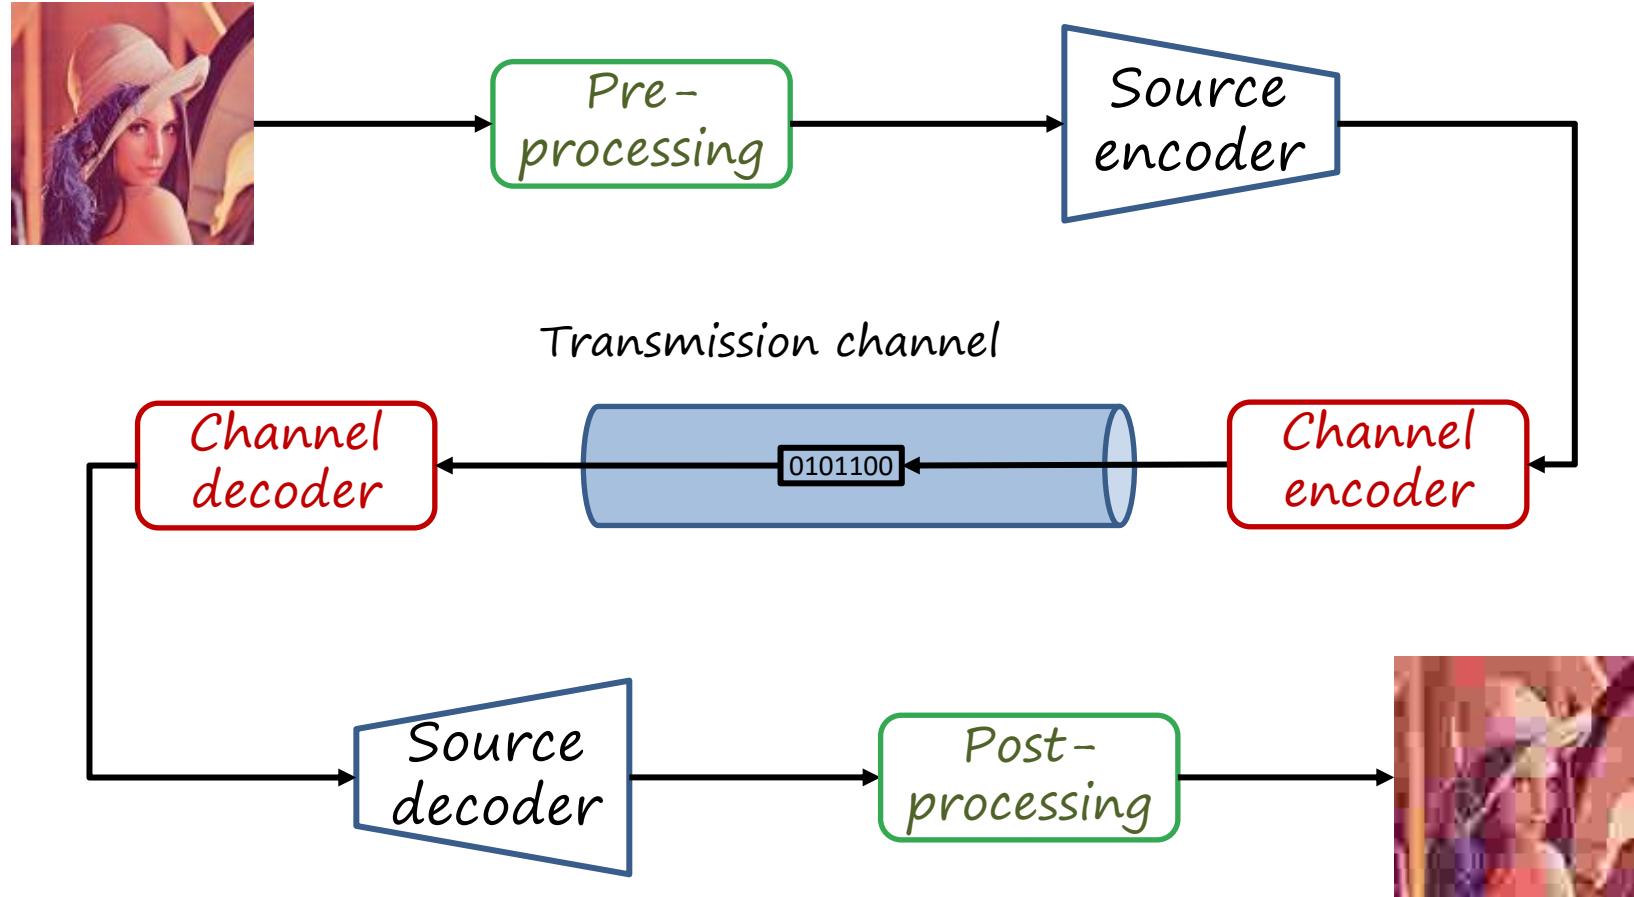
The visual communication problem



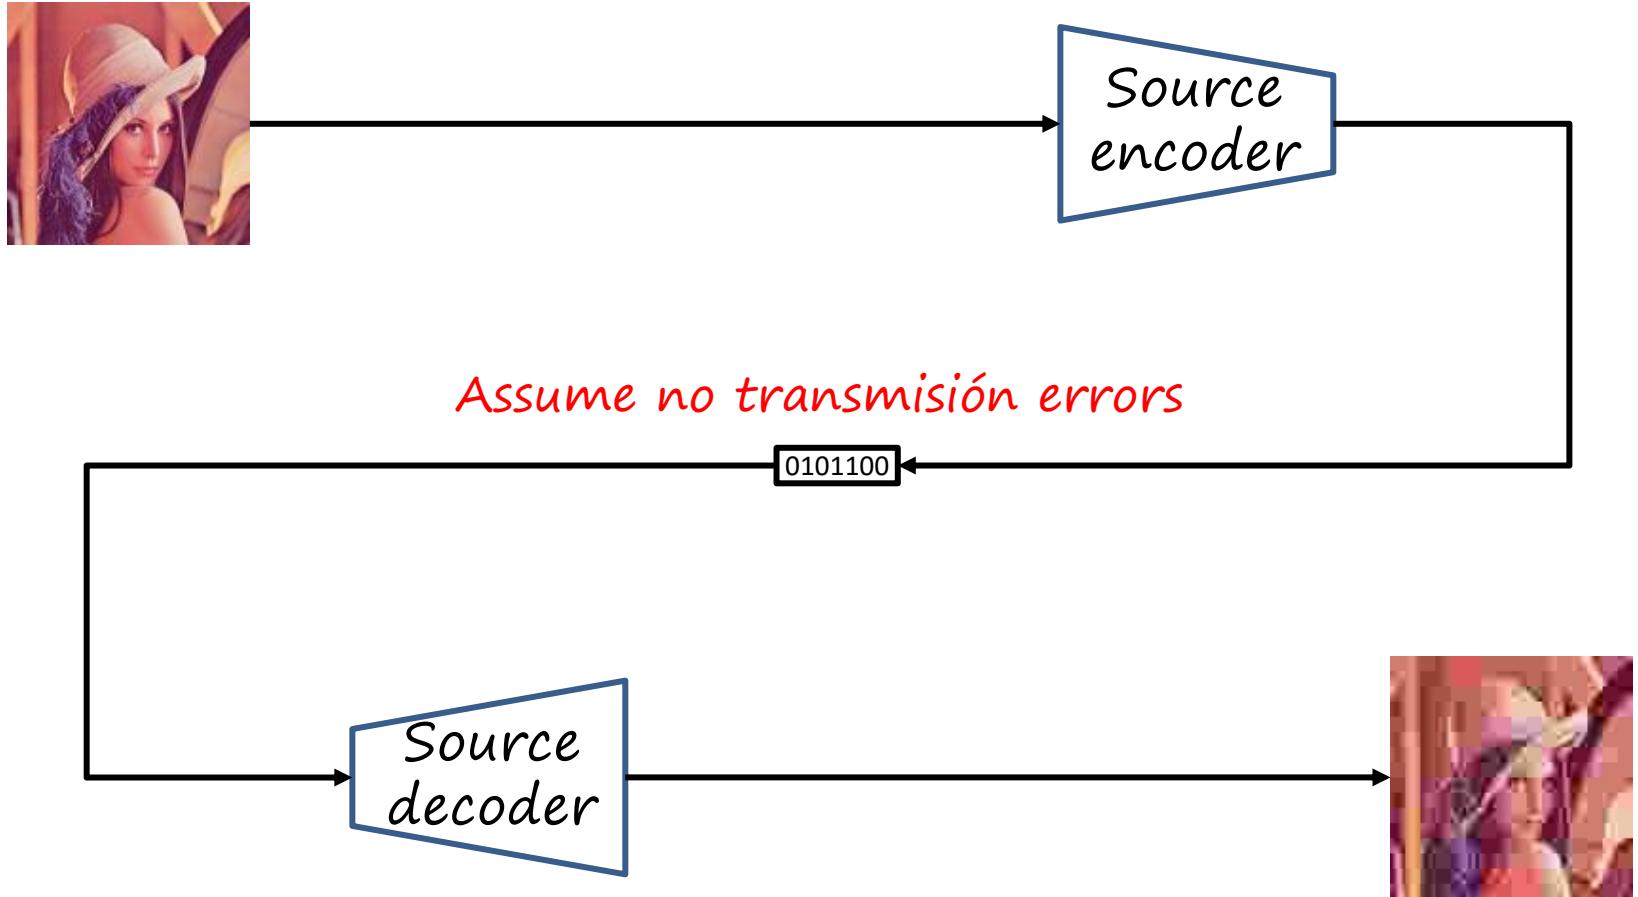
The visual communication problem



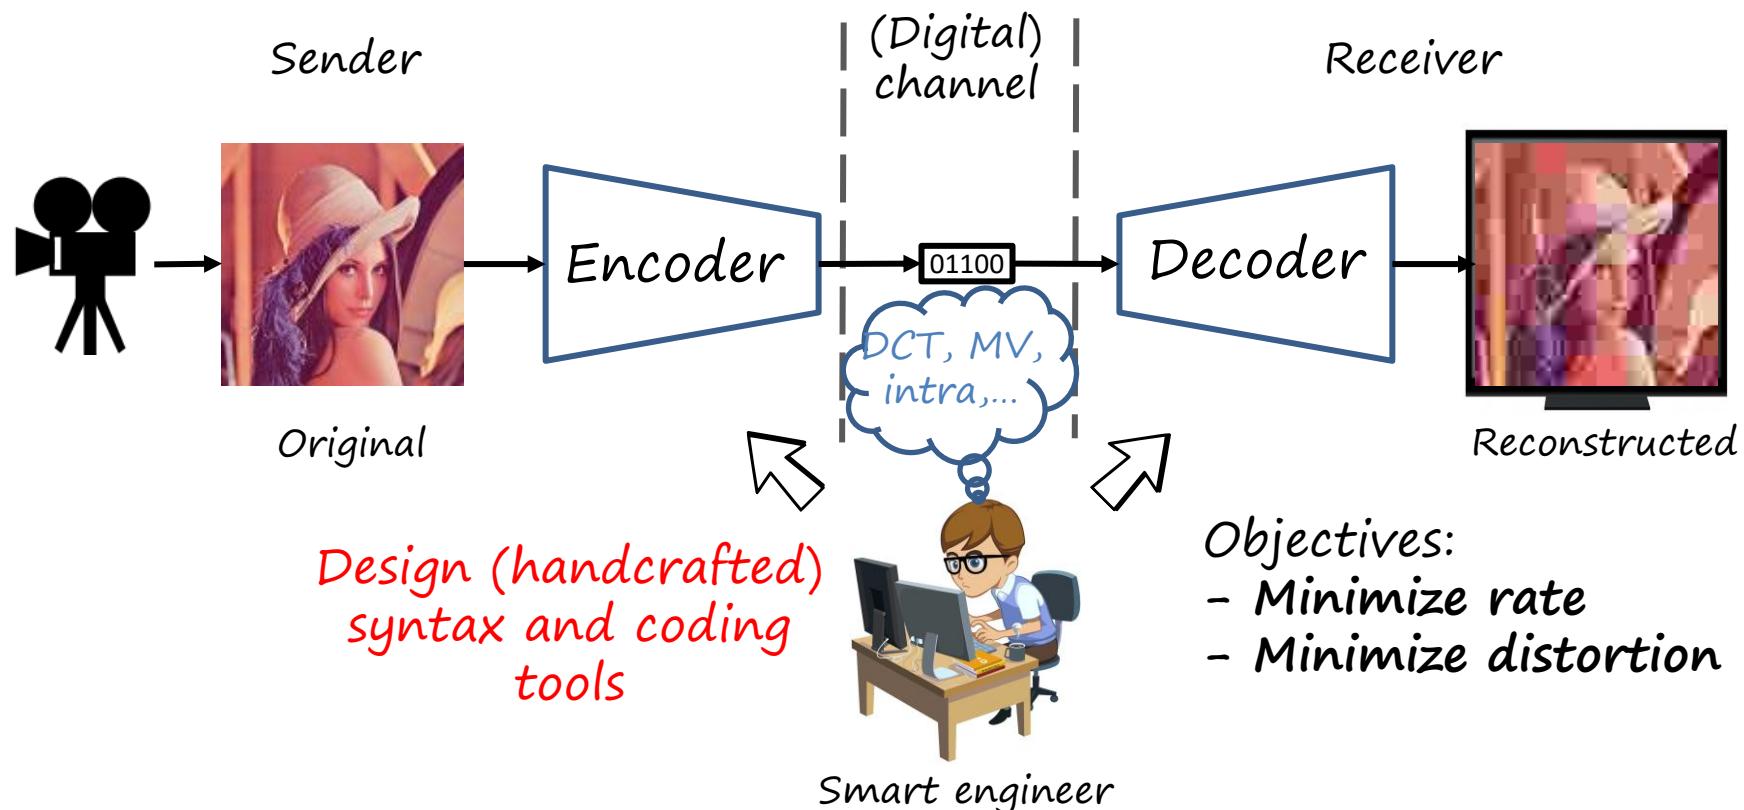
Pre/post-processing, source coding and channel coding



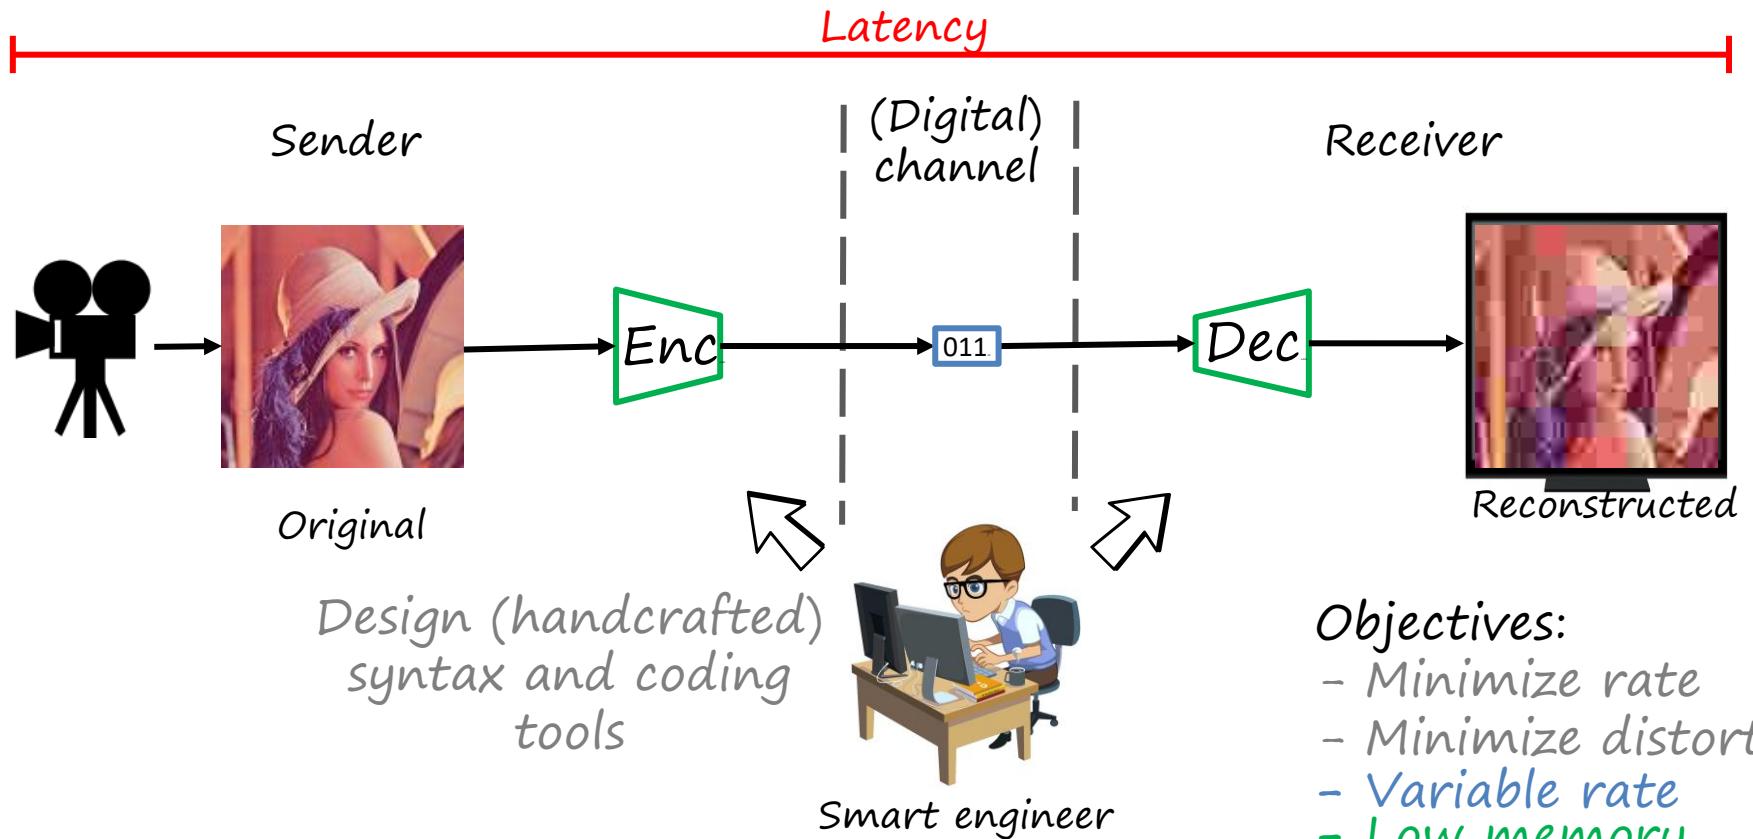
Source coding only



Developing traditional image/video codecs



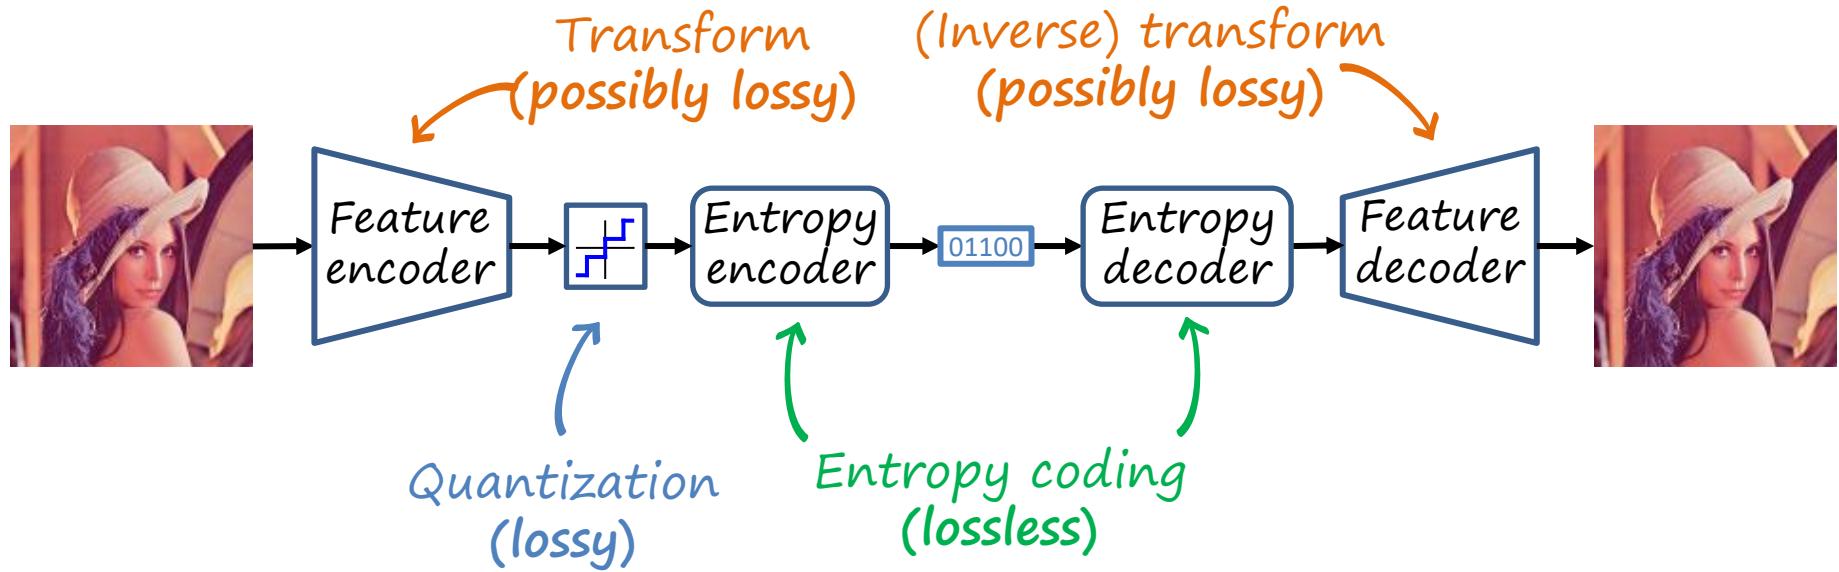
... for practical applications



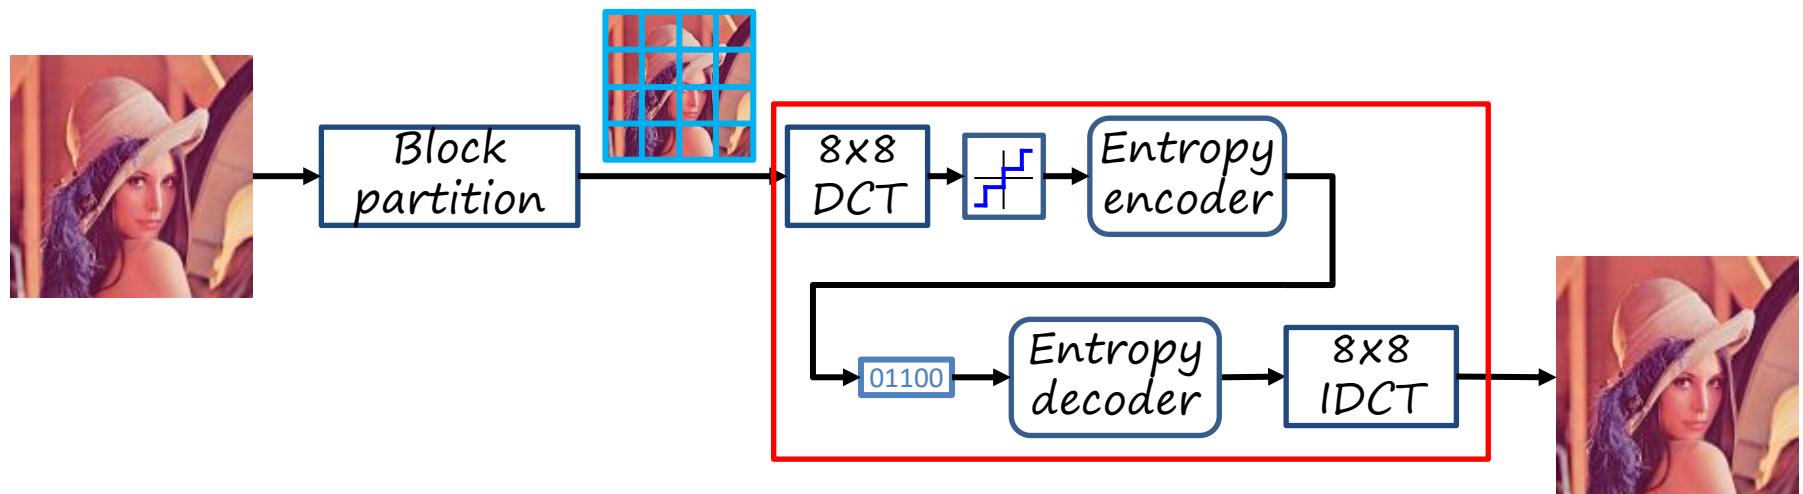
Objectives:

- Minimize rate
- Minimize distortion
- Variable rate
- Low memory
- Low computation
- Low latency
- Compatibility
- Domain-specific

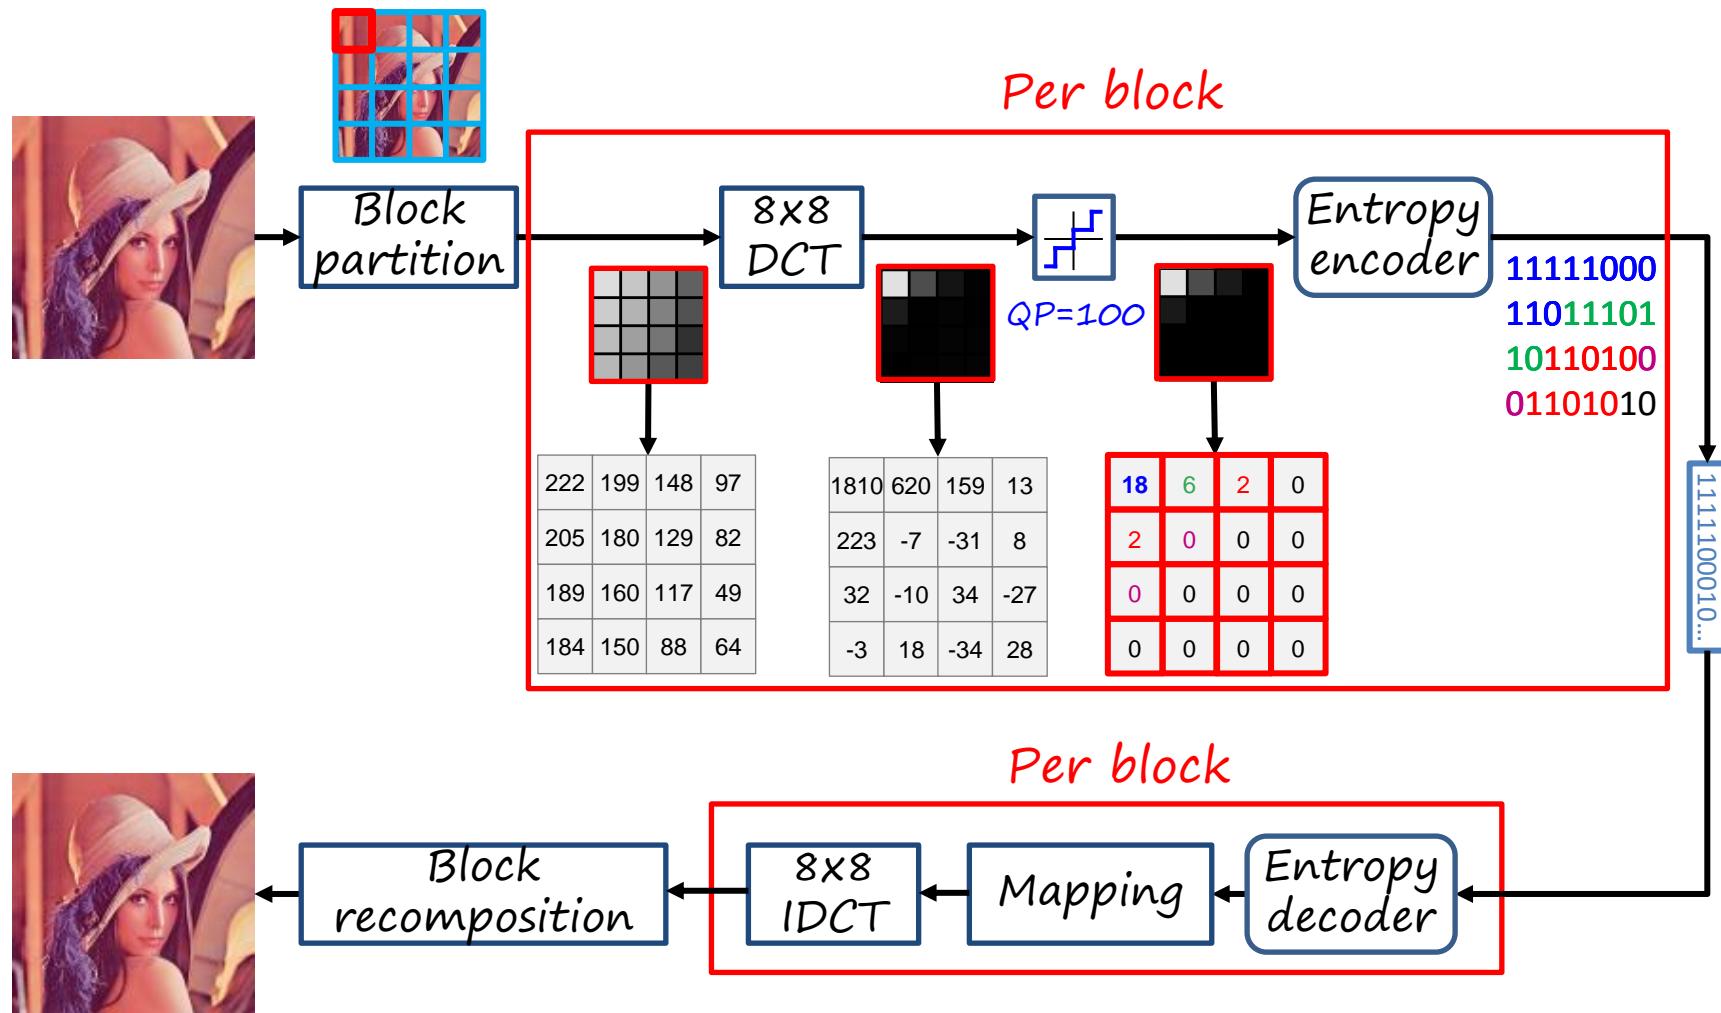
Transform coding pipeline



Example: block-based transform coding (e.g. JPEG, MPEG-2, H.264)

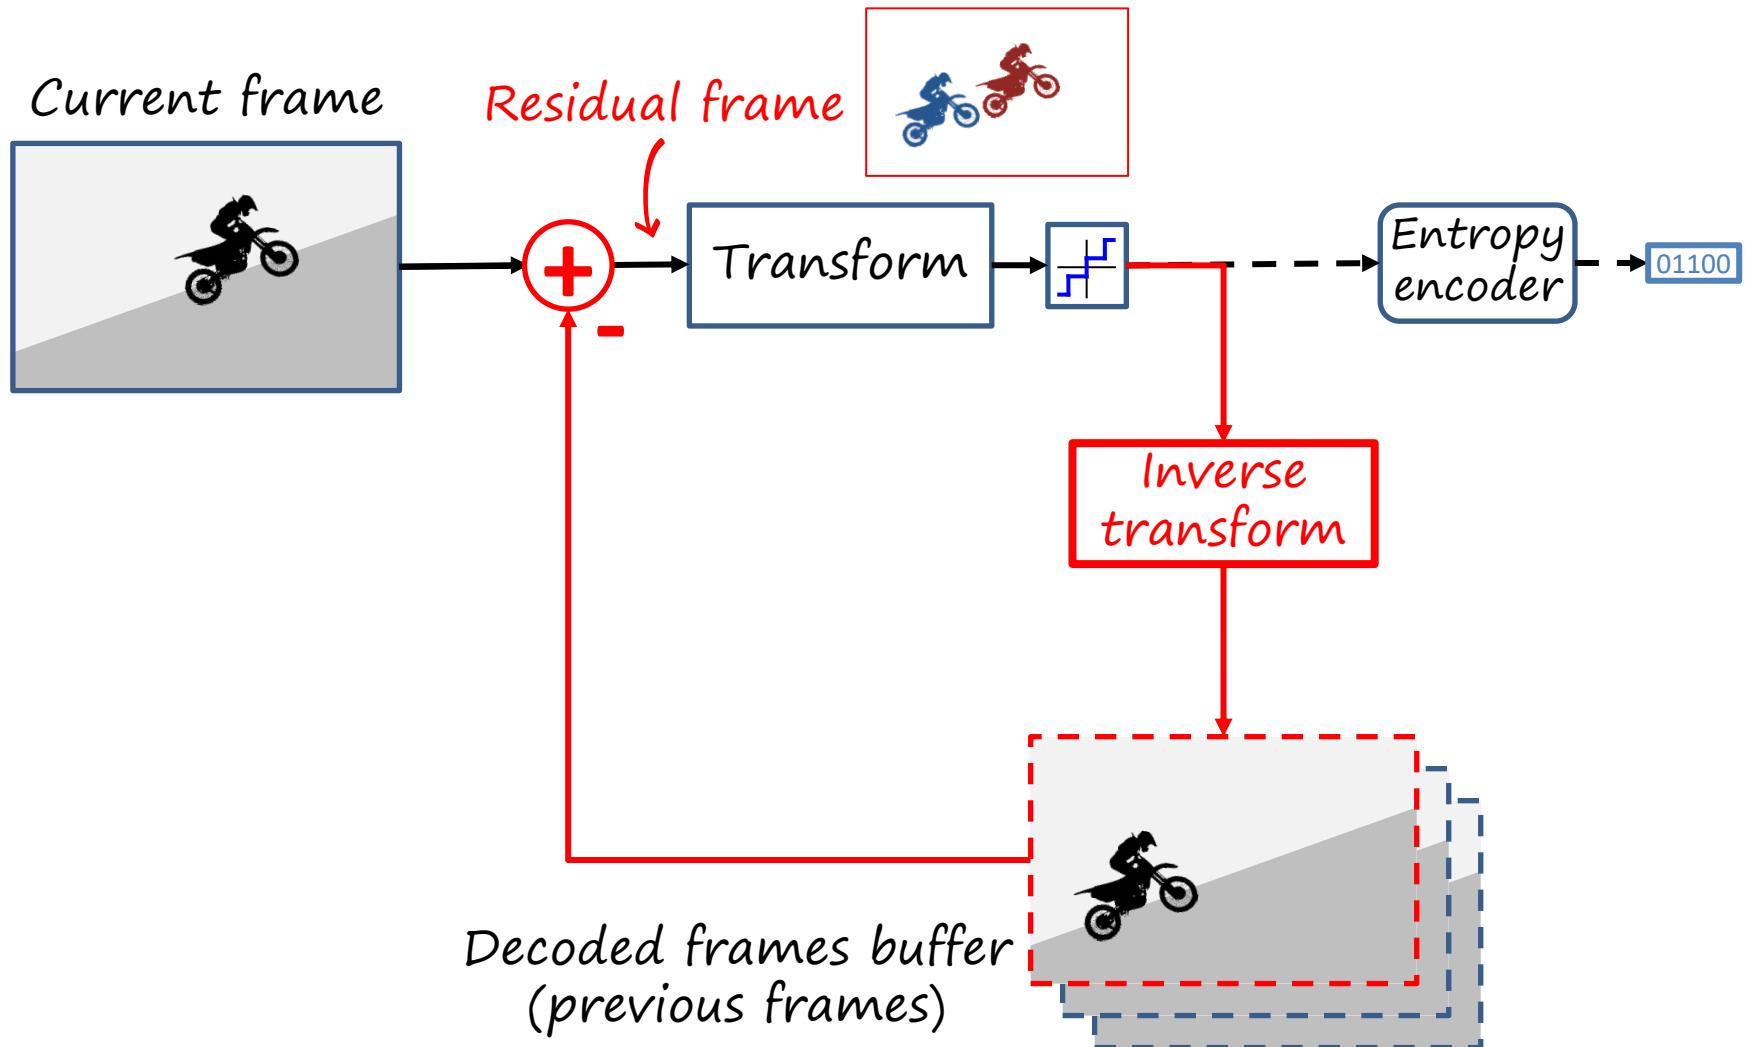


Transform coding pipeline: JPEG



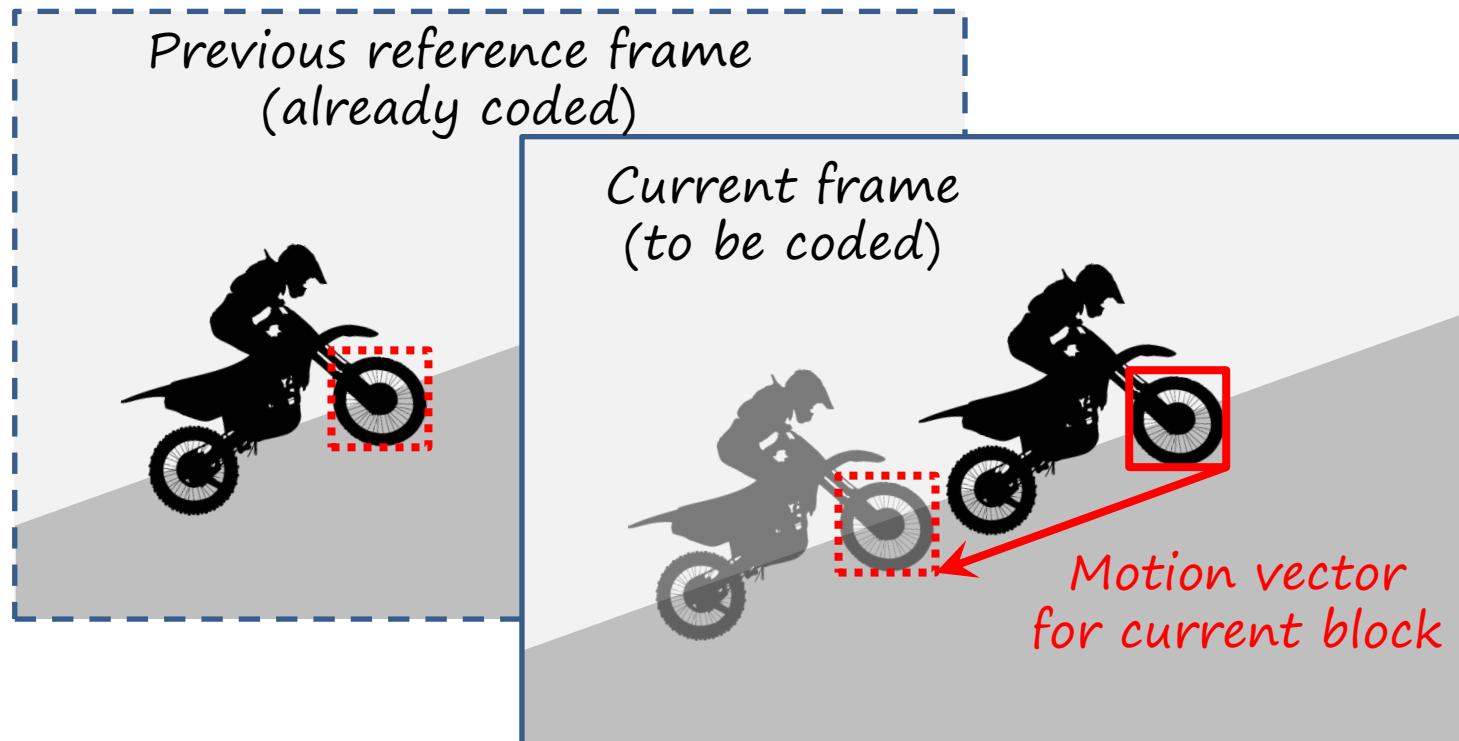
Coding video: temporal redundancy

Estimate current frame from previous coded ones



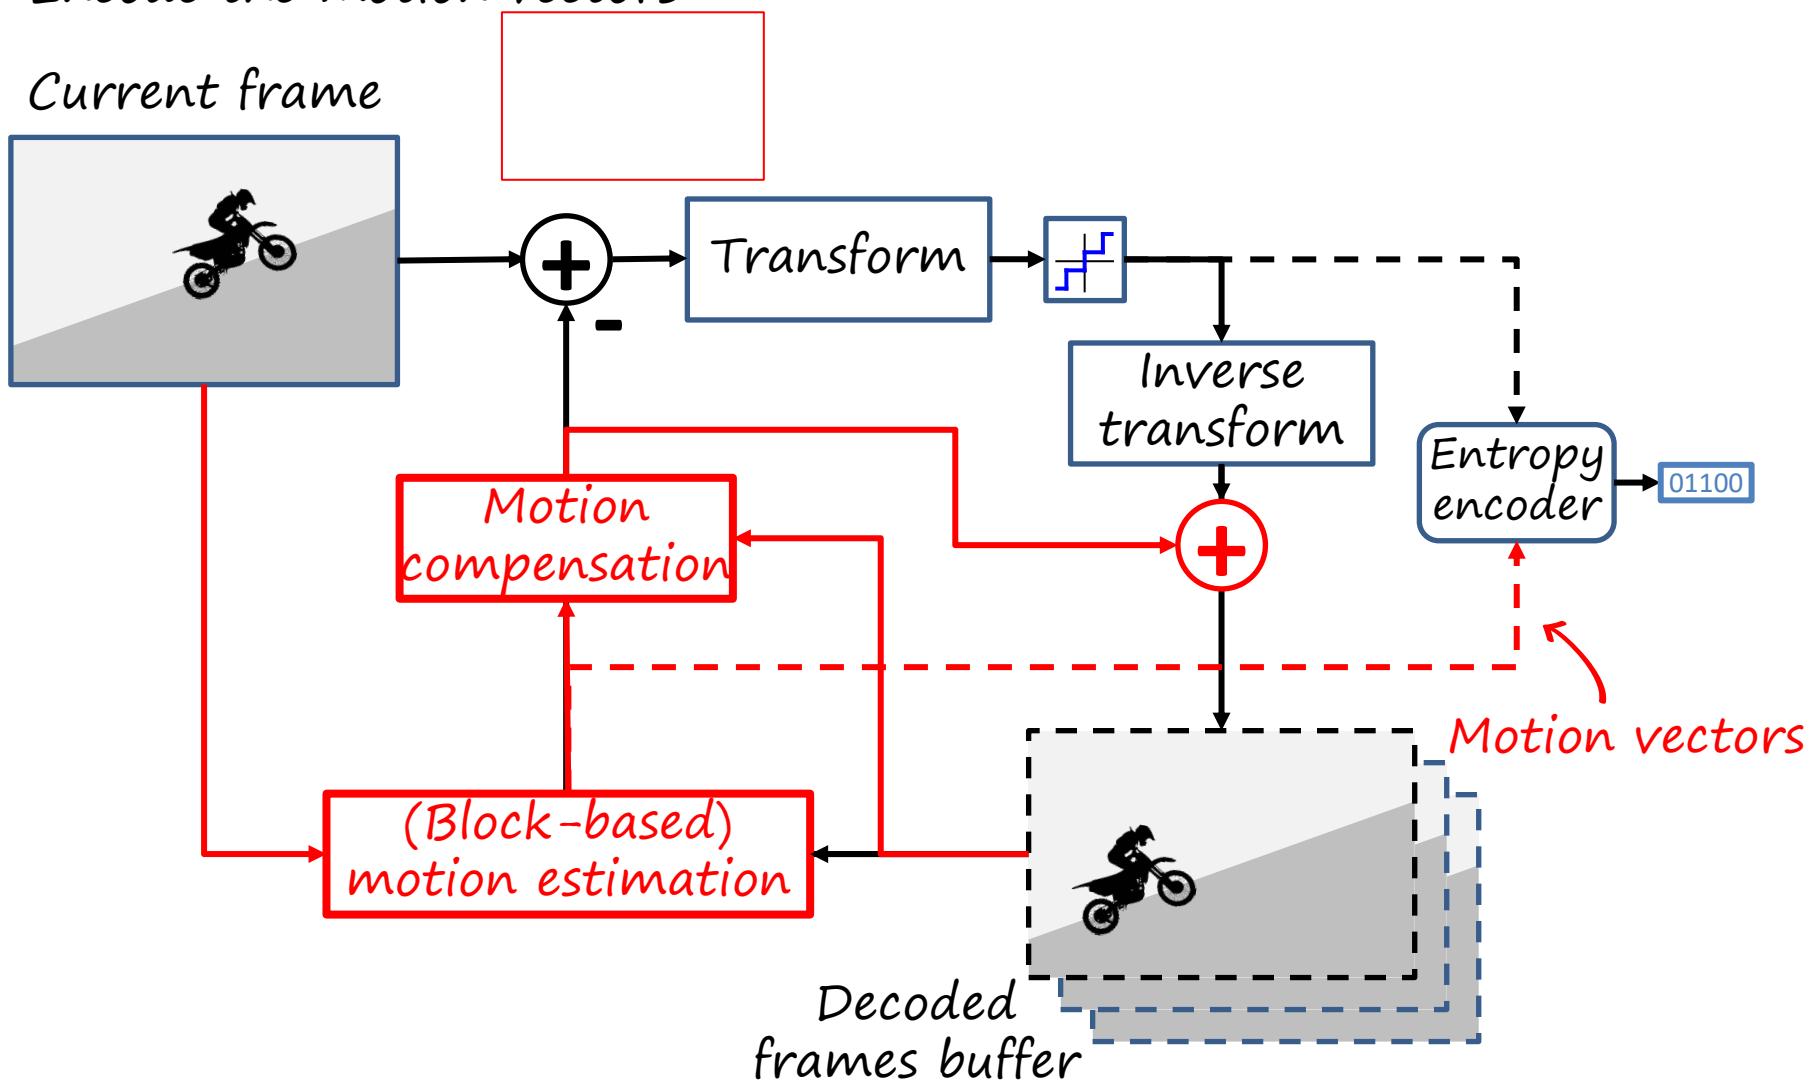
Motion-compensated prediction

Try to align frames: find most similar blocks in the reference frame



Motion-compensated video coding

Estimate current frame from previous coded ones
Encode the motion vectors

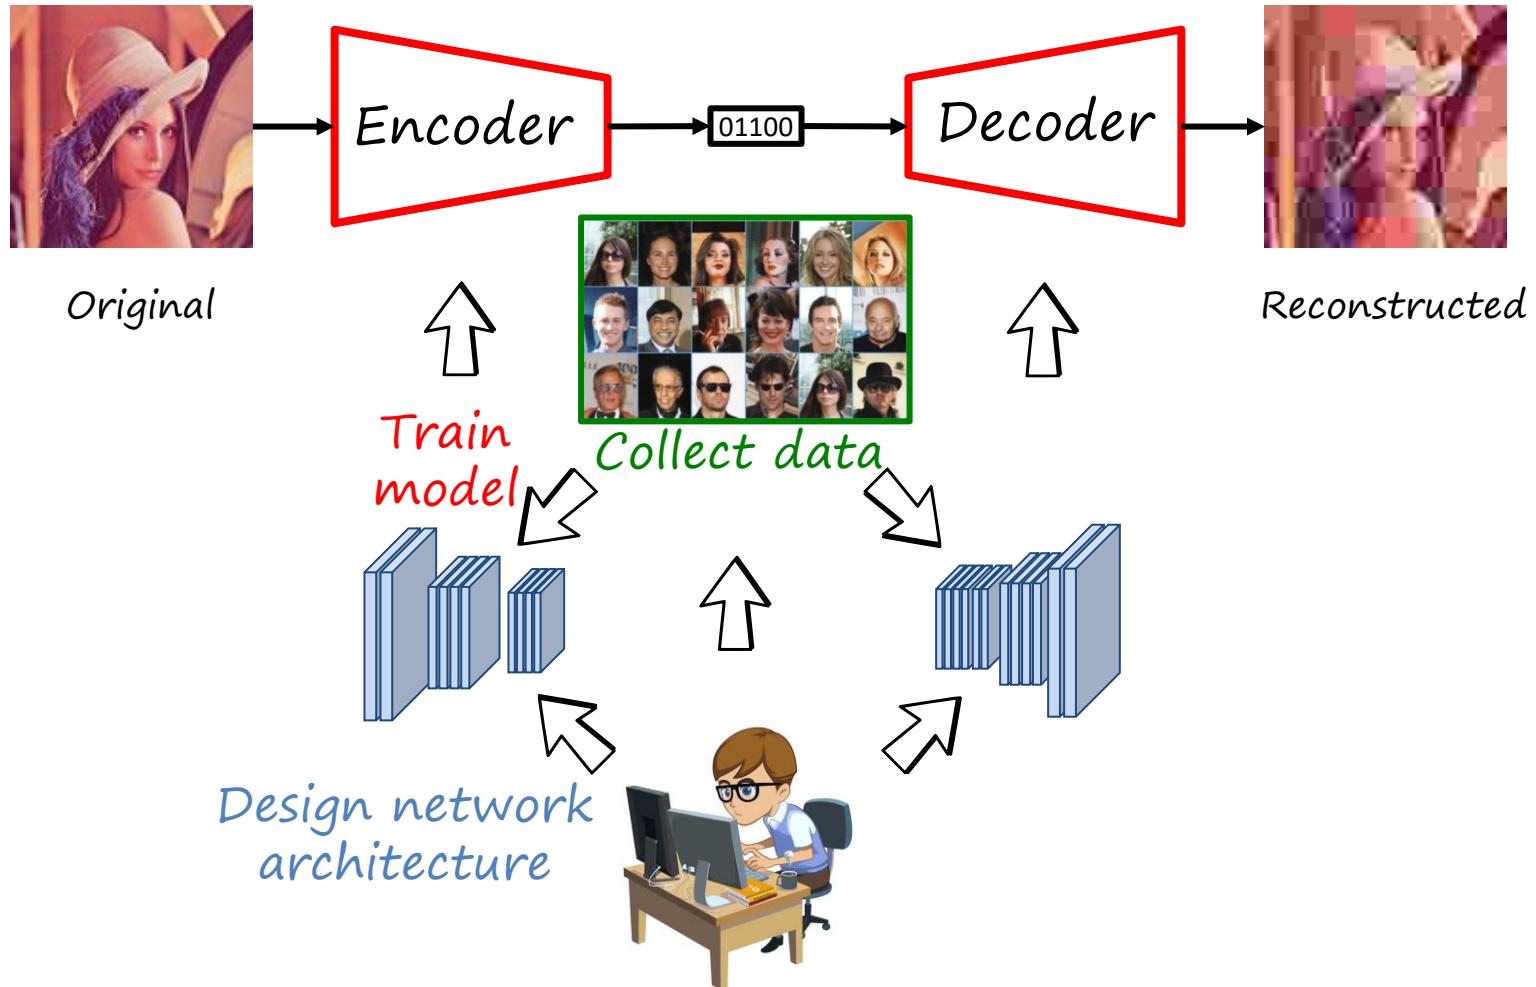


Outline

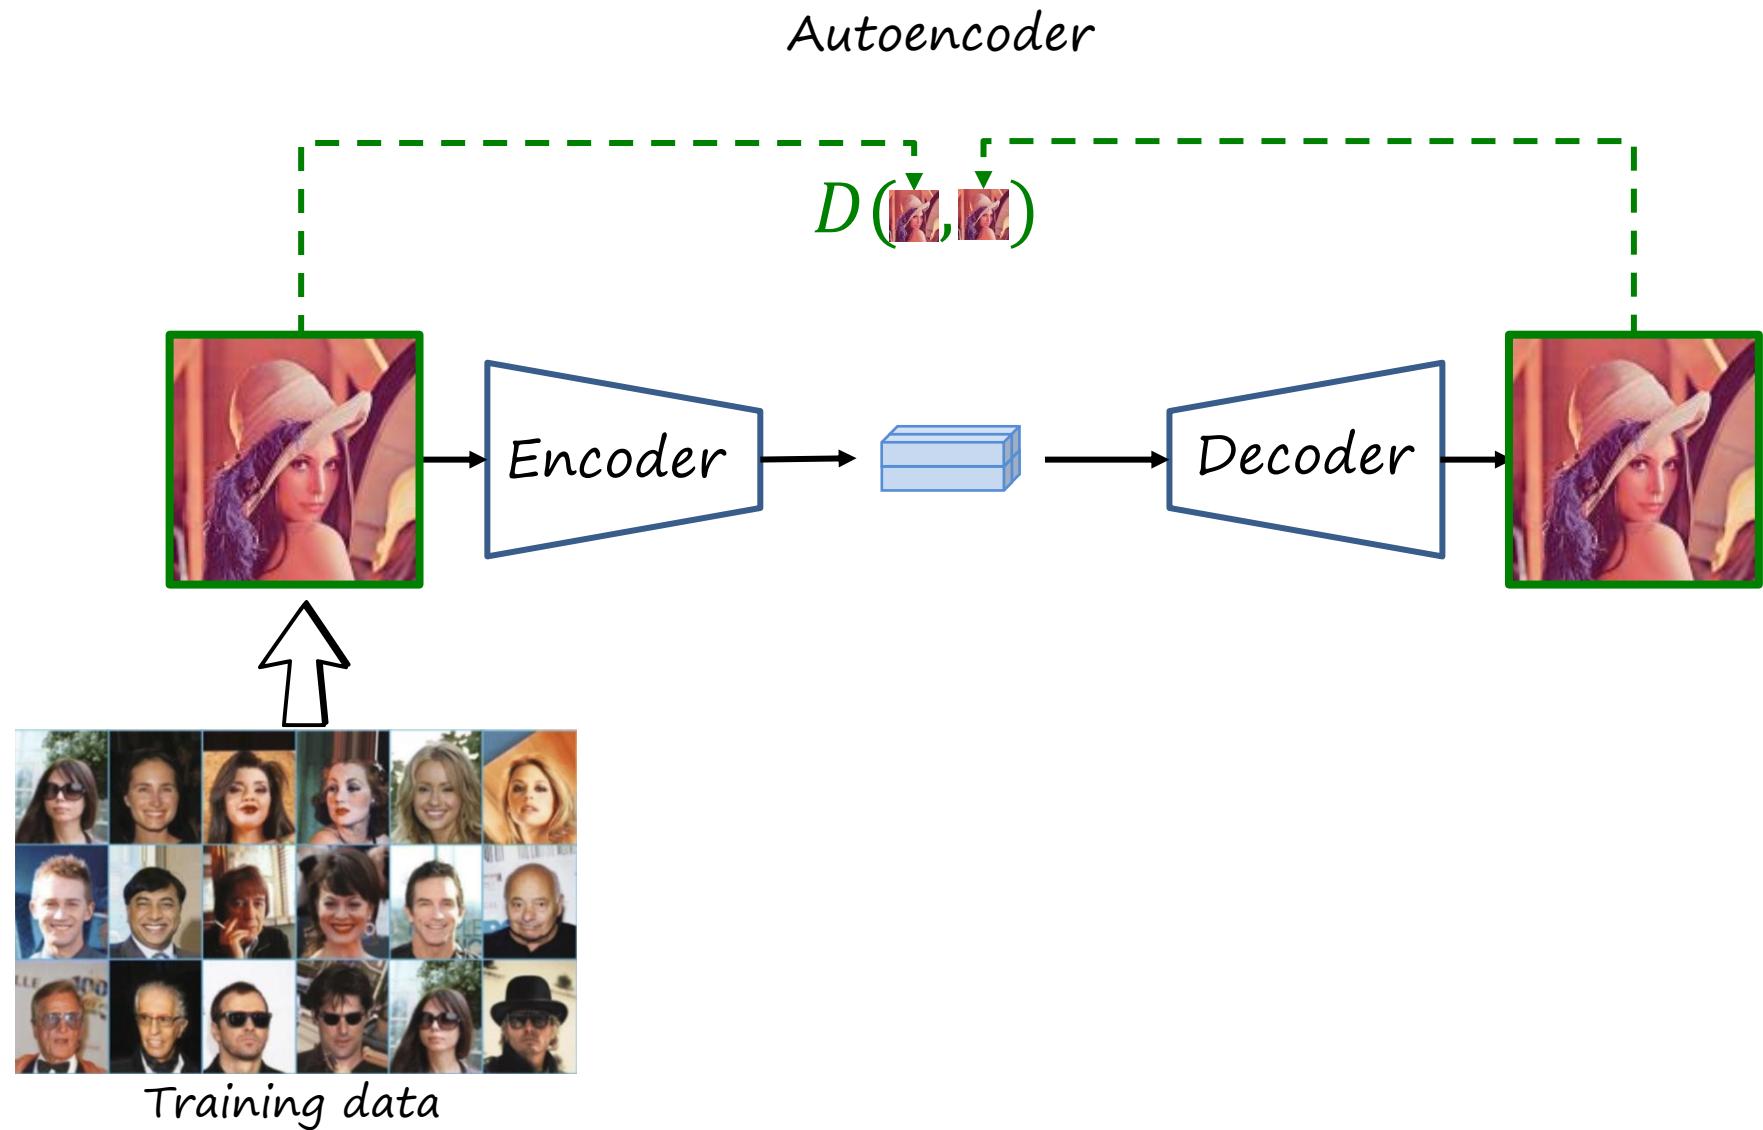
- Introduction: image coding
- **Compression with neural networks**
- Towards practical image compression

Neural image codecs

- Coding tools and syntax are **parametric** and **learned**
- Encoders/decoders and probability models **deep neural networks**

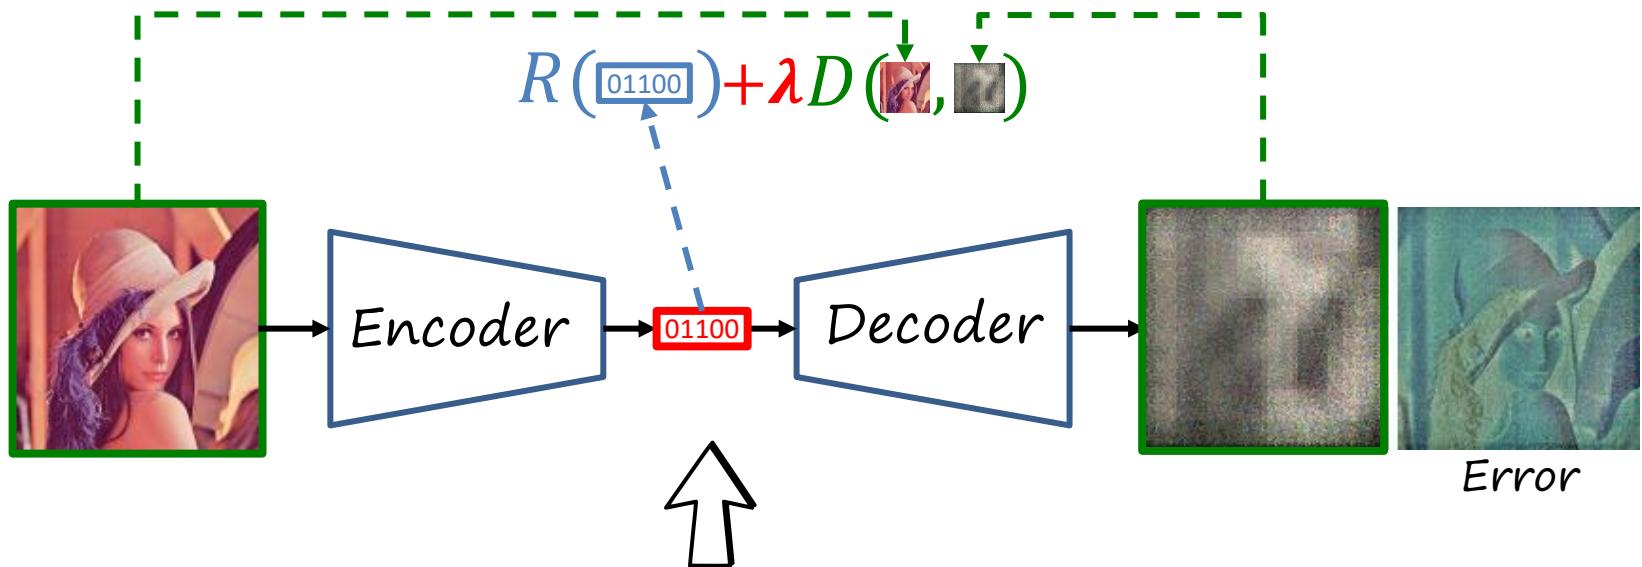


Neural image compression



Neural image compression

Compressive autoencoder (CAE) [Theis2017, Balle2017]
(autoencoder+binary latent representation)



Optimize a weighted
rate-distortion loss
(λ controls the tradeoff)

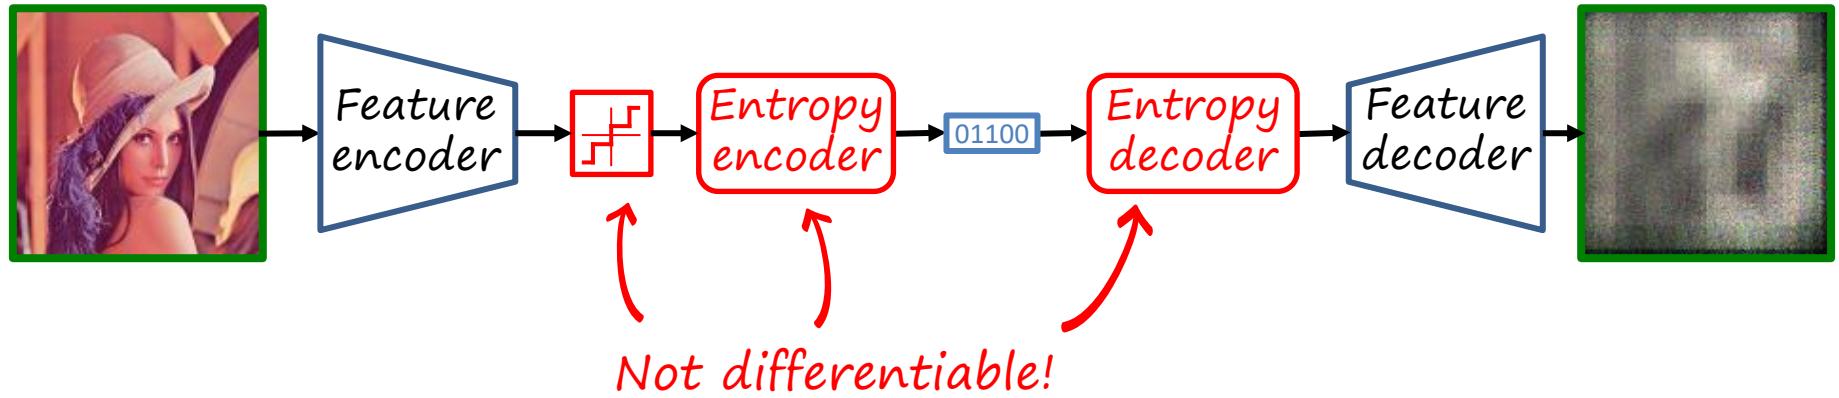


Distortion is typically
mean square error (MSE)

Training data

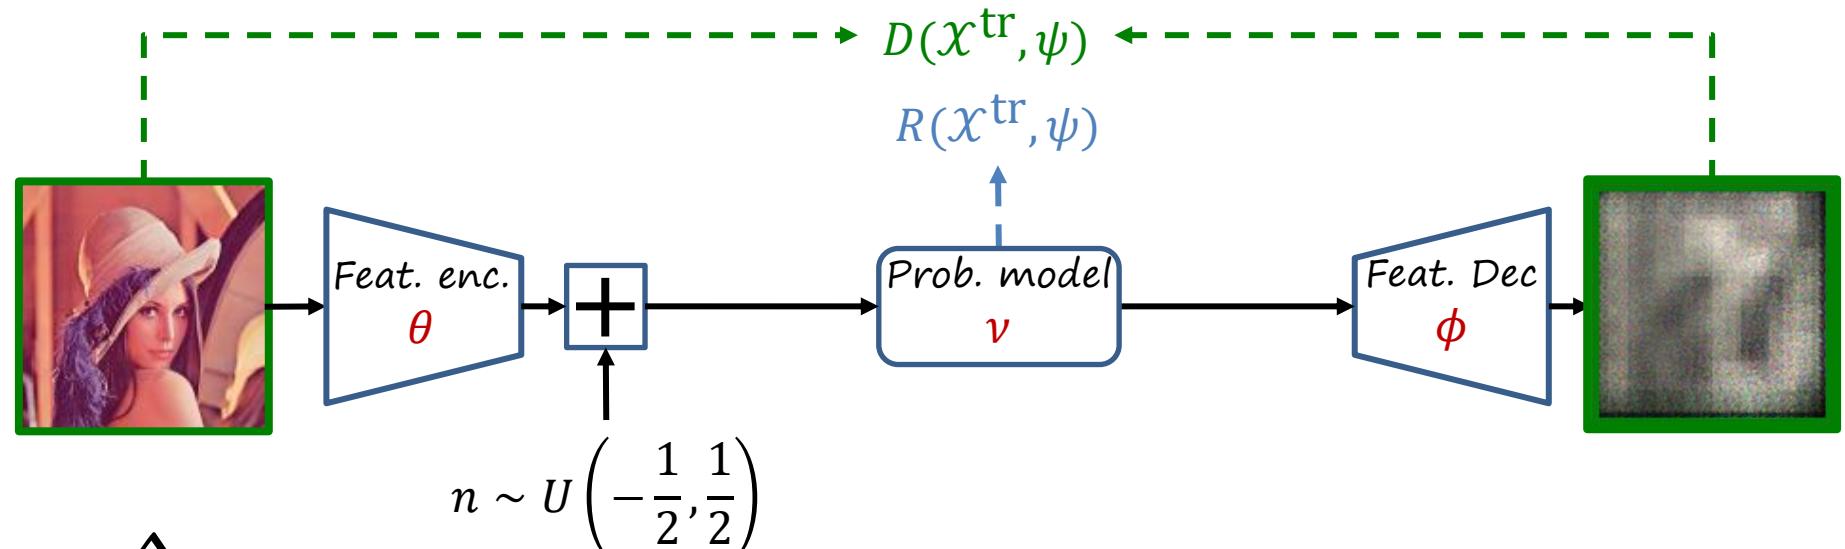
Typical pipeline

Compressive autoencoder (CAE) [Theis2017, Balle2017]
(autoencoder+quantization+entropy coding)



Architecture (training)

Use differentiable proxies for end-to-end training



Model parameters

$$\psi = (\theta, \phi, \nu)$$

Loss

$$J(\mathcal{X}^{\text{tr}}, \psi; \lambda) = R(\mathcal{X}^{\text{tr}}, \psi) + \lambda D(\mathcal{X}^{\text{tr}}, \psi)$$

Optimization problem

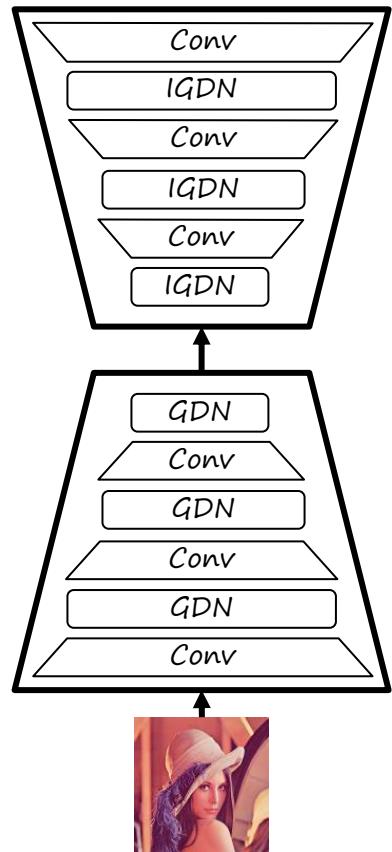
$$\psi^* = \min_{\psi} J(\mathcal{X}^{\text{tr}}, \psi; \lambda)$$



Training data \mathcal{X}^{tr}

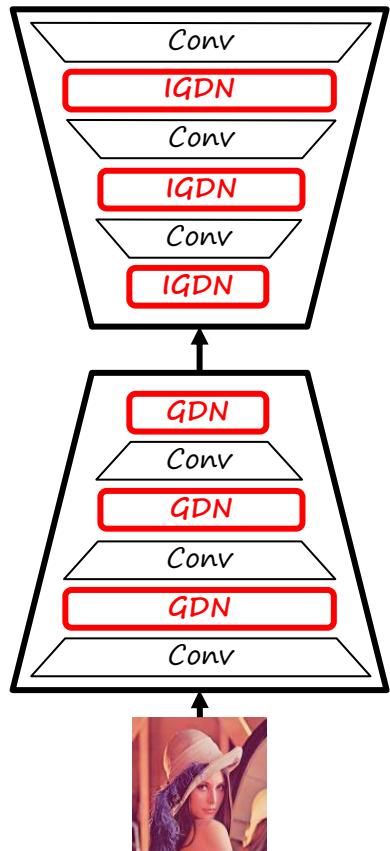
Autoencoder architecture

Balle et al.
[ICLR2017]



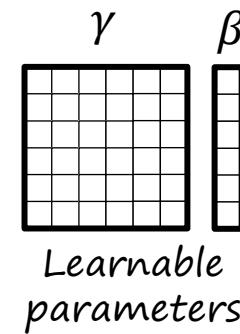
Autoencoder architecture

Balle et al.
[ICLR2017]

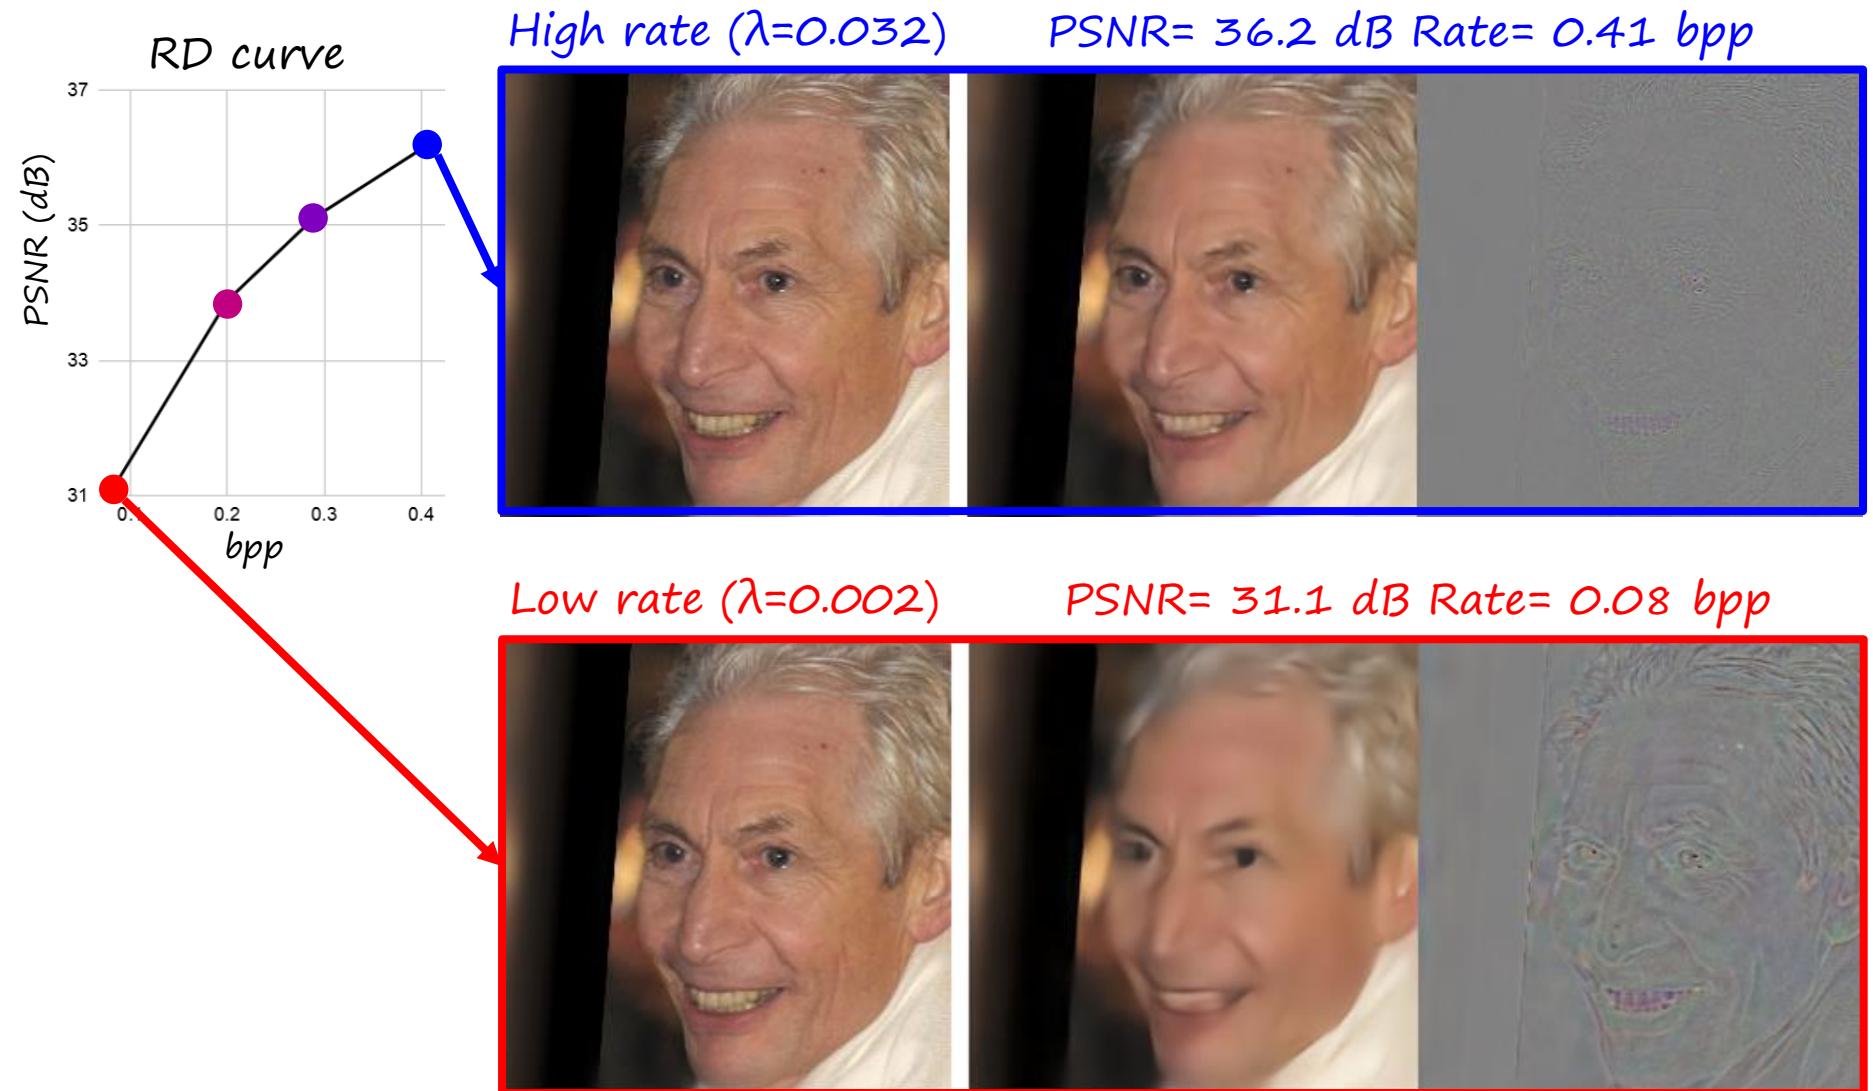


Generalized divisive normalization (GDN) [Balle2016]

$$\hat{y}_i = \frac{y_i}{(\beta_i + \sum_j \gamma_{ij} y_j^2)^{1/2}}$$



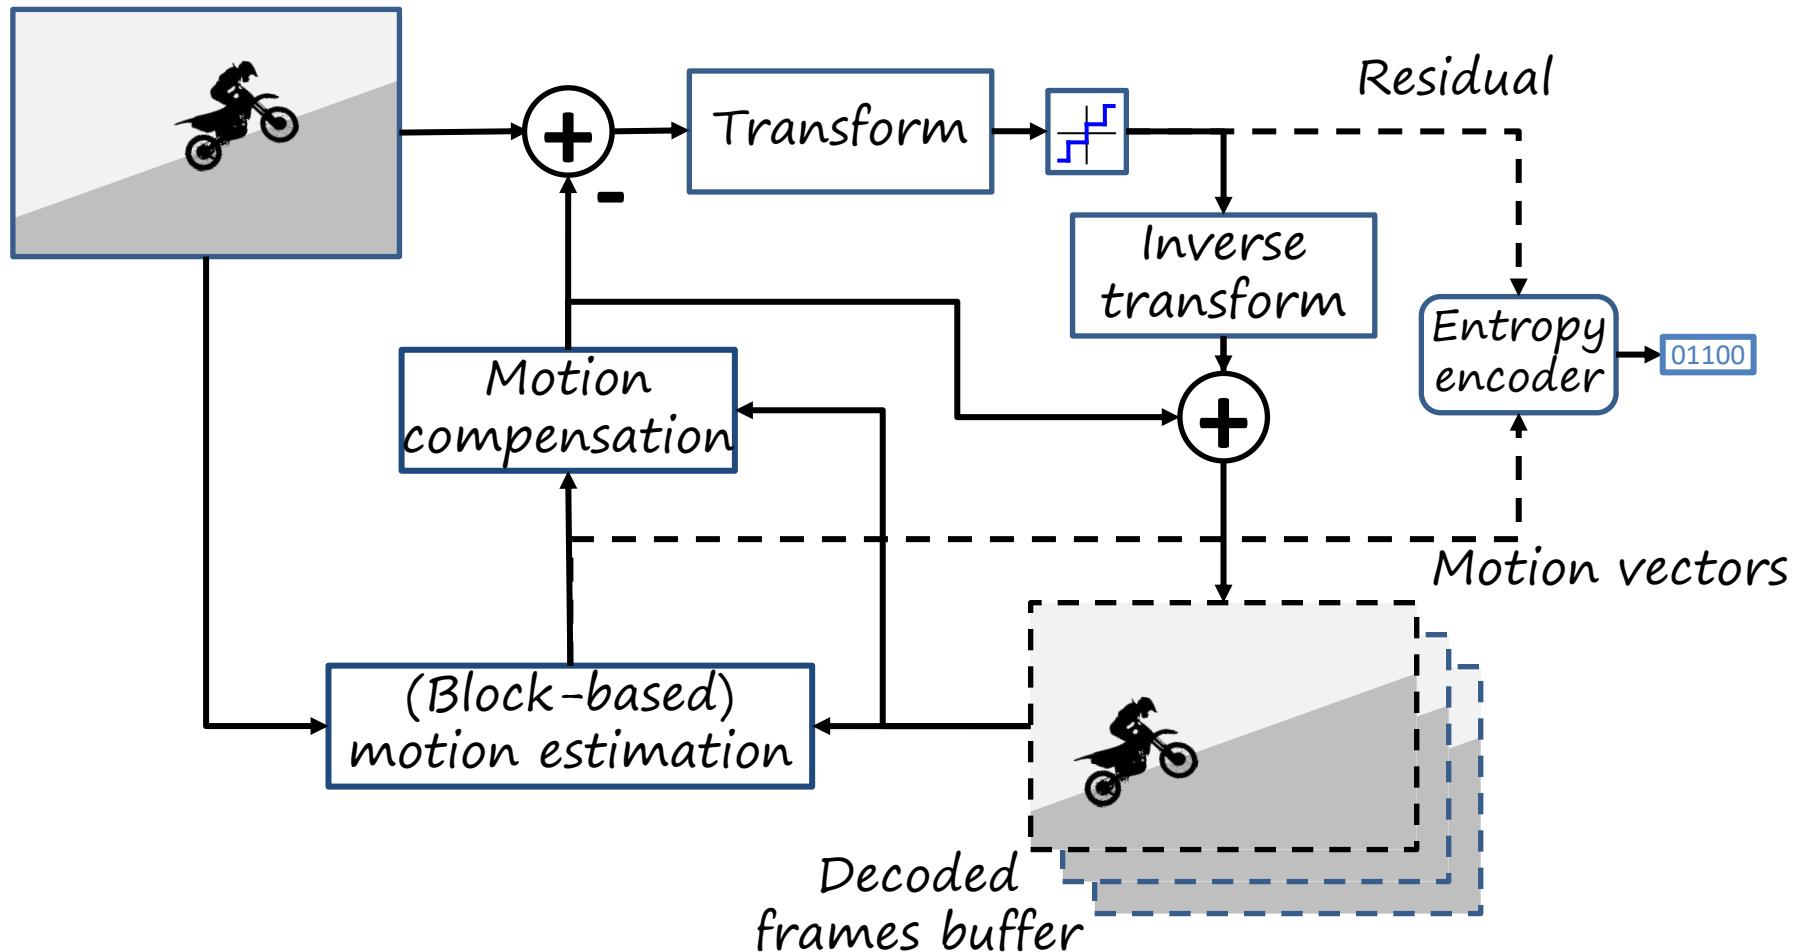
Rate-distortion tradeoff λ



Traditional video compression

Replace modules by trainable neural networks

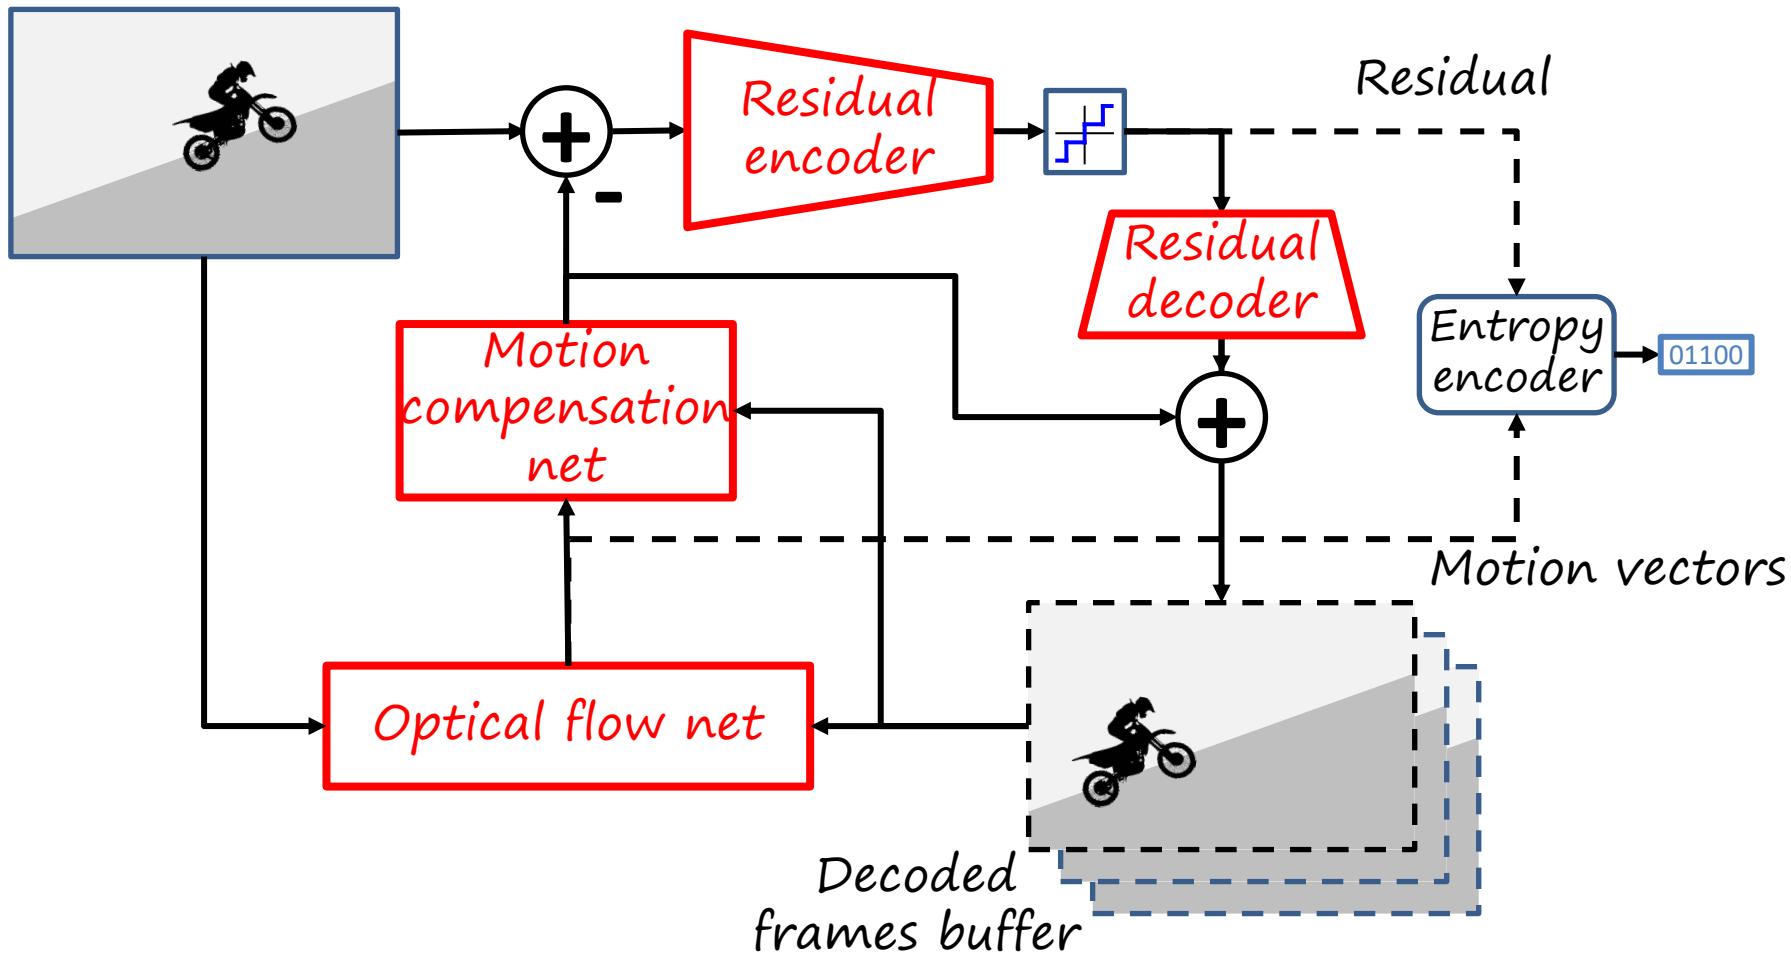
Current frame



Neural video compression

Replace modules by trainable neural networks

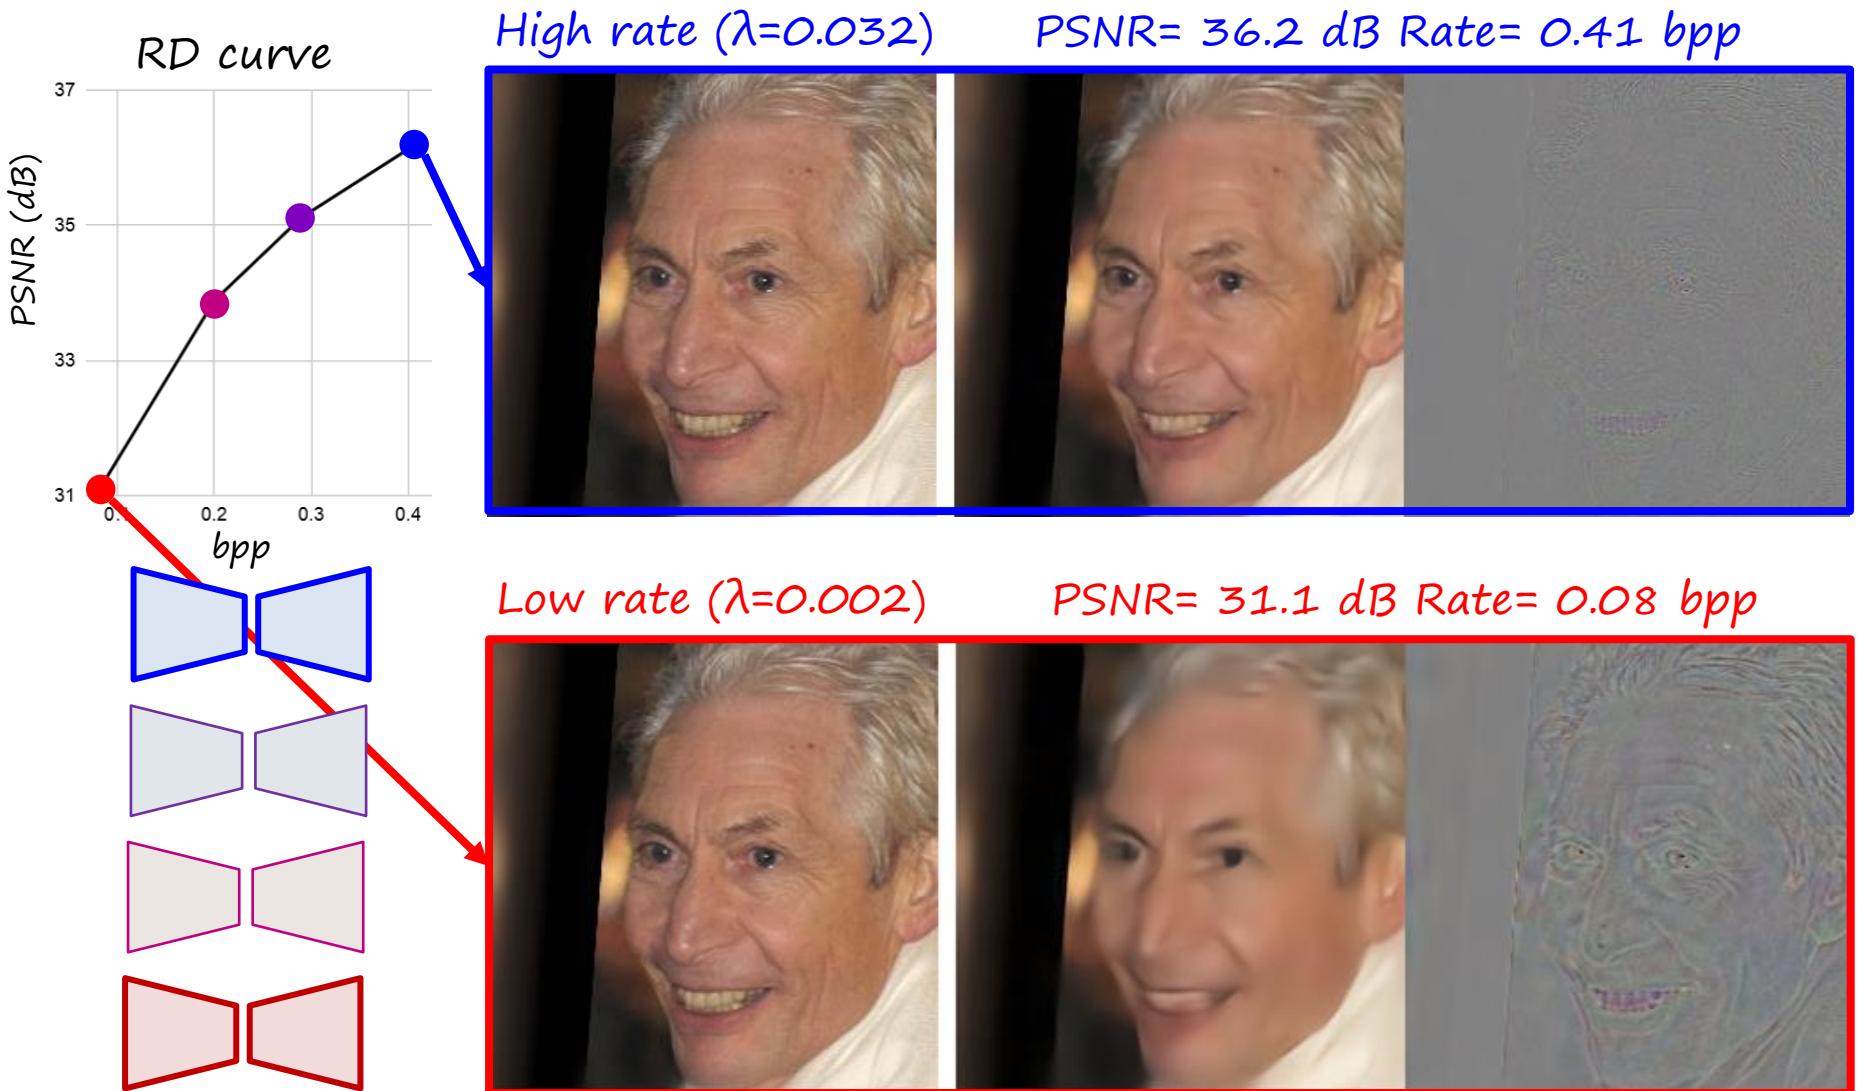
Current frame



Outline

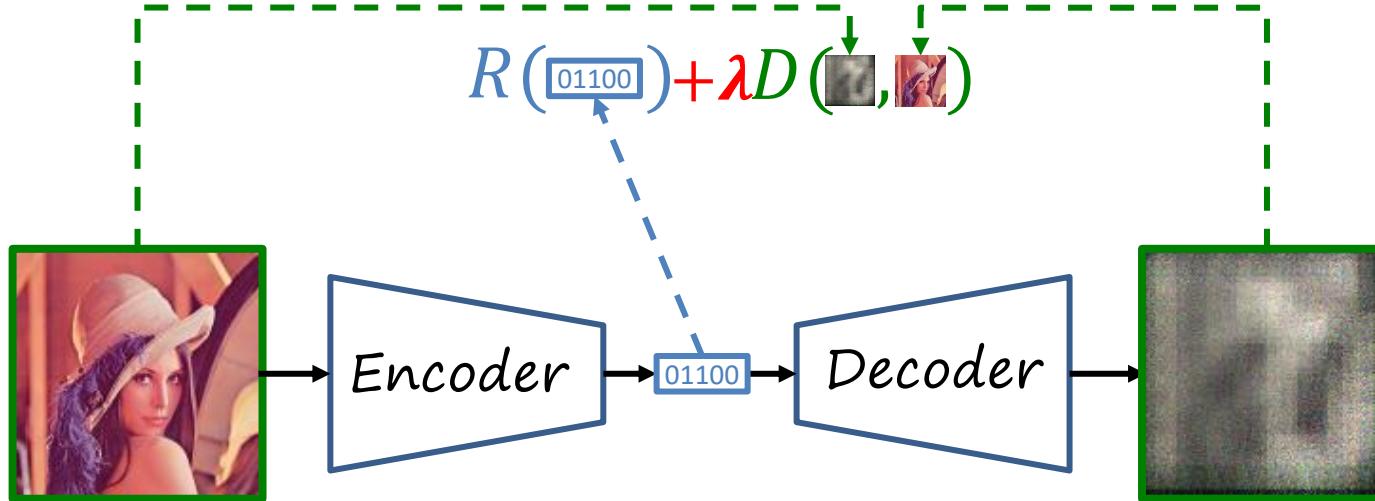
- Introduction: image coding
- Compression with neural networks
- Towards practical image compression

Rate-distortion tradeoff λ



Problems: total memory, total training time

Is neural image compression practical?



Limitations

- λ is fixed
- Heavy encoders/decoders

Practical neural image compression?

- Minimize rate ✓
- Minimize distortion ✓

- Variable rate ✗
- Low memory ✗
- Low computation ✗
- Low latency ✗

MAE
[SPL2020]

SlimCAE
[CVPR2021]

Other practical considerations

- Domain-specific codecs (e.g. videoconference, screencast)
- Back./forw. compatibility (with legacy encoders/decoders)

DANICE
[CLIC2021]

[SPL2020] [Variable Rate Deep Image Compression with Modulated Autoencoder](#), Signal Processing Letters 2020

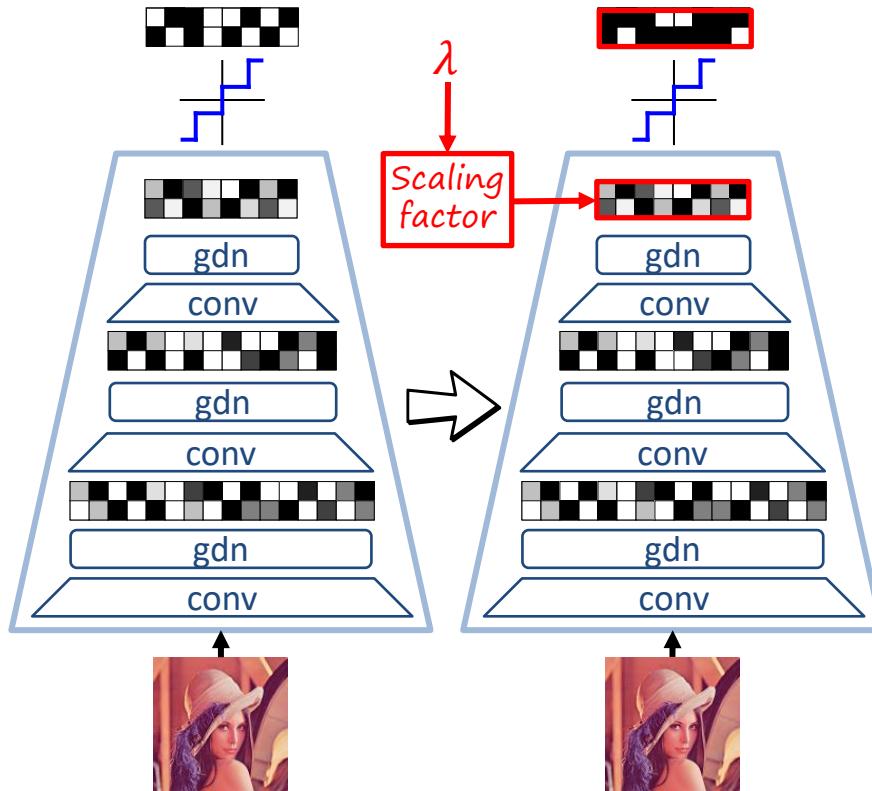
[CVPR2021] [Slimmable compressive autoencoders for practical image compression](#), CVPR 2021

[CLIC2021] [DANICE: Domain adaptation without forgetting in neural image compression](#), CLIC 2021 at CVPR 2021

Variable rate with modulated autoencoders

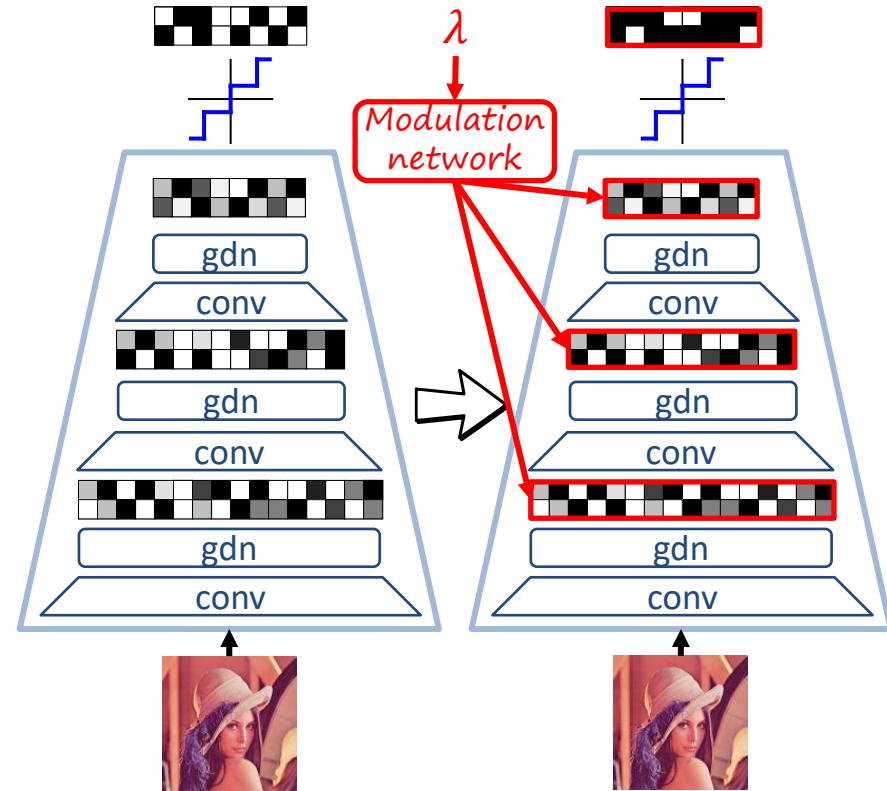
Objective: one single model for multiple λ

Bottleneck scaling [Theis2017]



- Minimize rate ✓
- Minimize distortion ✓
- Variable rate ✓

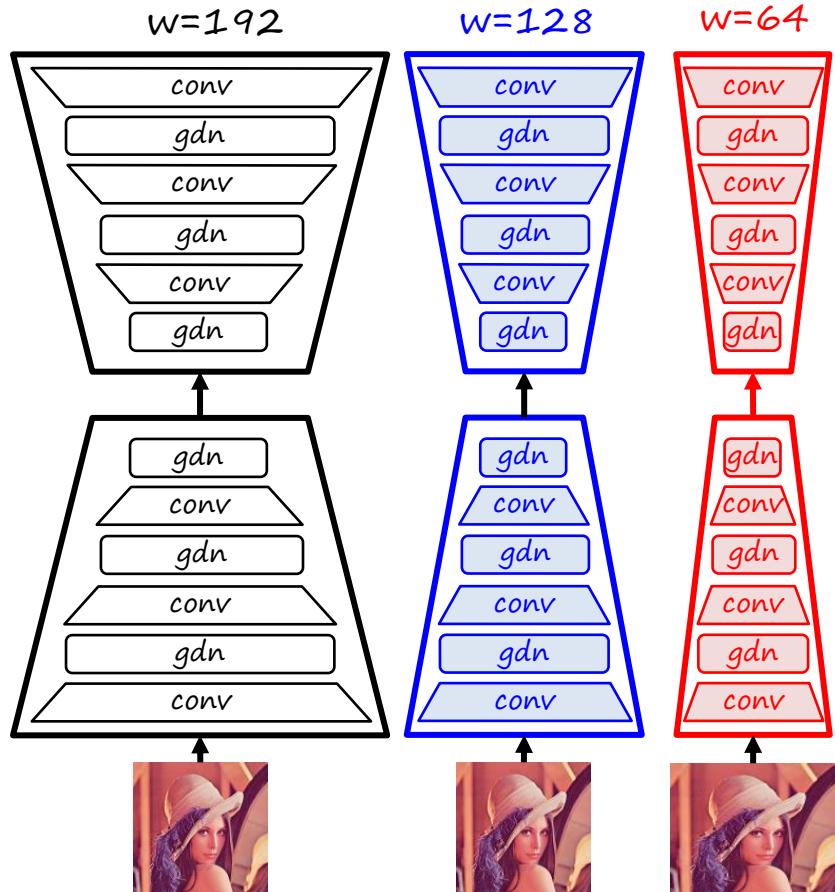
Feature modulation [MAE, cAE]



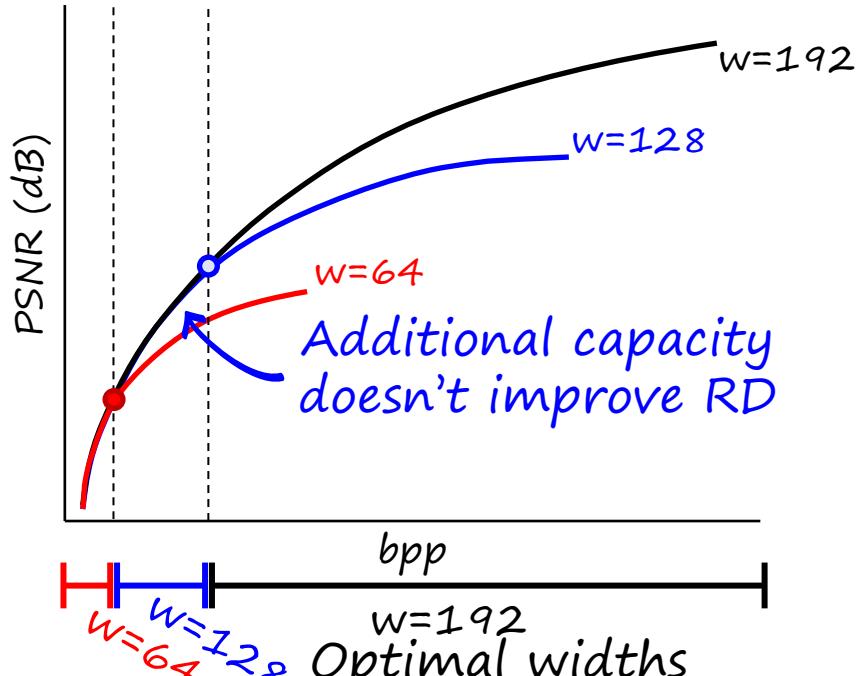
- Low memory ✗
- Low computation ✗
- Low latency ✗

Model capacity and rate-distortion

w =filters per layer



There is a minimal capacity
for every RD tradeoff

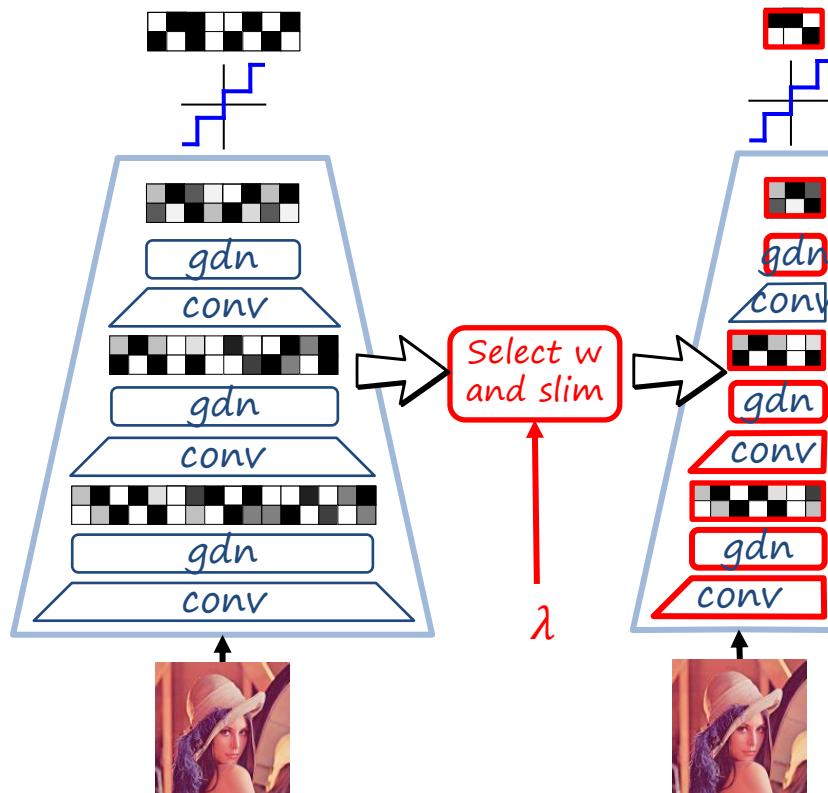


Lower w results in less
memory and computation!!

Slimmable compressive autoencoder

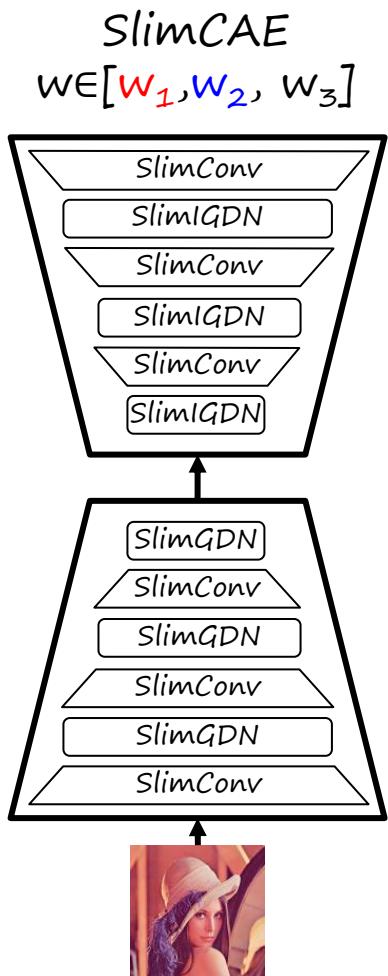
Approach: slim the network to the minimal capacity for a given λ

Slimming [SlimCAE]

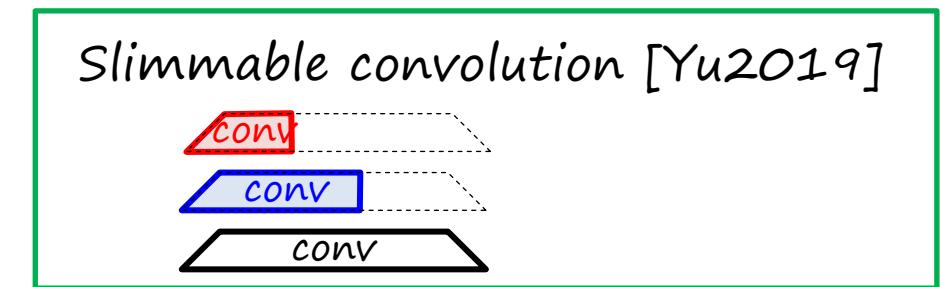
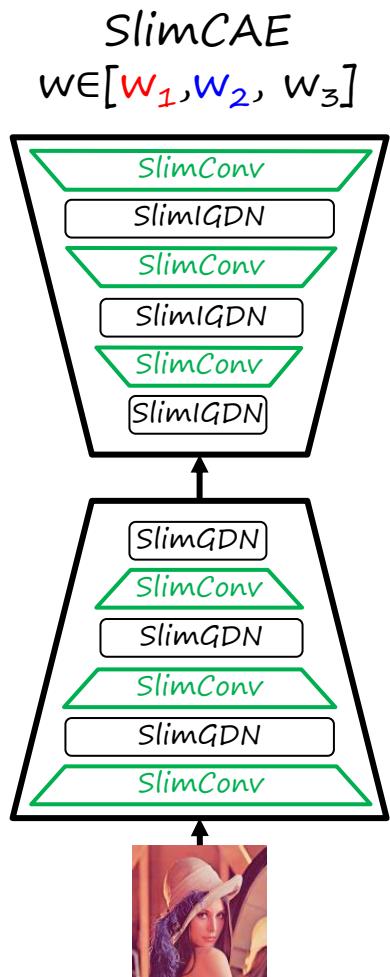


- Minimize rate ✓
 - Minimize distortion ✓
 - Variable rate ✓
 - Lower memory ✓
 - Lower computation ✓
 - Lower latency ✓
- (for low-mid rates)

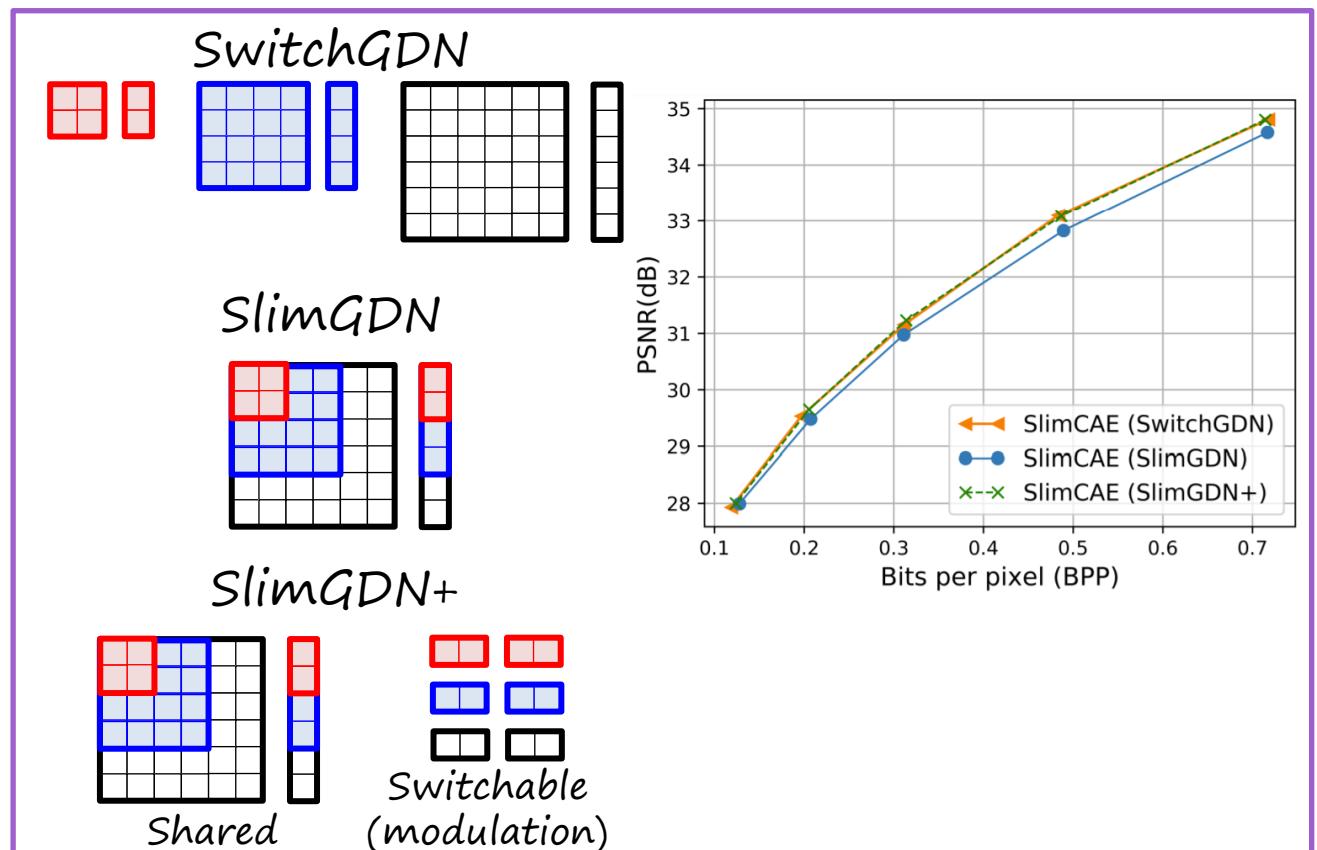
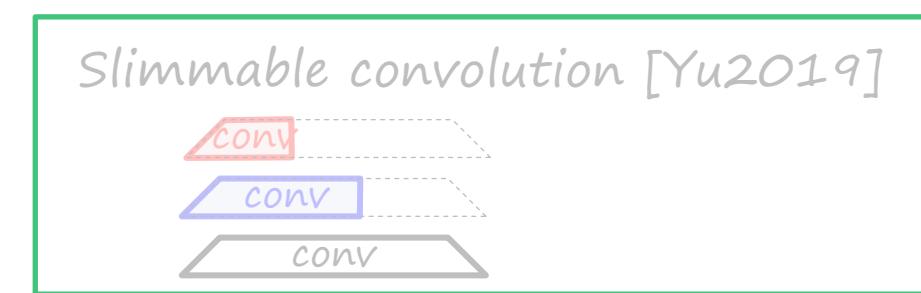
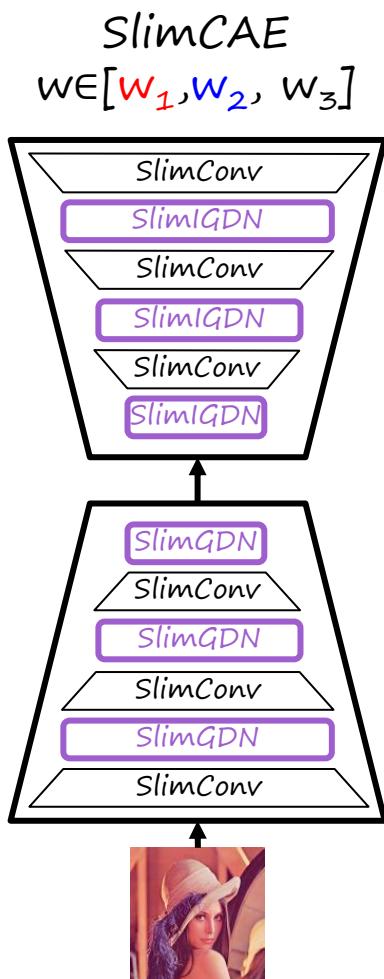
Slimmable layers in SlimCAE



Slimmable layers in SlimCAE

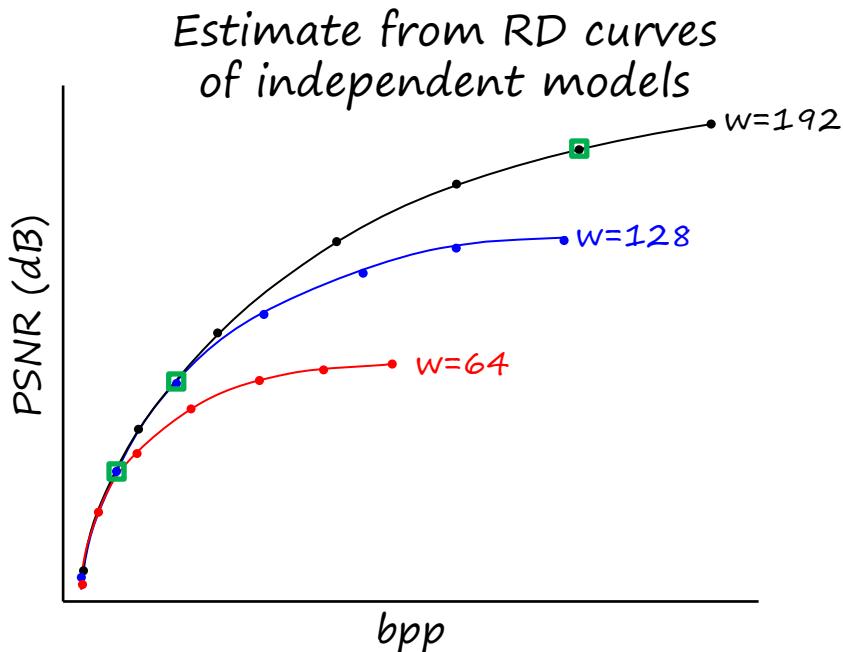


Slimmable layers in SlimCAE



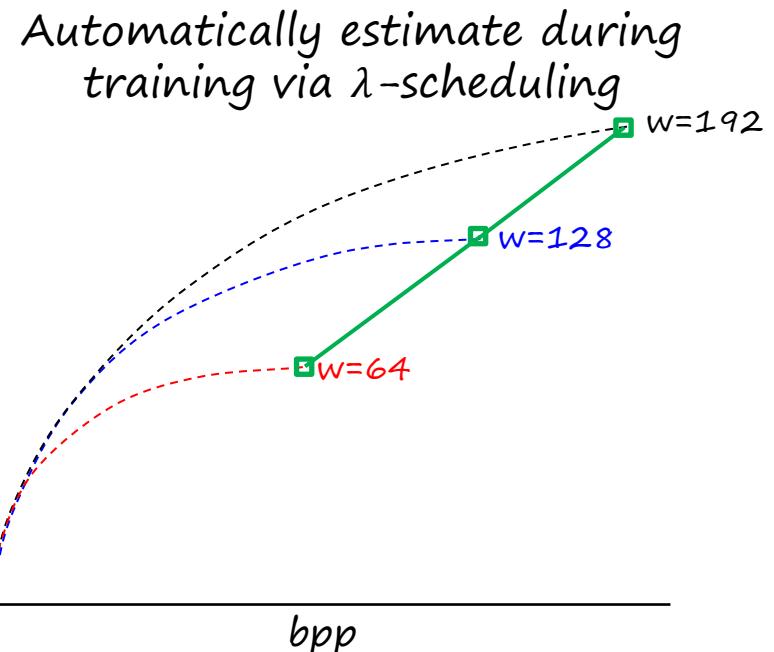
Training SlimCAE

Problem: we need the optimal λ s to train the SlimCAE



1. Train several independent models for different w
2. Plot RD curves and find critical points
3. Estimate optimal λ s from trained models

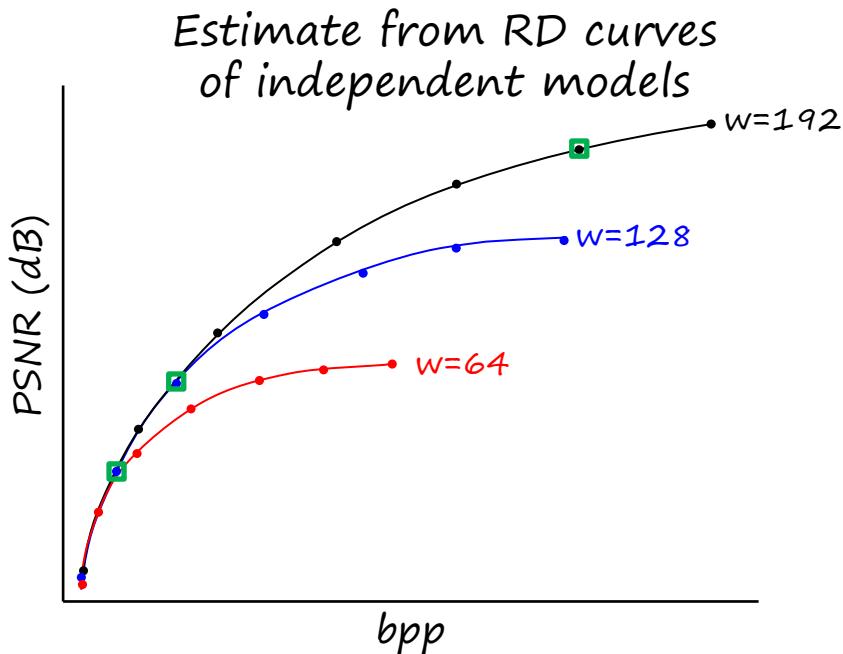
Problem: extremely expensive!



1. Train a SlimCAE with $\lambda_1 = \lambda_2 = \lambda_3$
2. While not converged do
 - Update λ s according to schedule
 - Optimize CAE

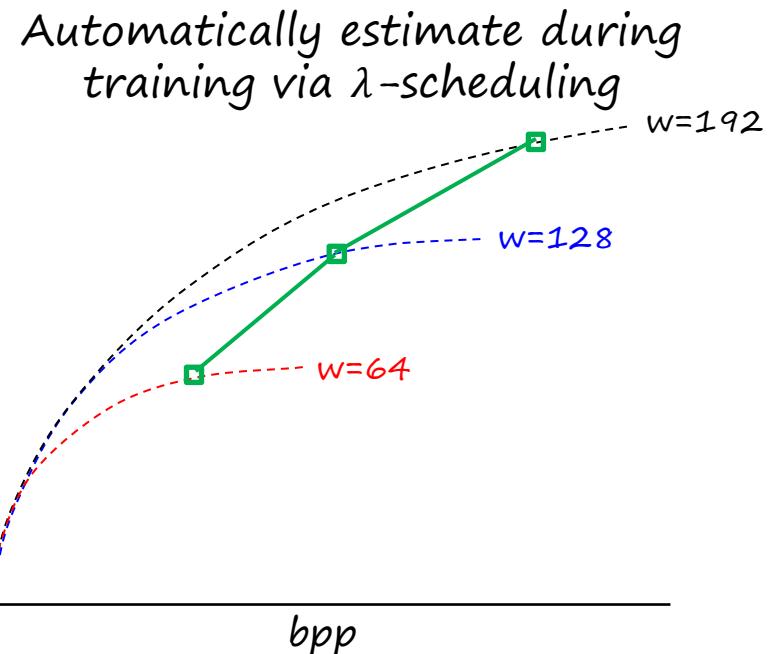
Training SlimCAE

Problem: we need the optimal λ s to train the SlimCAE



1. Train several independent models for different w
2. Plot RD curves and find critical points
3. Estimate optimal λ s from trained models

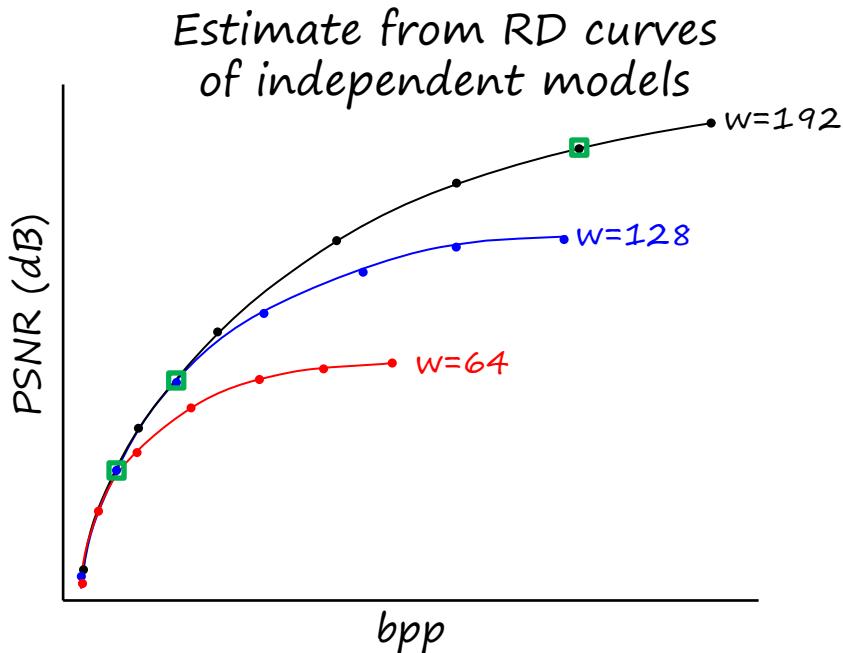
Problem: extremely expensive!



1. Train a SlimCAE with $\lambda_1 = \lambda_2 = \lambda_3$
2. While not converged do
 - Update λ s according to schedule
 - Optimize CAE

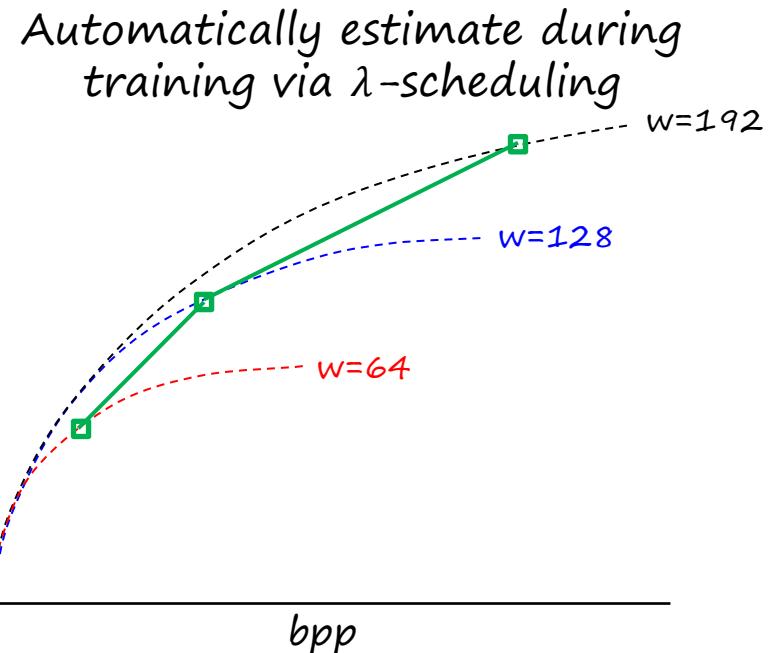
Training SlimCAE

Problem: we need the optimal λ s to train the SlimCAE



1. Train several independent models for different w
2. Plot RD curves and find critical points
3. Estimate optimal λ s from trained models

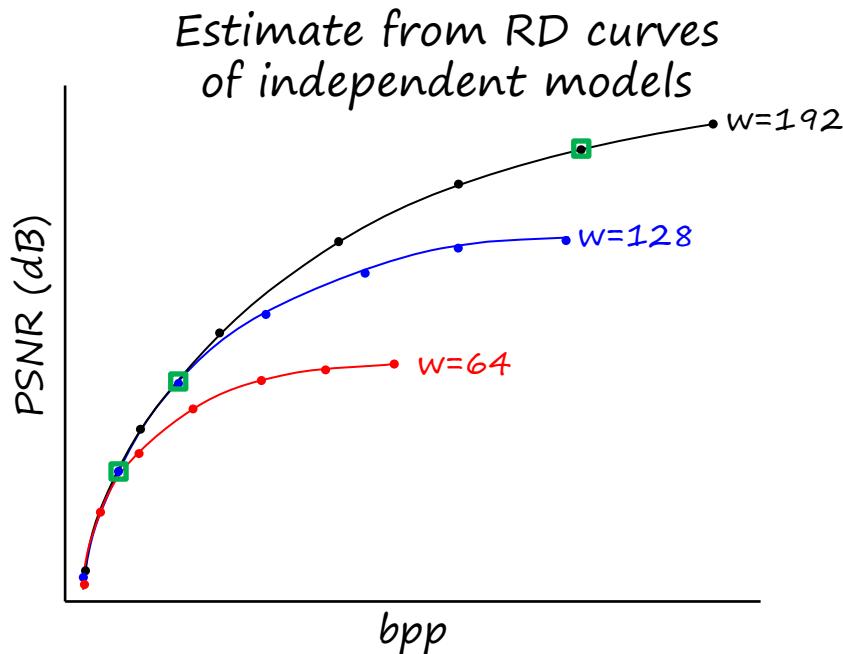
Problem: extremely expensive!



1. Train a SlimCAE with $\lambda_1 = \lambda_2 = \lambda_3$
2. While not converged do
 - Update λ s according to schedule
 - Optimize CAE

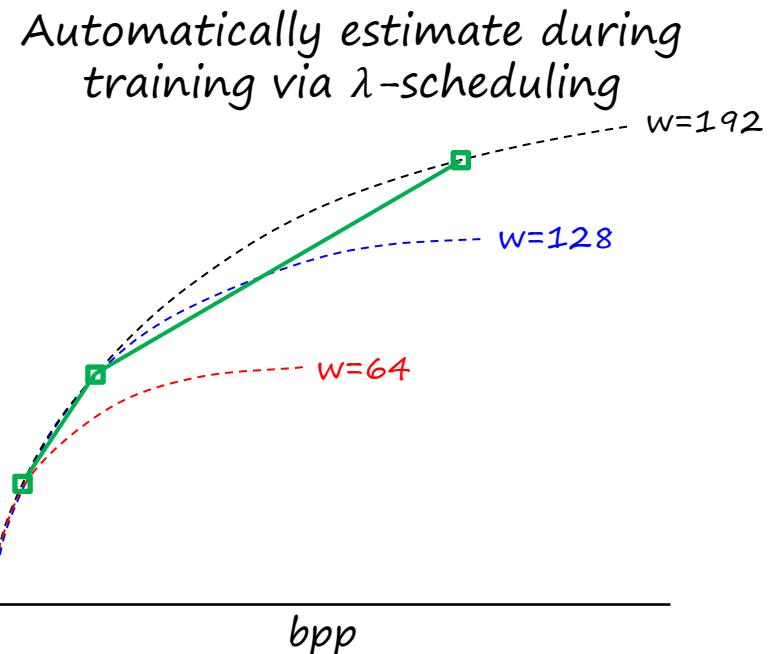
Training SlimCAE

Problem: we need the optimal λ s to train the SlimCAE



1. Train several independent models for different w
2. Plot RD curves and find critical points
3. Estimate optimal λ s from trained models

Problem: extremely expensive!



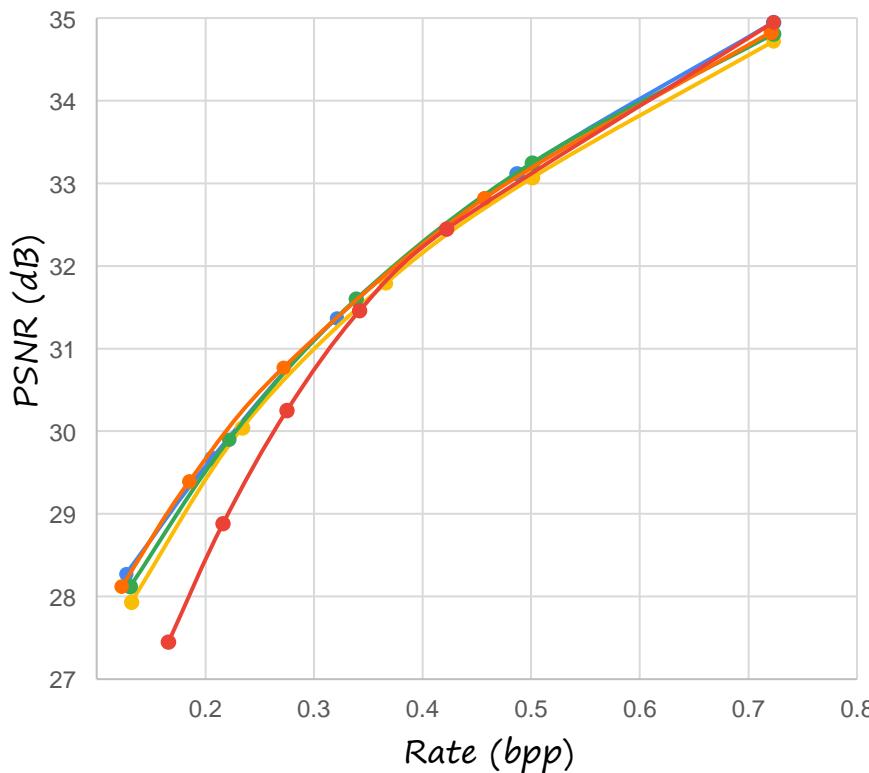
1. Train a SlimCAE with $\lambda_1 = \lambda_2 = \lambda_3$
2. While not converged do
 - Update λ s according to schedule
 - Optimize CAE

Directly train one model!

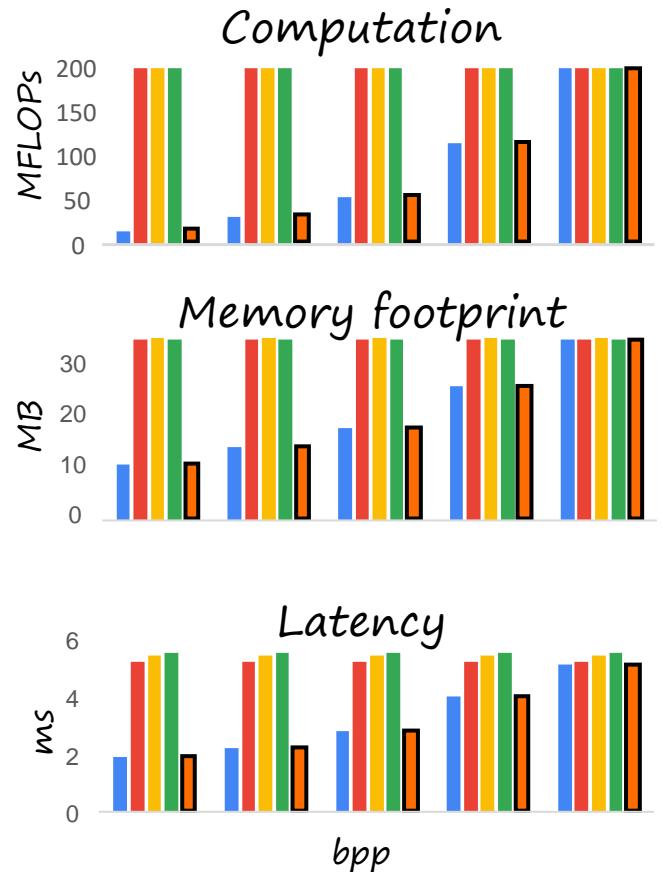
Performance comparison

Independent CAEs
(each with minimal capacity) Scaling [Theis2017] MAE [Yang2020] cAE [Choi2019] **SlimCAE (ours)**

Rate-distortion



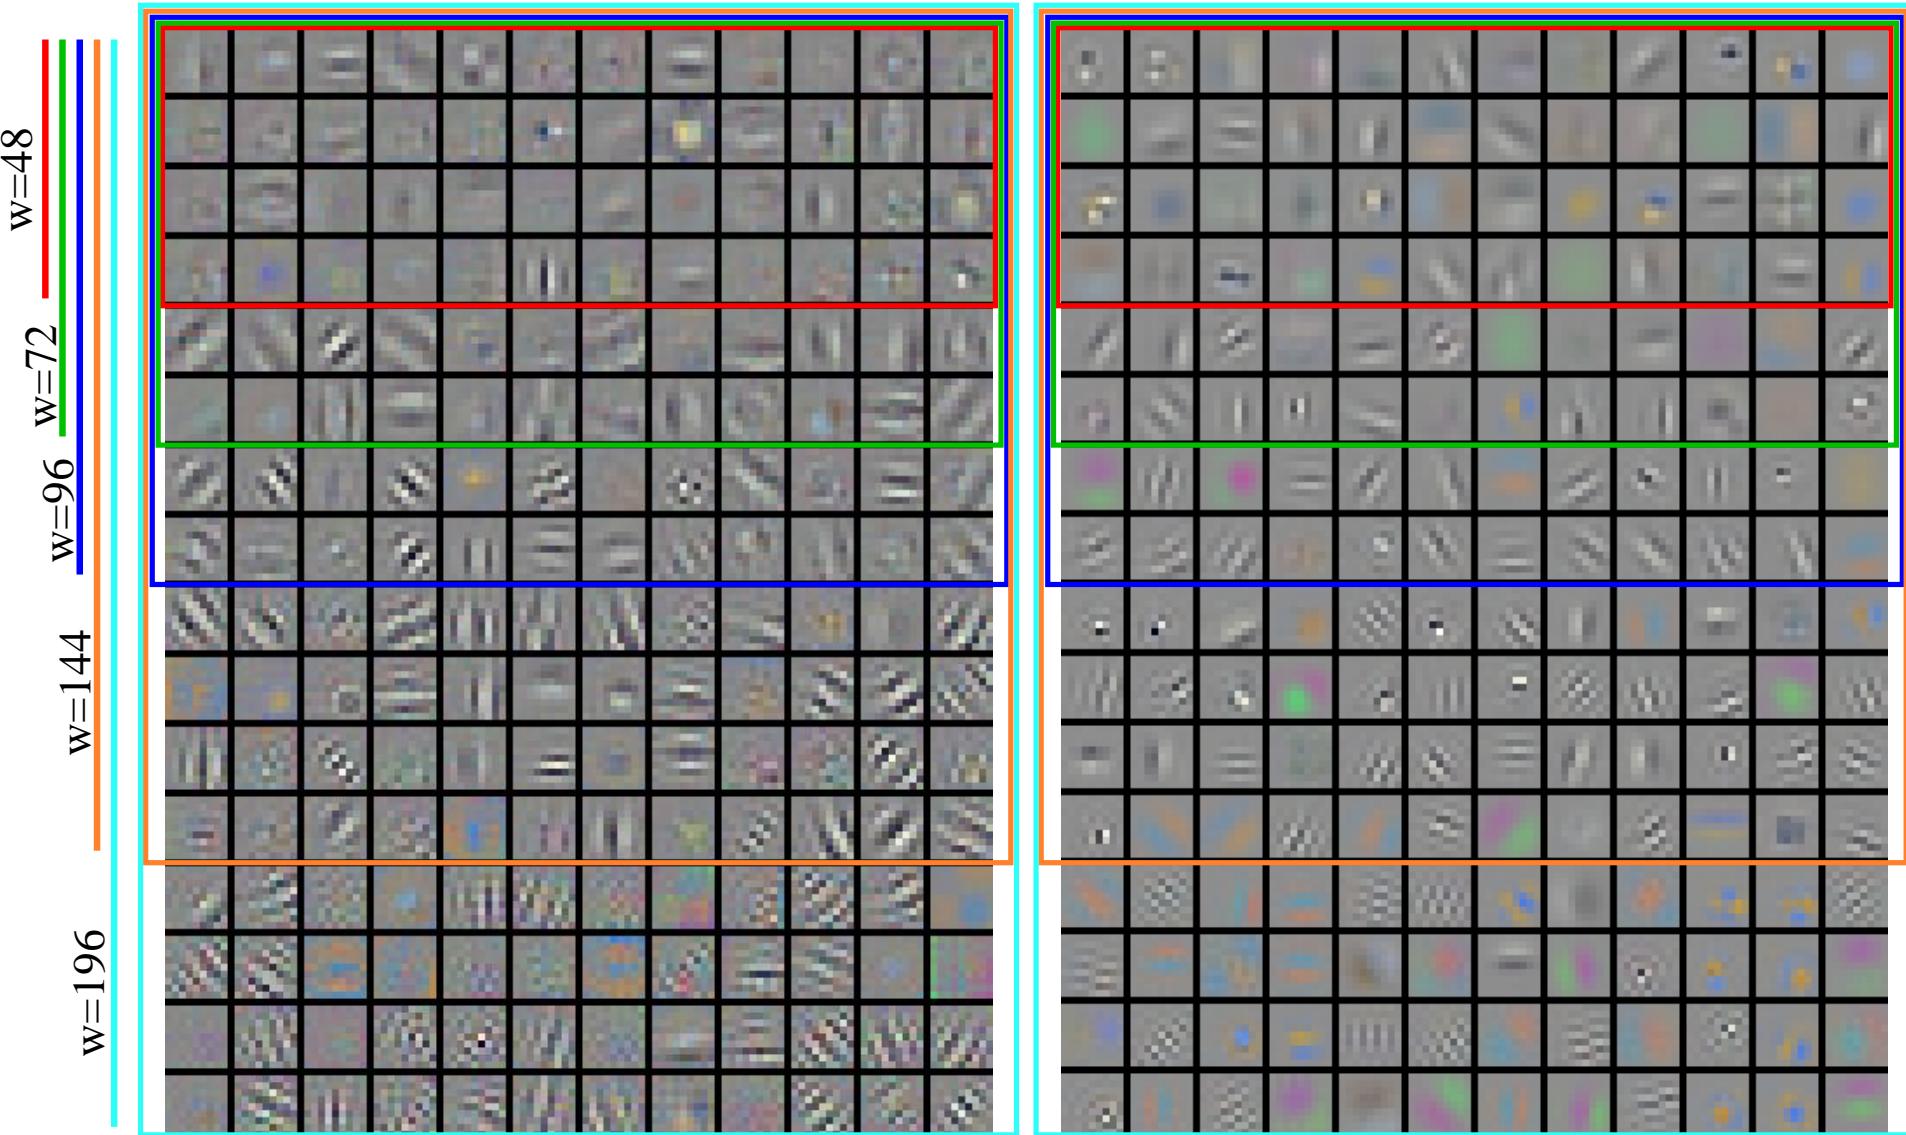
Encoder



Visualizing some parameters

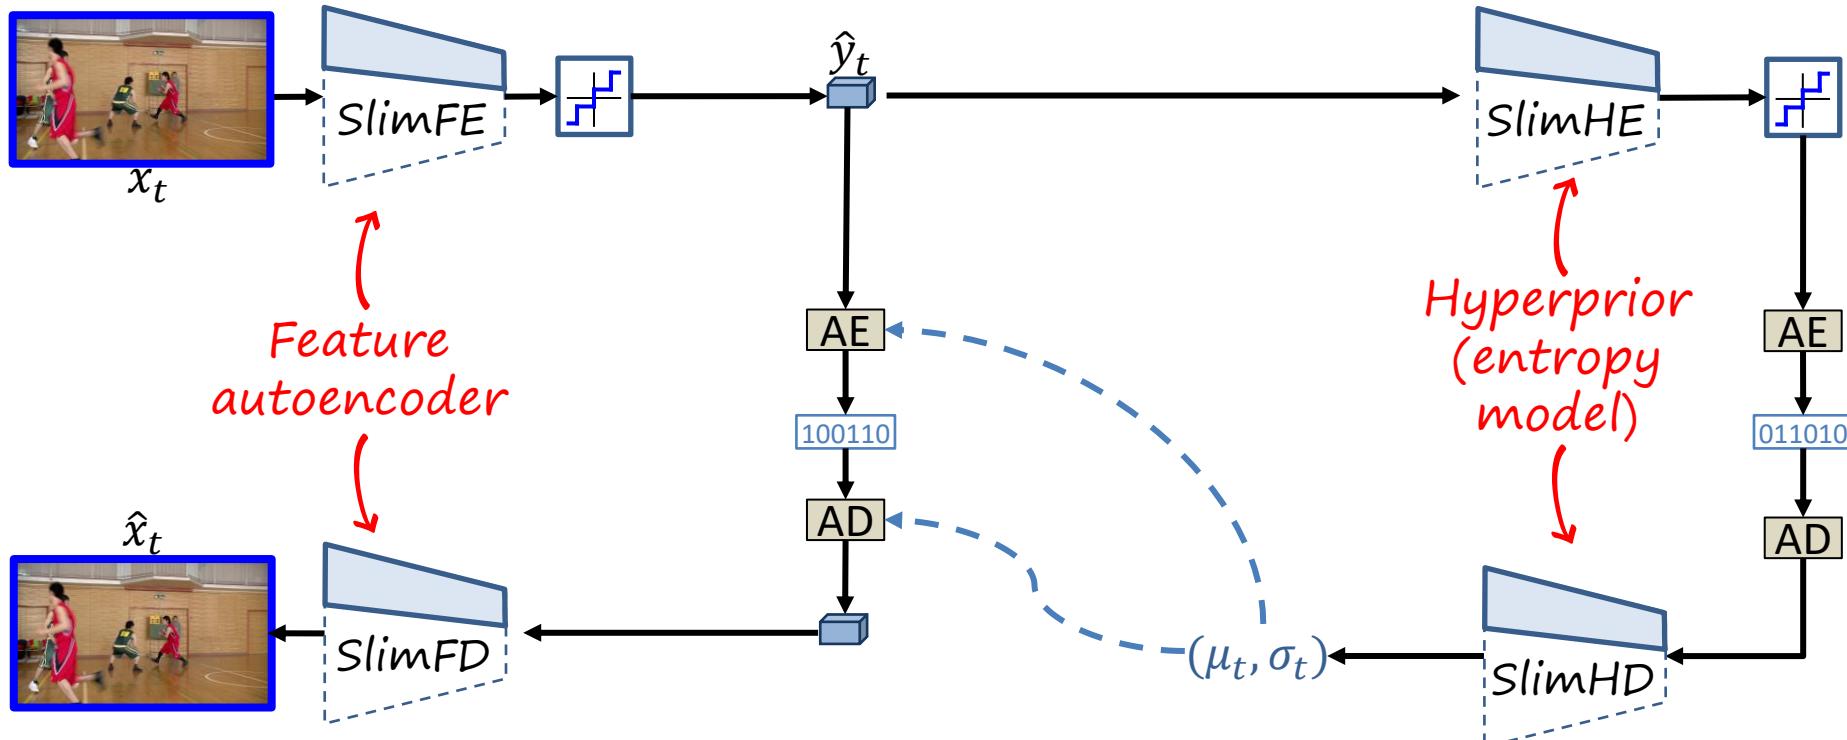
Encoder (first conv layer)

Decoder (last conv layer)



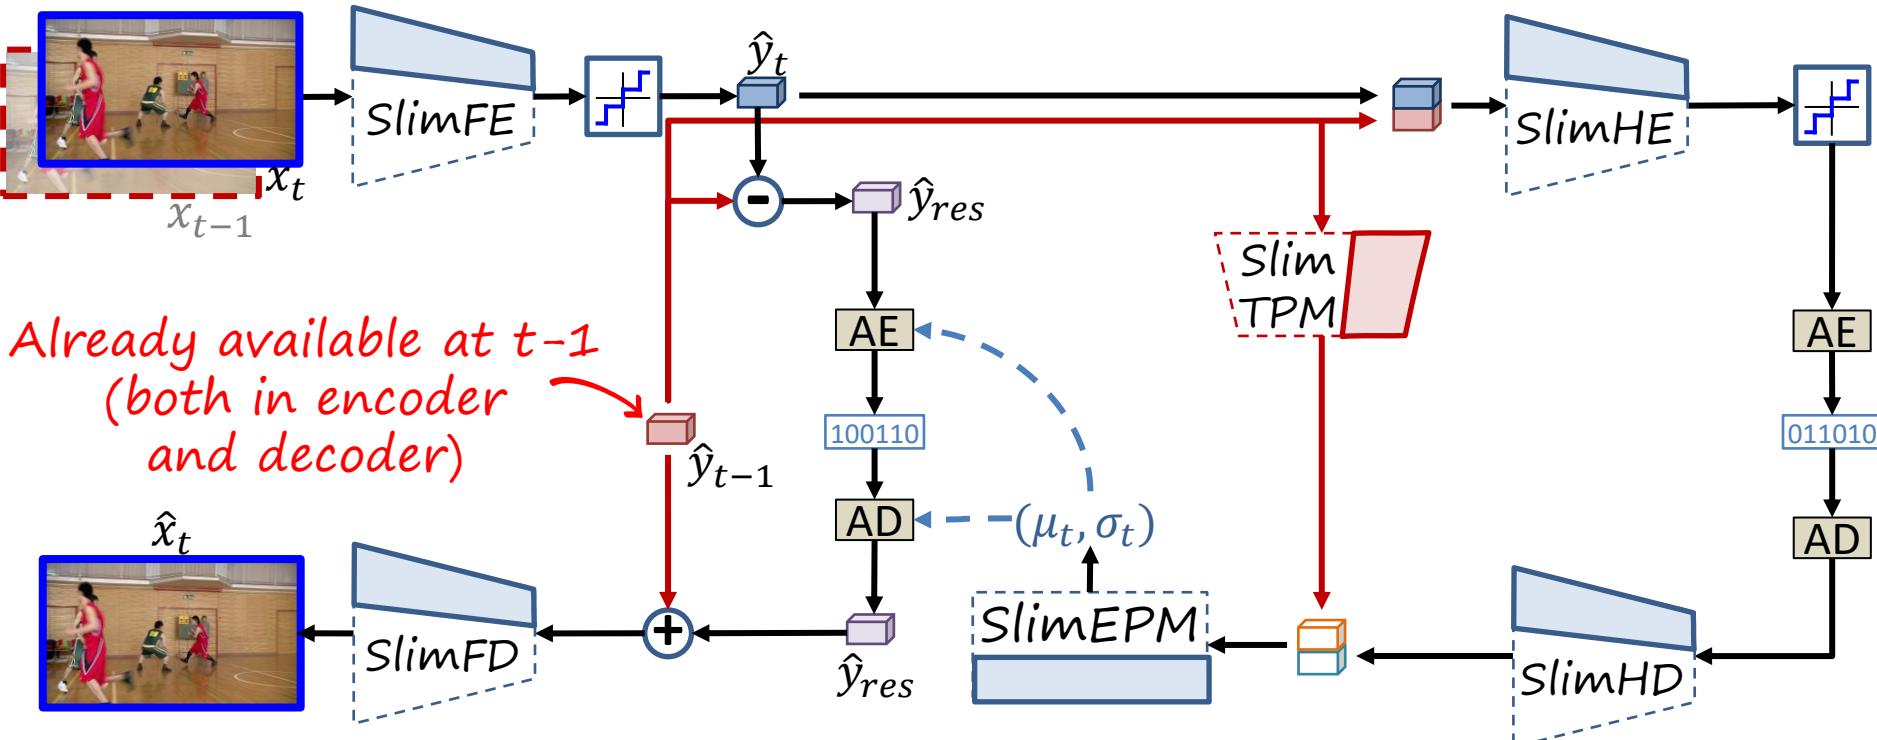
Slimmable video codec (SlimVC)

Extending SlimCAE to video



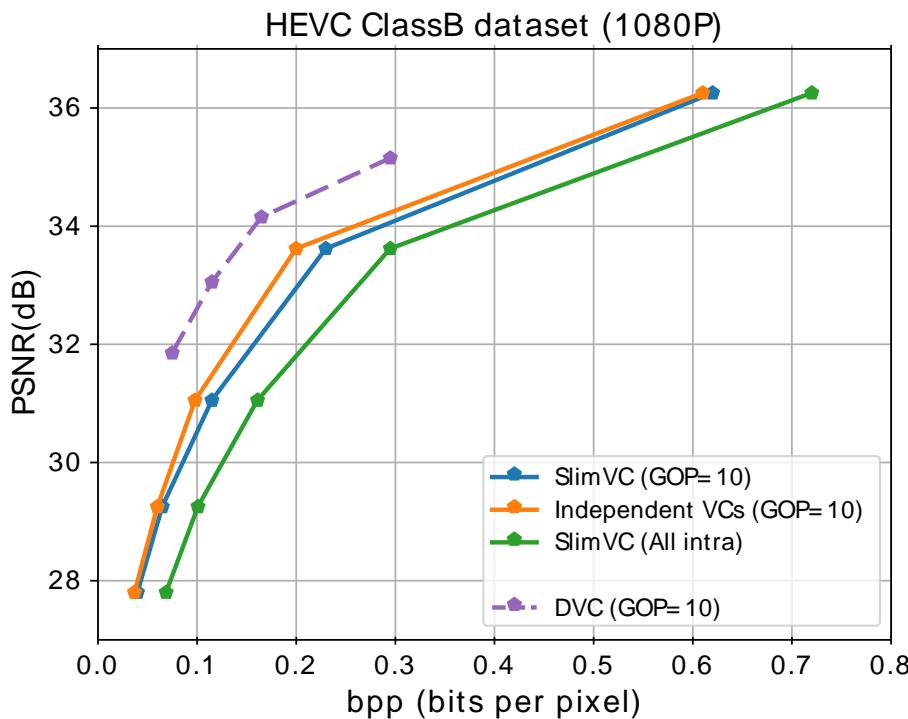
Slimmable video codec (SlimVC)

Extending SlimCAE to video

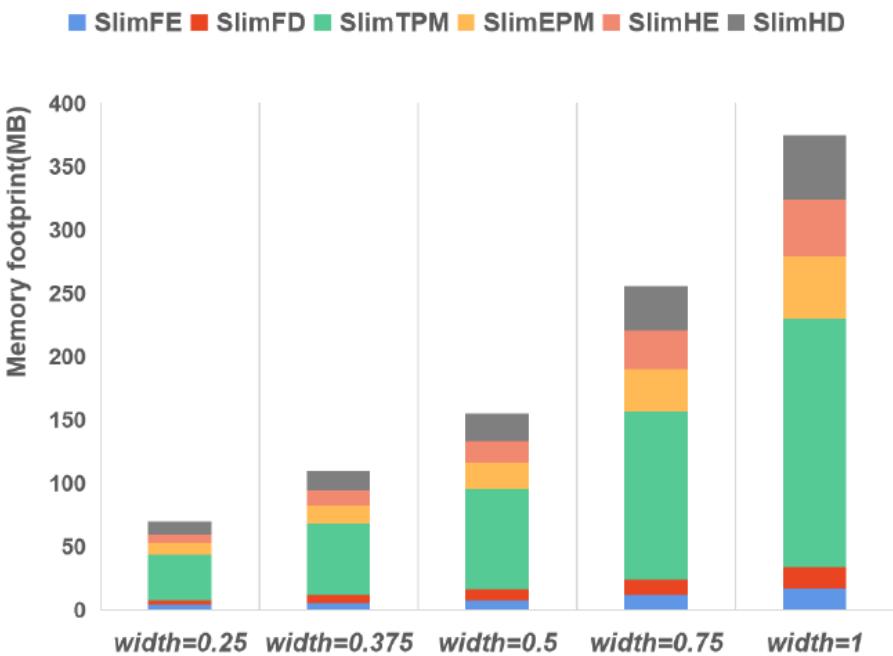


Slimmable video codec (SlimVC)

RD performance

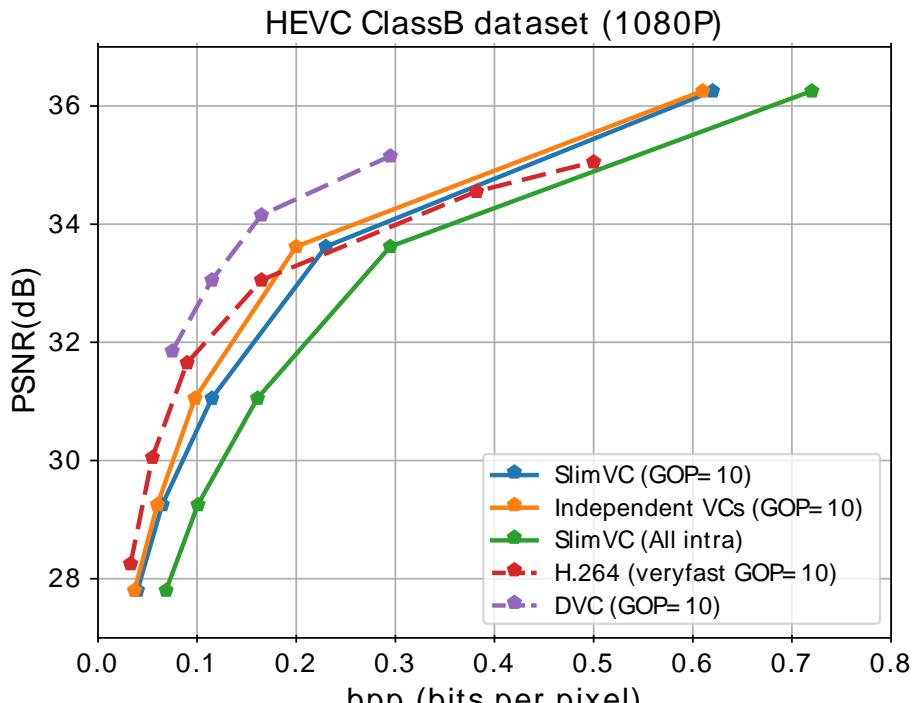


Memory footprint

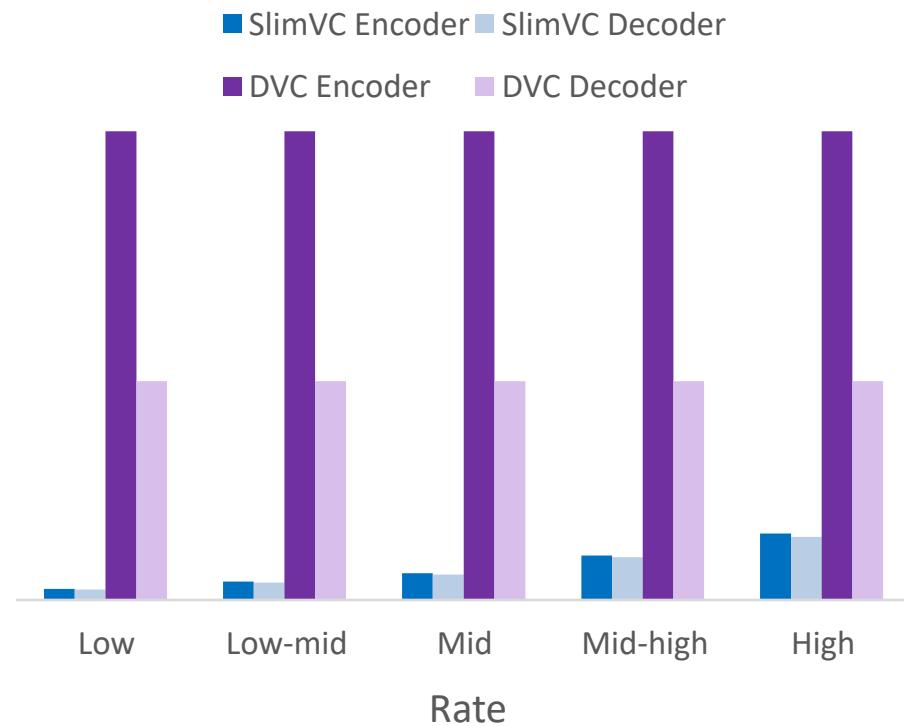


Slimmable video codec (SlimVC)

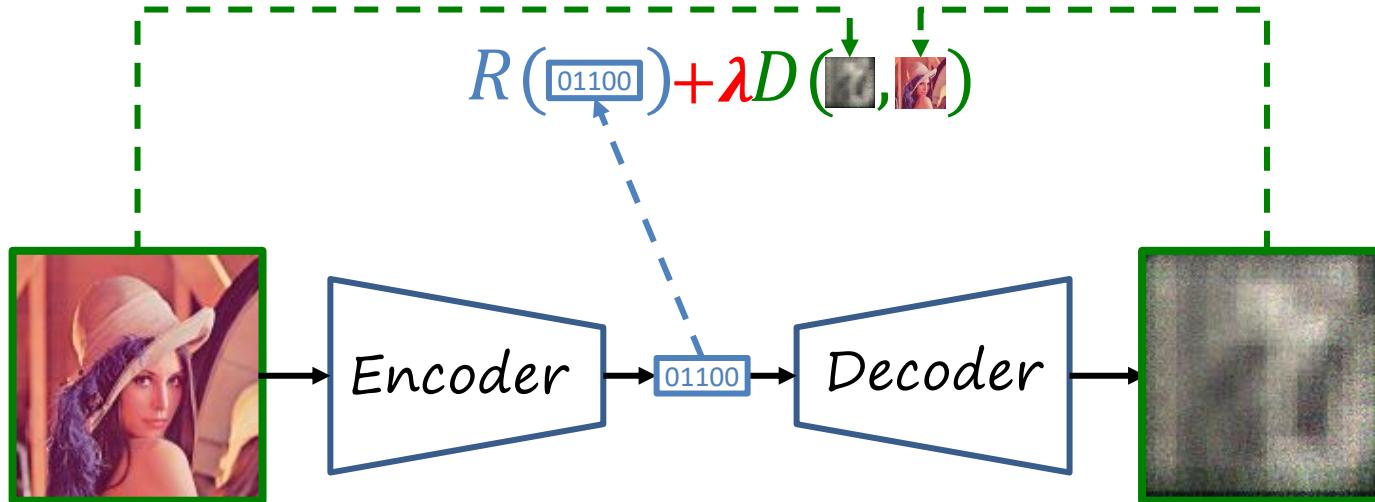
RD performance



Computational cost (GFLOPs)



Is neural image compression practical?



Limitations

- λ is fixed
- Heavy encoders/decoders

Practical neural image compression?

- Minimize rate ✓
- Minimize distortion ✓

- Variable rate ✗
- Low memory ✗
- Low computation ✗
- Low latency ✗

MAE
[SPL2020]
SlimCAE
[CVPR2021]

Other practical considerations

- Domain-specific codecs (e.g. videoconference, screencast)
- Back./forw. compatibility (with legacy encoders/decoders)

DANICE
[CLIC2021]

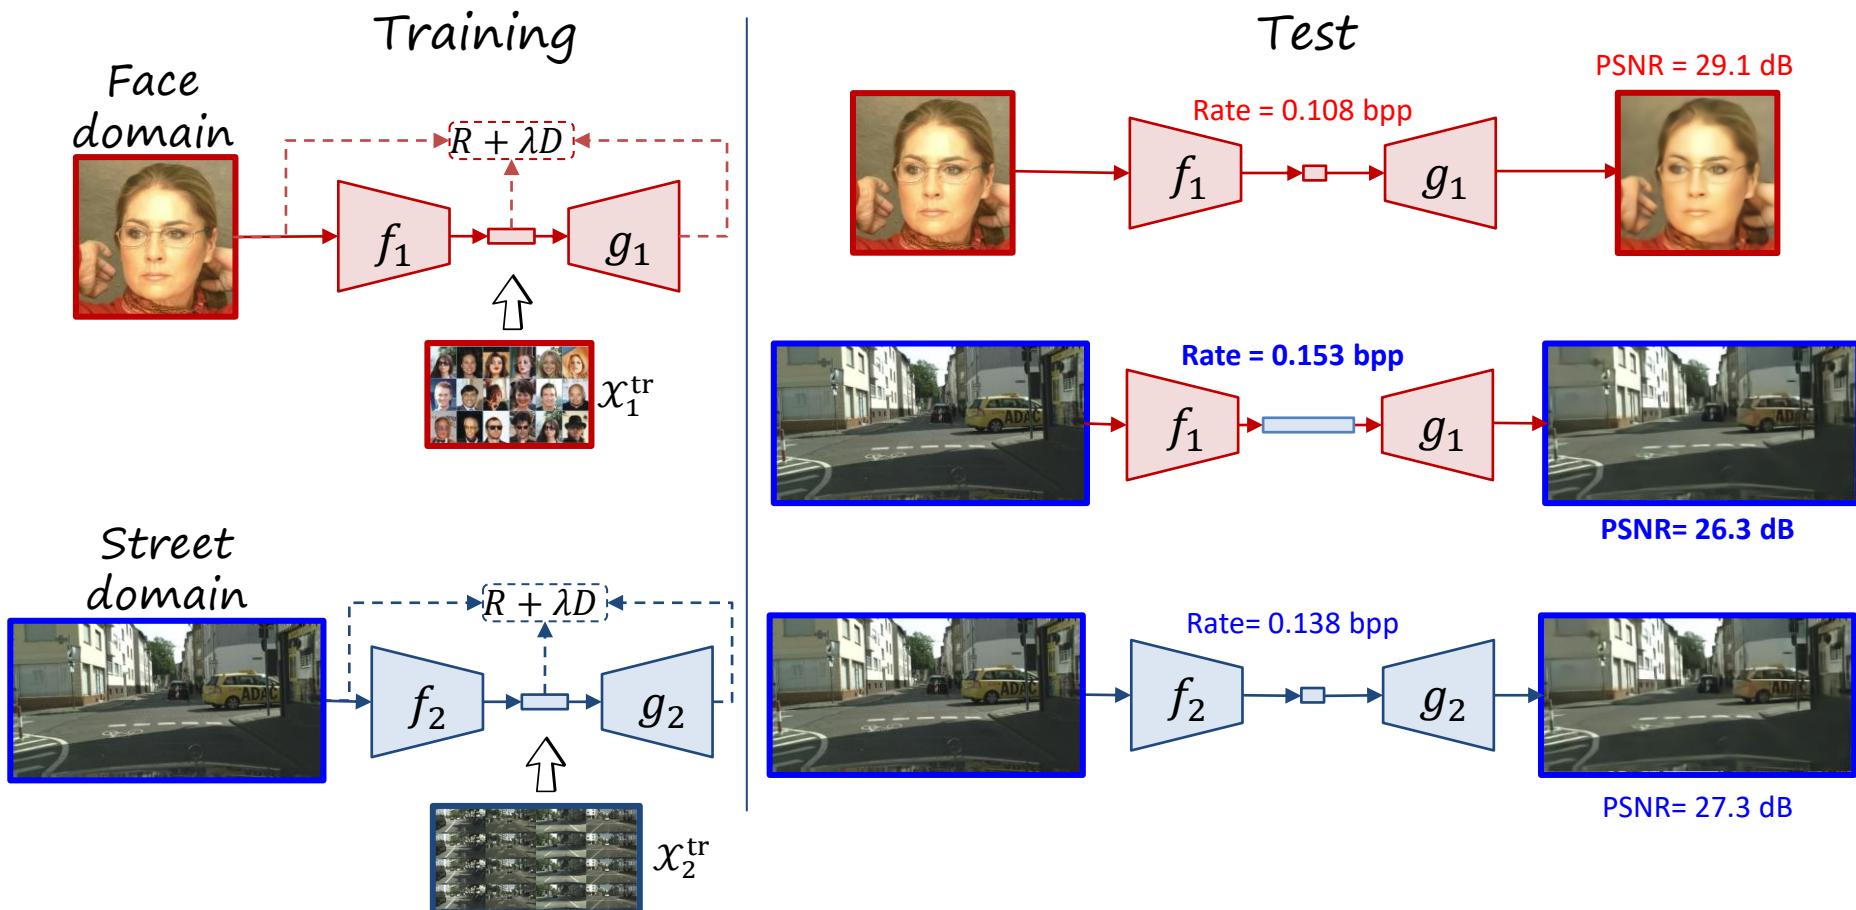
[SPL2020] [Variable Rate Deep Image Compression with Modulated Autoencoder](#), Signal Processing Letters 2020

[CVPR2021] [Slimmable compressive autoencoders for practical image compression](#), CVPR 2021

[CLIC2021] [DANICE: Domain adaptation without forgetting in neural image compression](#), CLIC 2021 at CVPR 2021

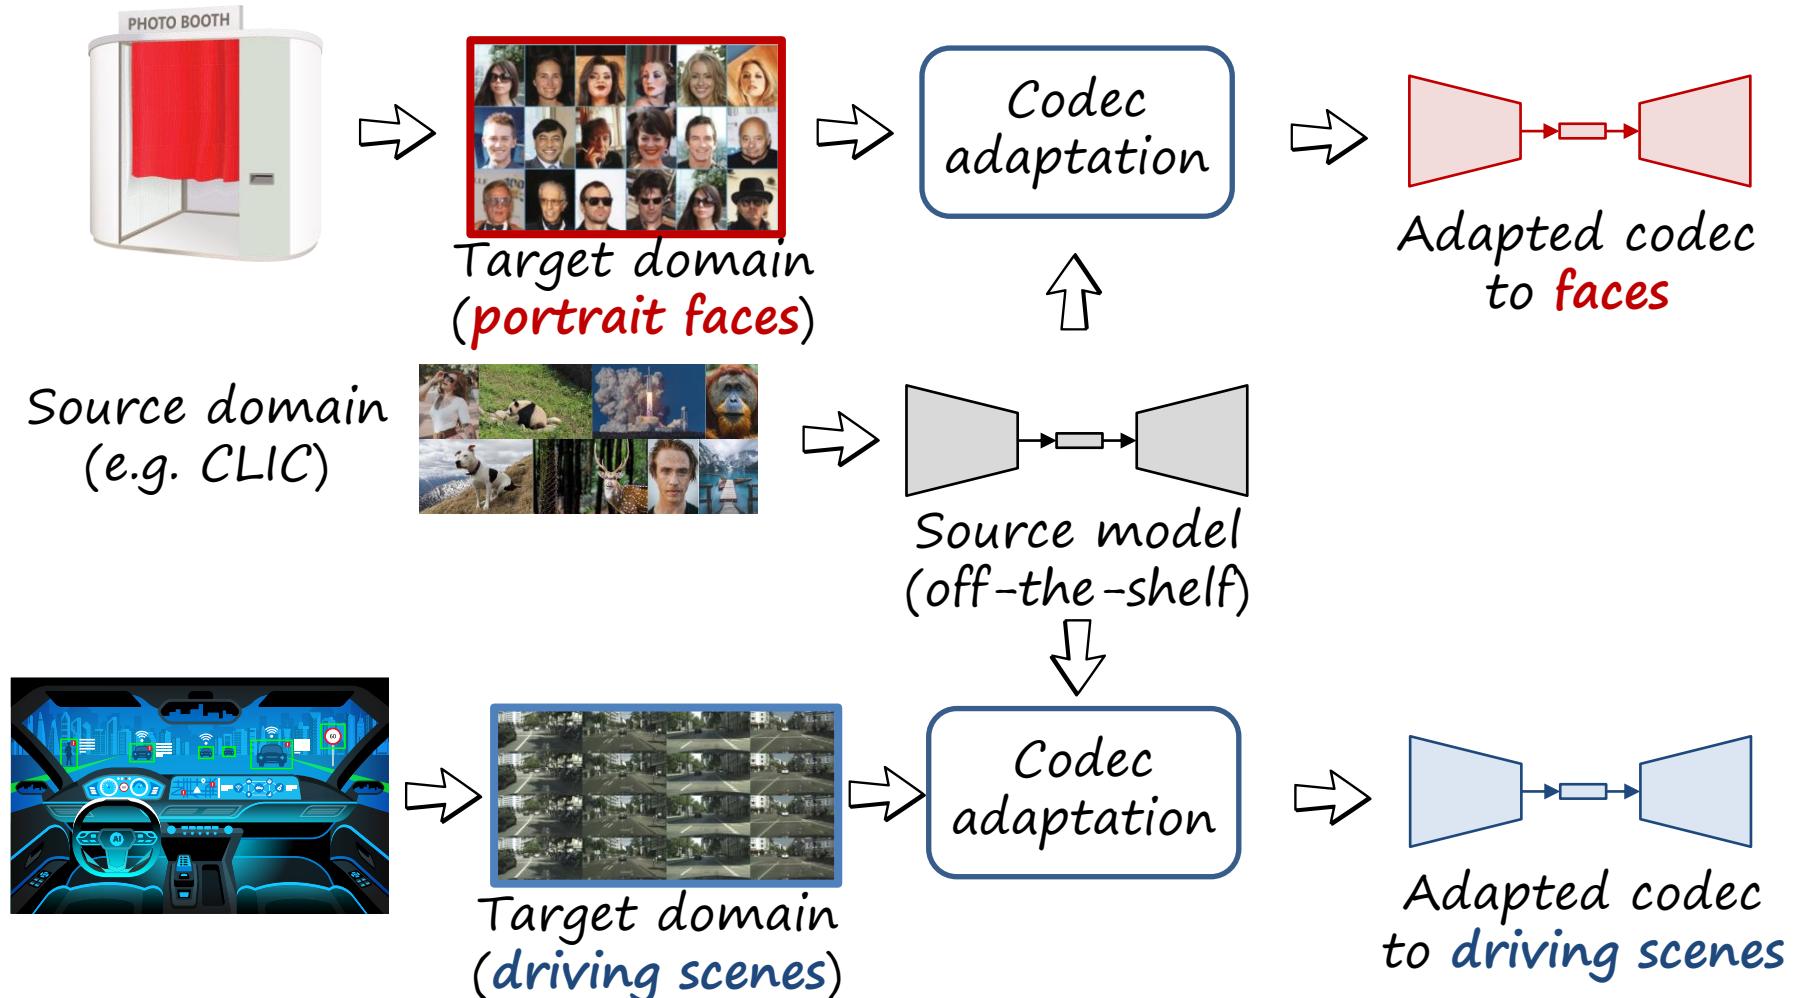
Rate-distortion optimality of learned codecs

Learned codecs are only optimal in the domain of the training data



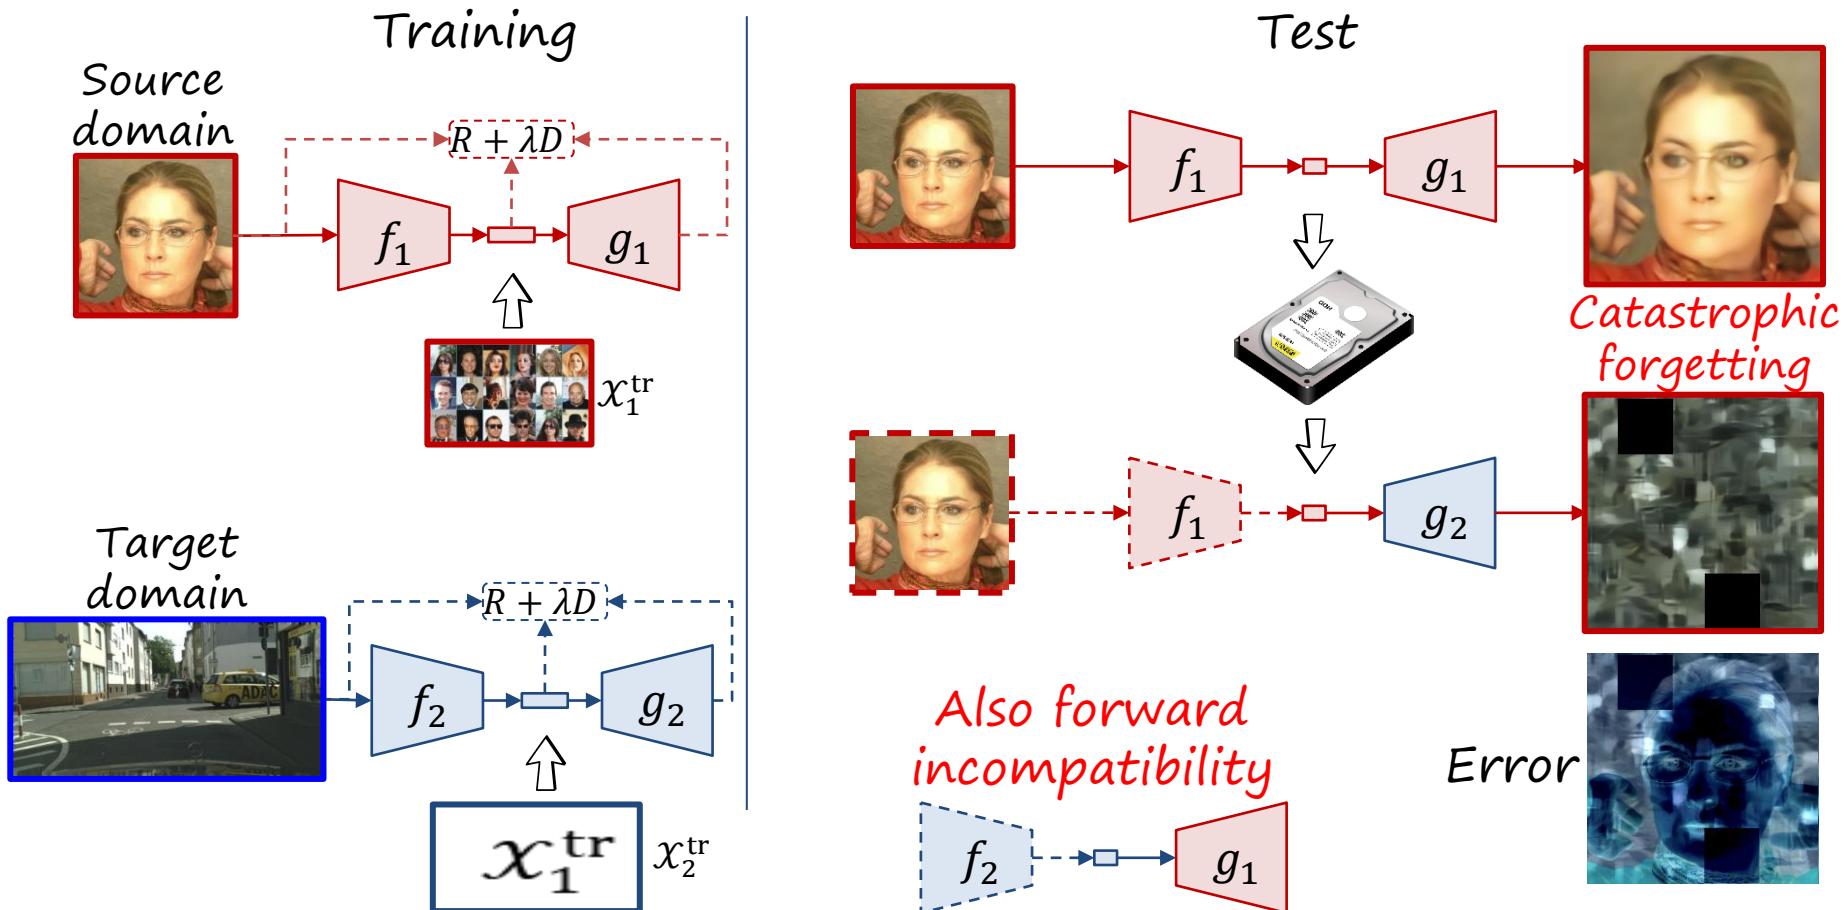
Domain Adaptation in Neural Image ComprEsson (DANICE)

Learned codecs can be customized with user content to specific domains
Problem: usually not enough custom data; training is expensive
Solution: transfer pre-trained codecs



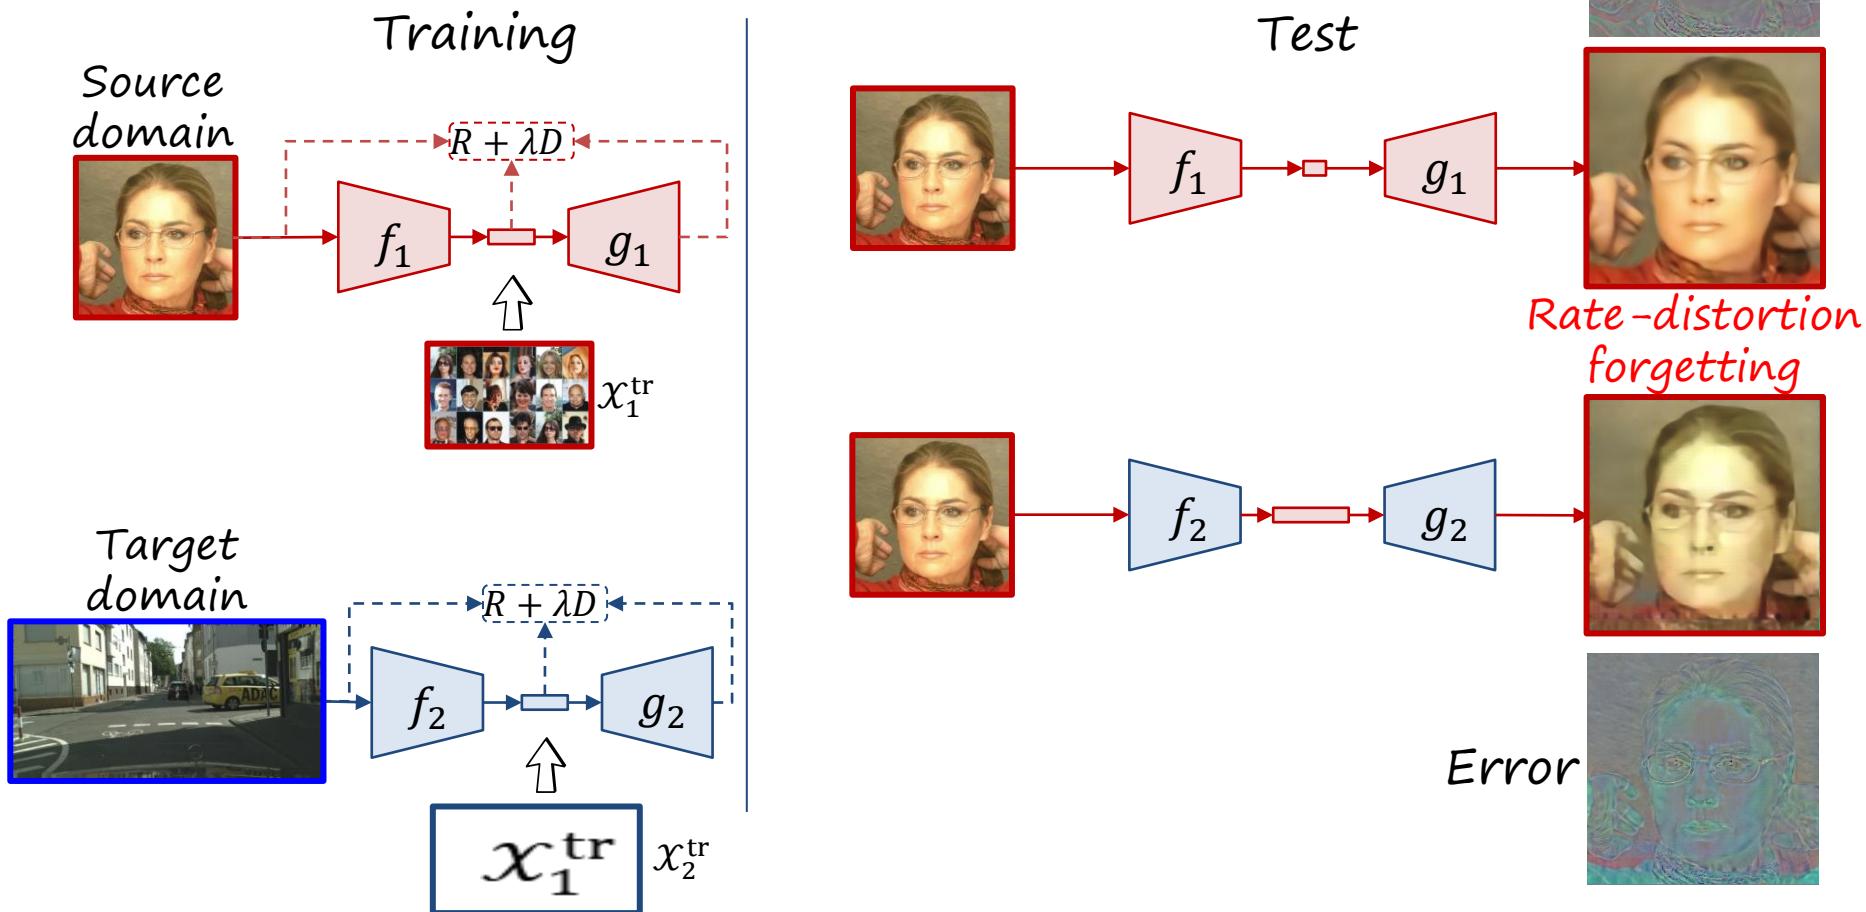
Backward incompatibility with legacy bitstreams: catastrophic forgetting

Misalignment between encoding-decoding latent spaces
(i.e. bitstream syntax incompatible)



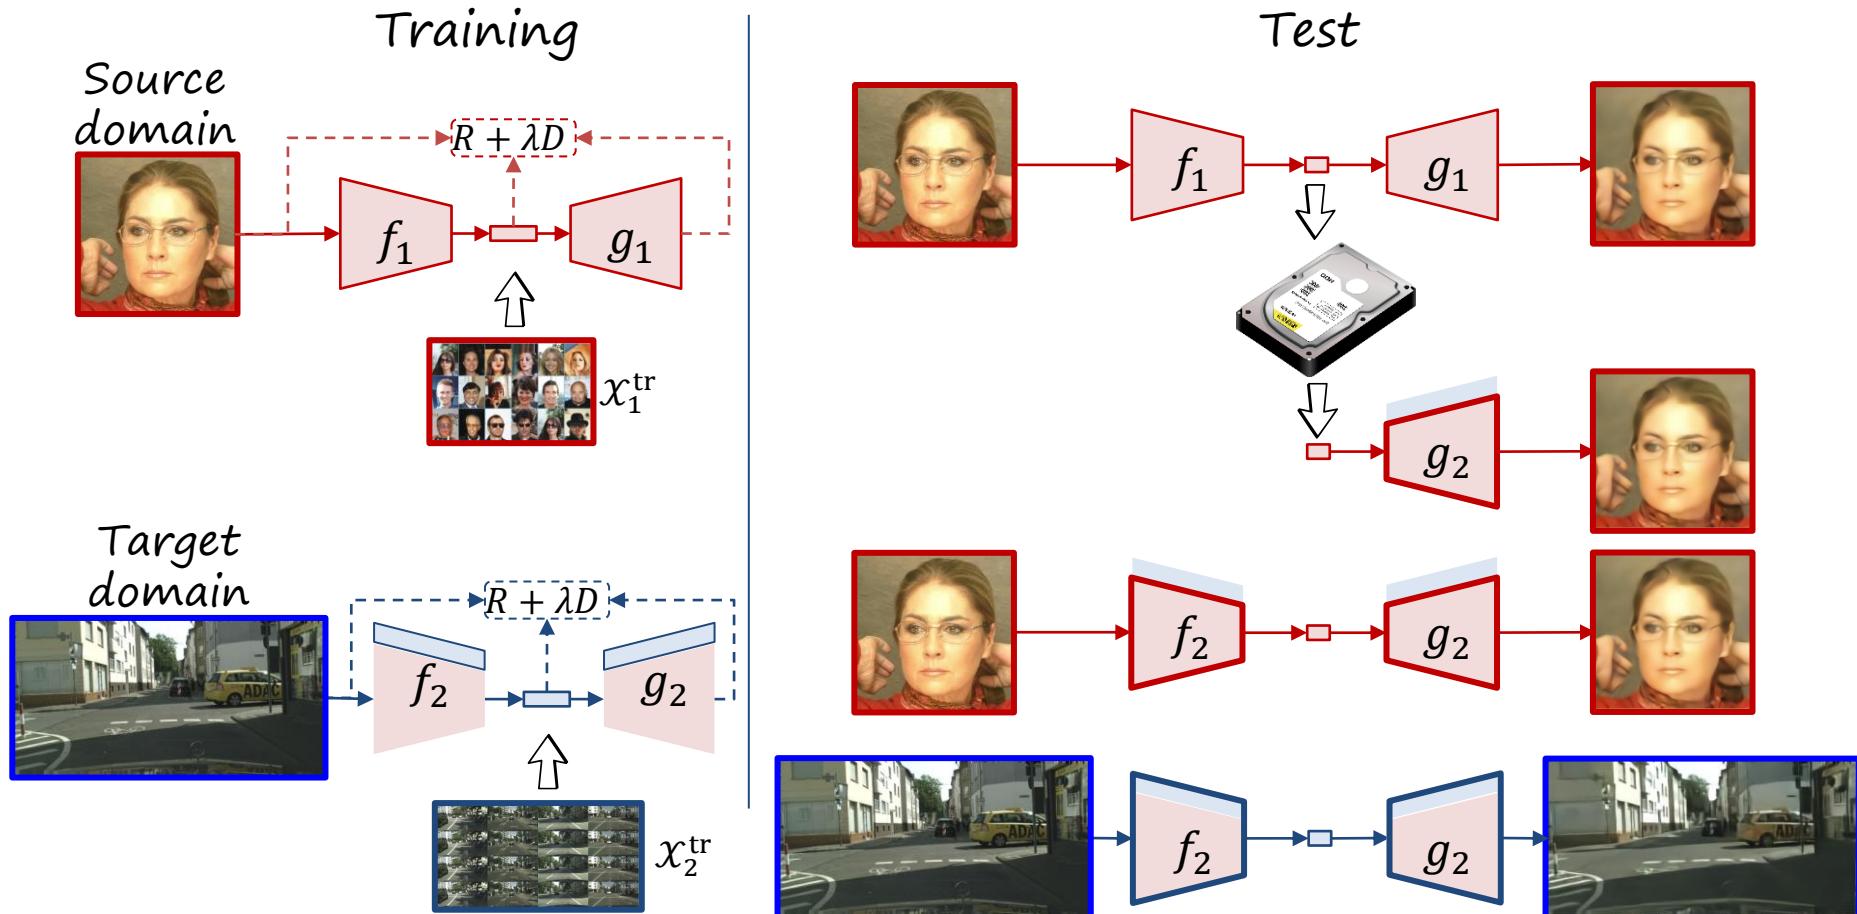
Rate-distortion forgetting

Encoding-decoding latent spaces aligned, but suboptimal
(i.e. bitstream syntax compatible, yet degraded)



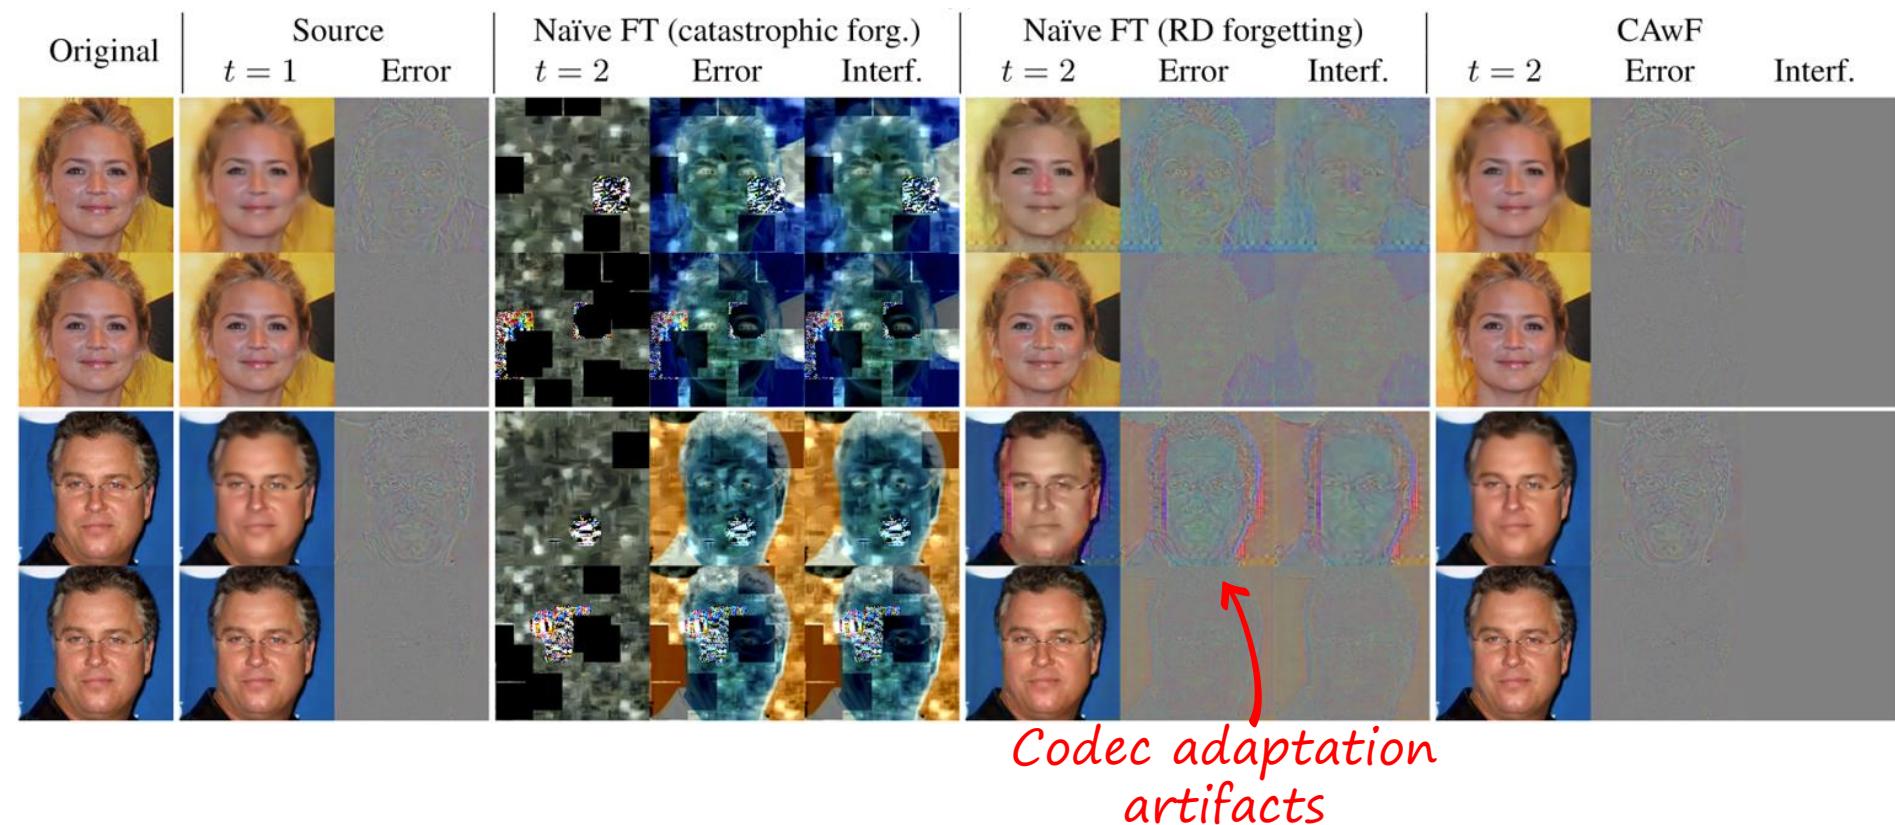
Codec adaptation without forgetting (CAwF)

Freeze source codec, and learn target codec as an enhancement layer
Drawback: adds additional parameters



Codec adaptation without forgetting (CAwF)

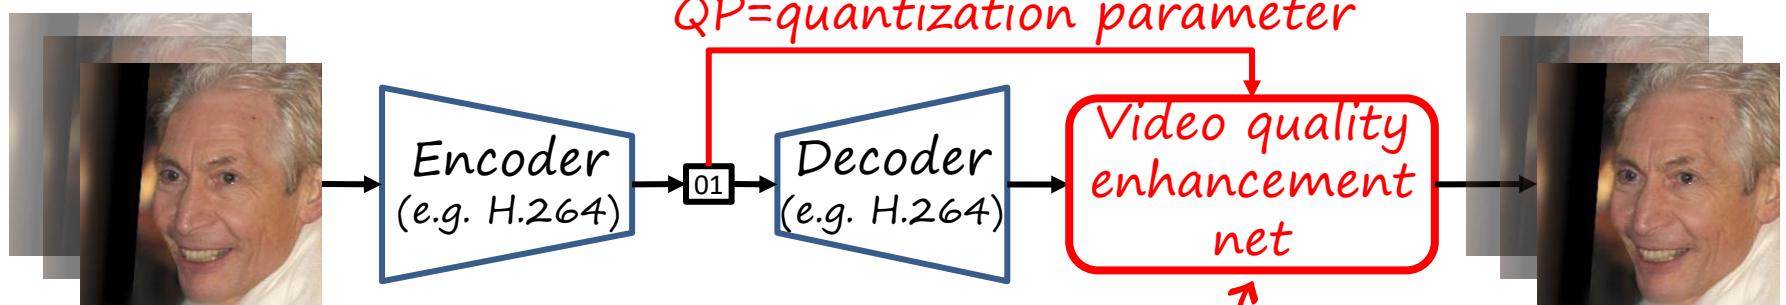
*CelebA→Cityscapes
(source domain)*



Outline

- Introduction: image coding
- Compression with neural networks
- Towards practical image compression
- **Briefly: other works**

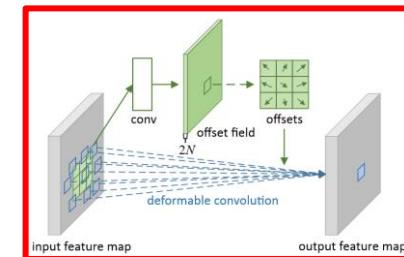
Video quality enhancement and artifact removal



Typical approach:

- Align several frames
- Aggregate the aligned information to alleviate noise/artifacts

Deformable convolution



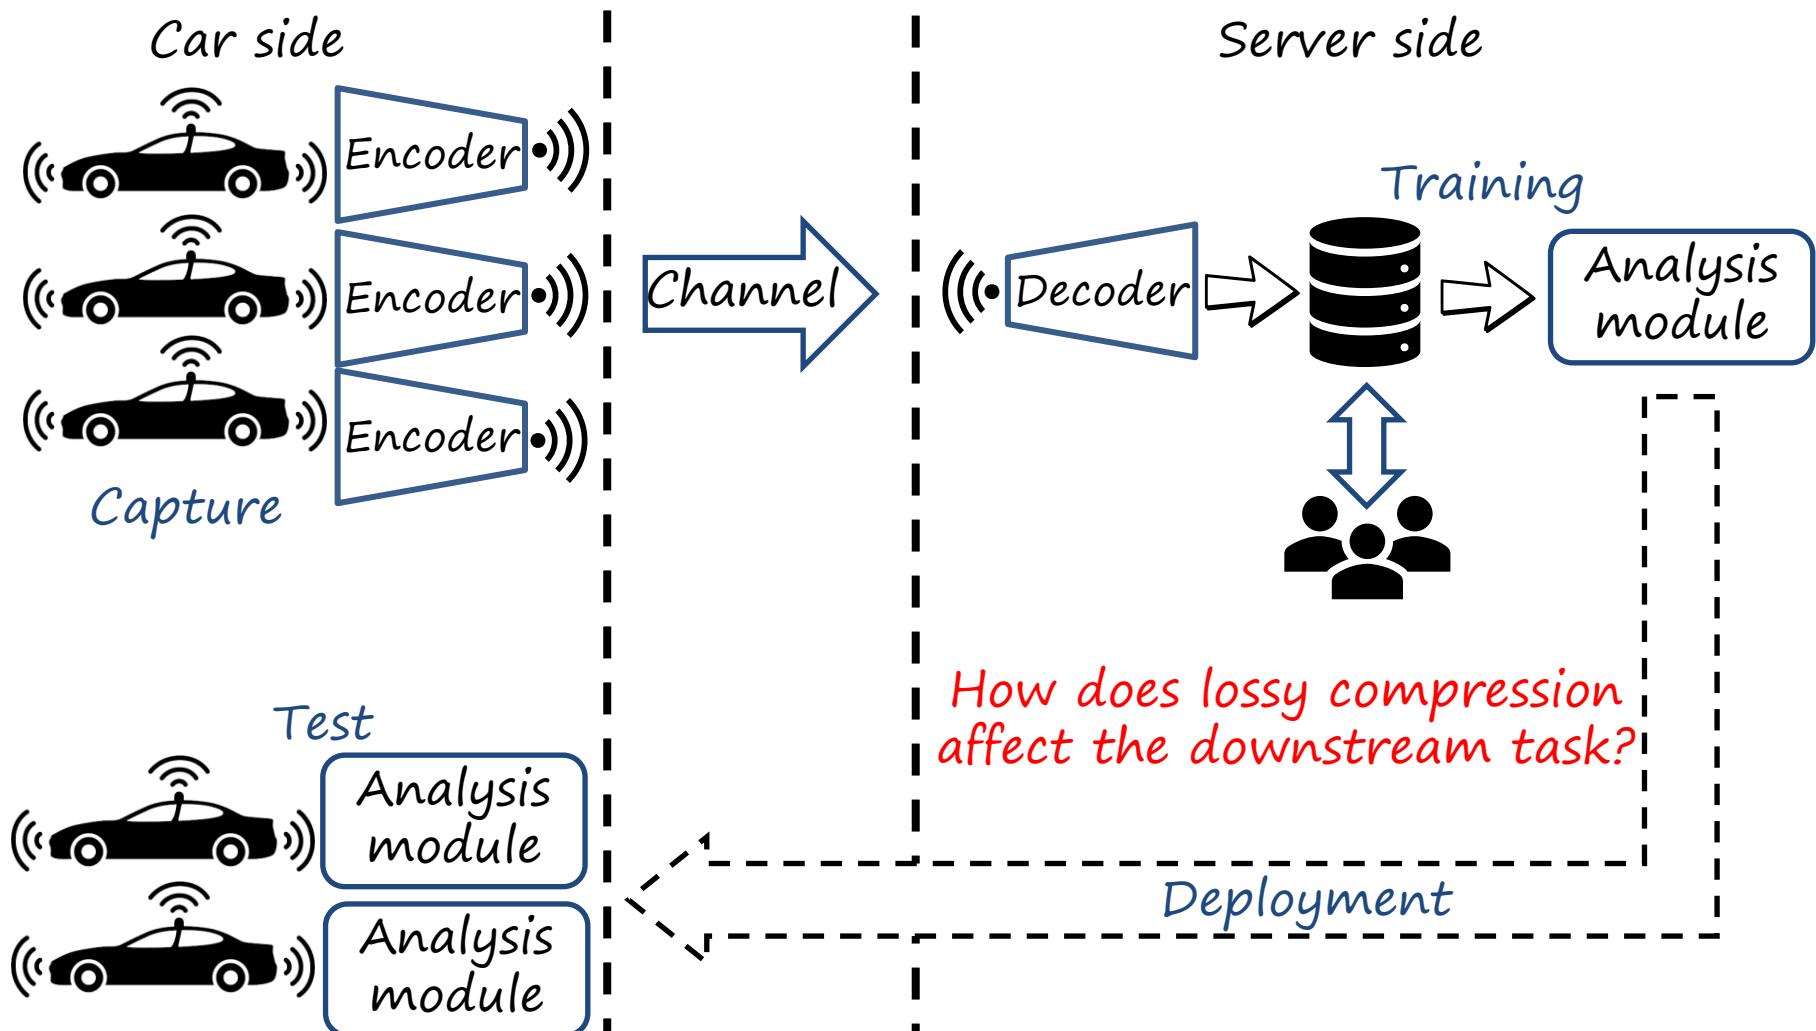
Our specific contribution:

- Use **deformable convolutions** for multiframe alignment
- **QP-conditional quality enhancement network**

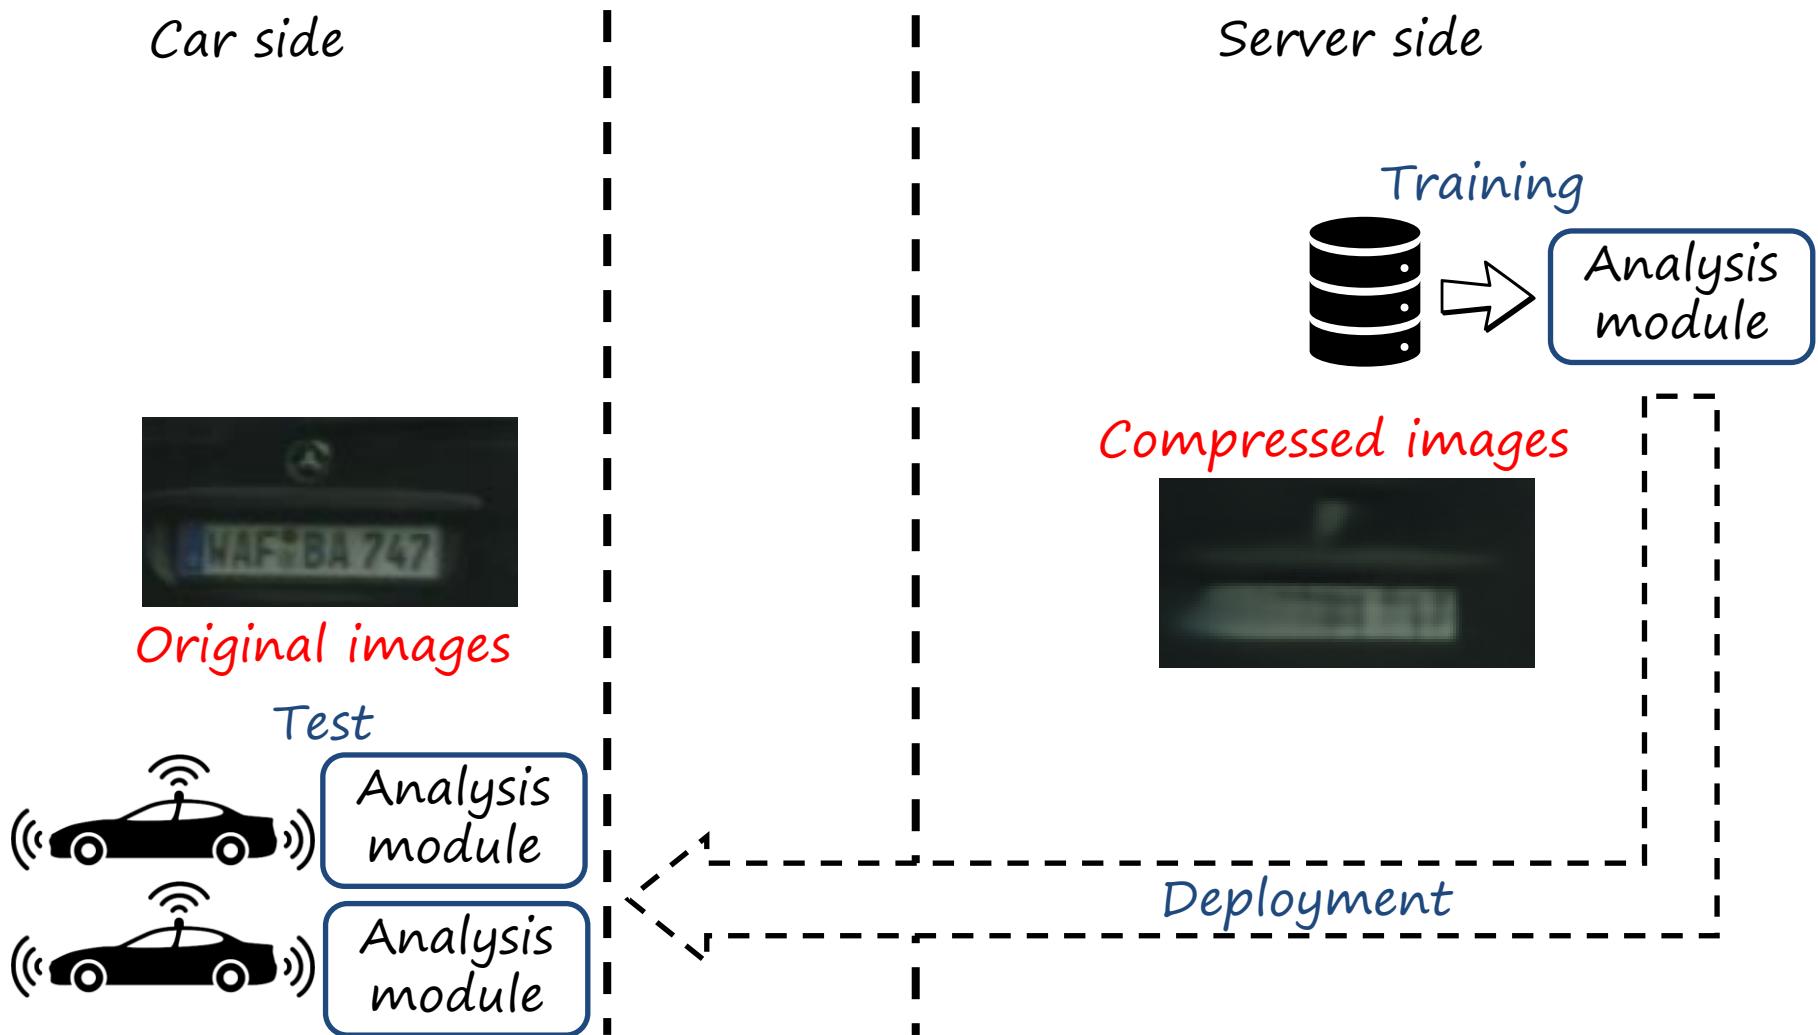
[Dai et al., Deformable Convolutional Networks, ICCV 2017](#)

[DCNGAN: A deformable convolution-based GAN with QP adaptation for perceptual quality enhancement of compressed video, ICASSP 2022](#)

Distributed data collection for autonomous driving



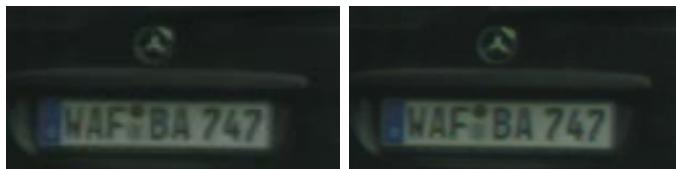
Distributed data collection



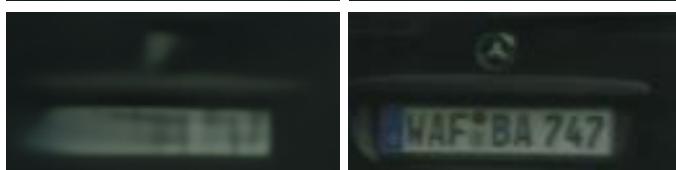
Effect on downstream task

Training

OO (ideal)



CO



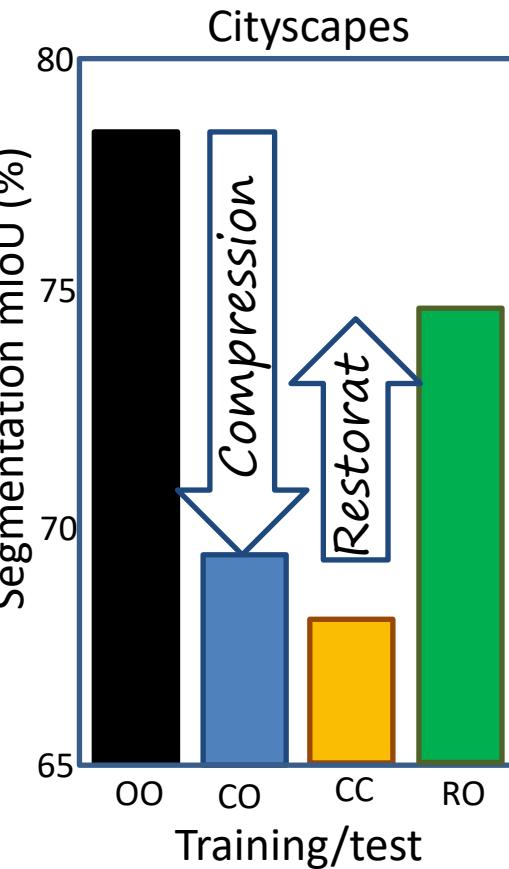
CC



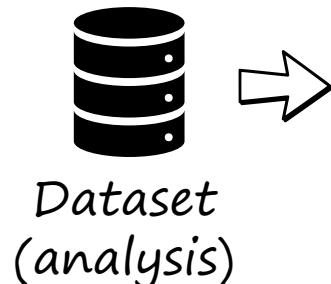
RO



Test



Proposed approach:
dataset restoration



Restoration
model



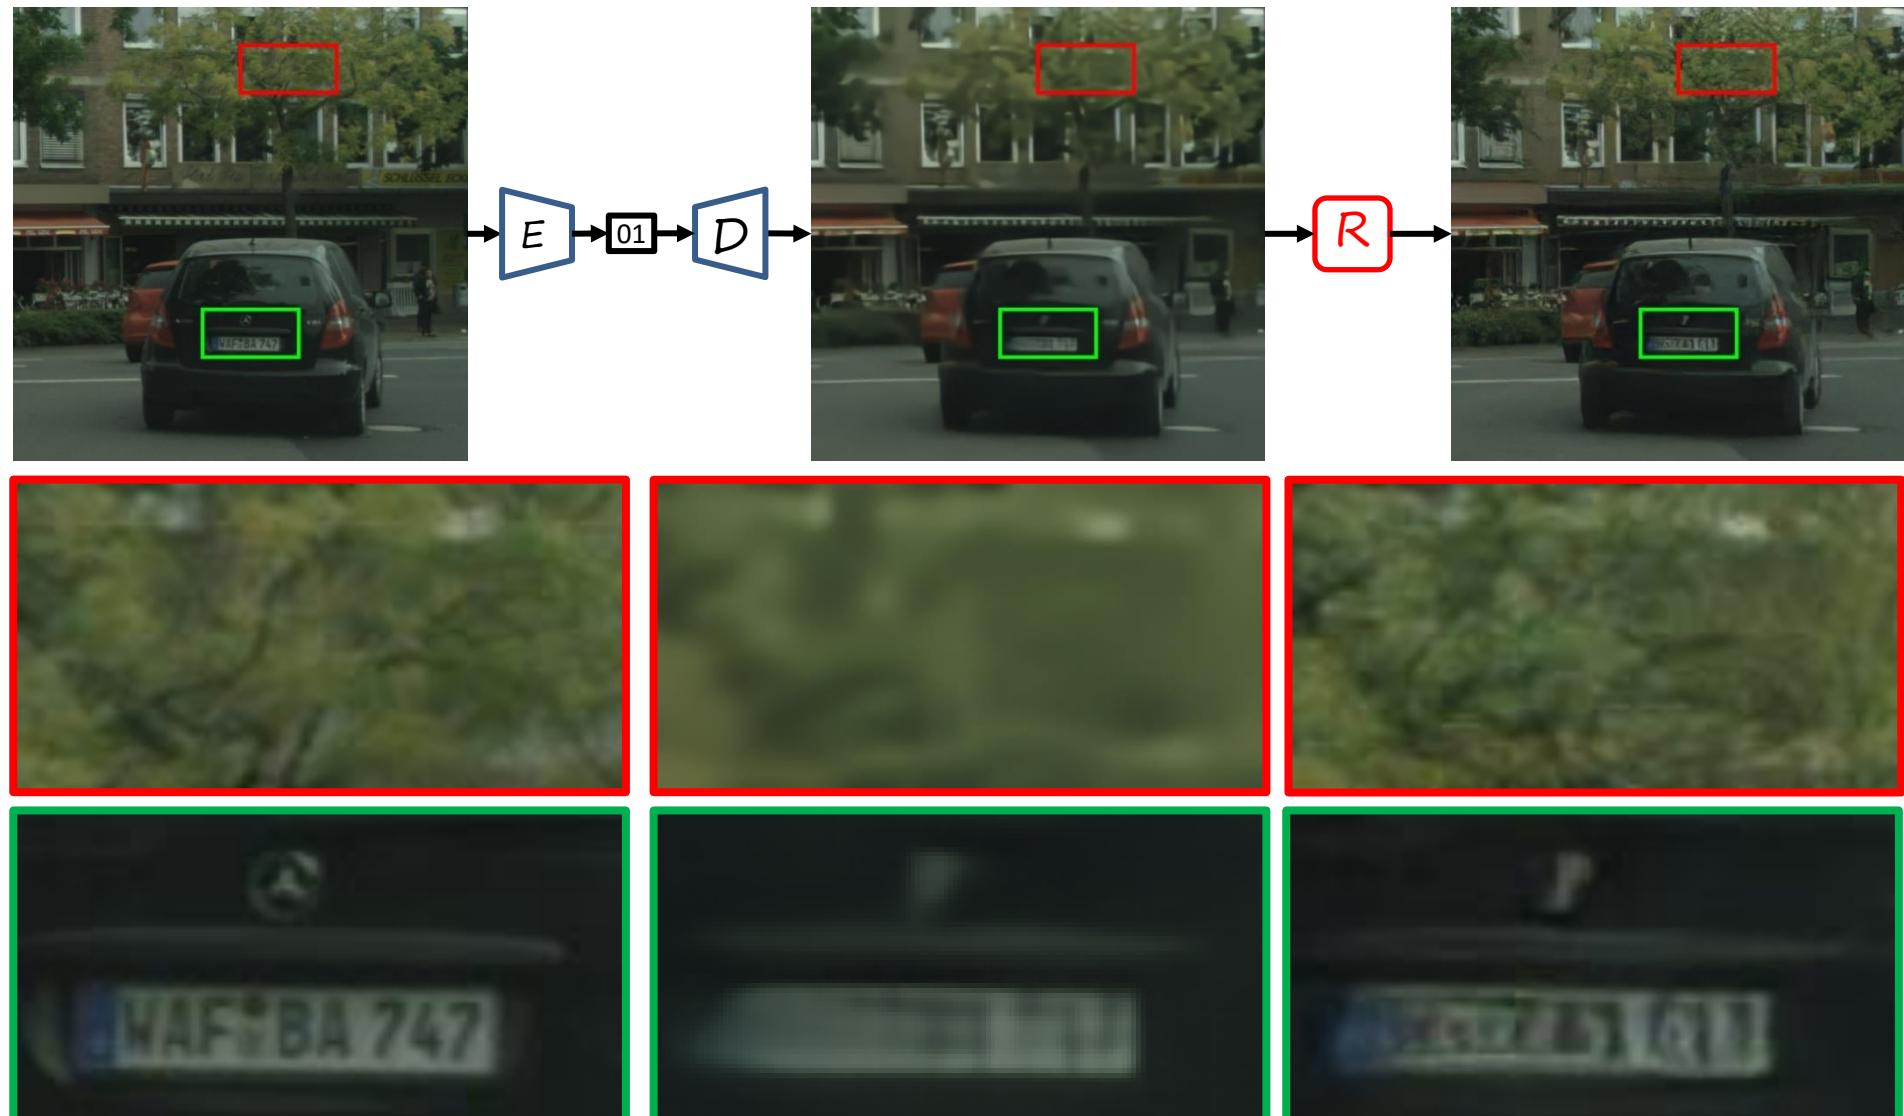
Restored
dataset

Training images vs test images

Original (test)

Compressed

Restored



References

General references

- [Balle et al. End-to-end Optimized Image Compression](#), ICLR 2017
- [Theis et al., Lossy Image Compression with Compressive Autoencoders](#), ICLR 2017
- [Dai et al., Deformable Convolutional Networks](#), ICCV 2017
- [Sun et al., Spatiotemporal Entropy Model is All You Need for Learned Video Compression](#), arxiv 2021

Works by our group and collaborators (with Marta Mrak's group at BBC R&D, London, UK and Shuai Wan's group at Northwestern Politechnic University/Xidian Uni., Xi'an, China)

- [Yang et al., Variable Rate Deep Image Compression with Modulated Autoencoder](#), Signal Processing Letters 2020
- [Yang et al., Slimmable compressive autoencoders for practical image compression](#), CVPR 2021
- [Katakol et al., DANICE: Domain adaptation without forgetting in neural image compression](#), CLIC 2021 at CVPR 2021
- [Zhang et al., DCNGAN: A deformable convolution-based GAN with QP adaptation for perceptual quality enhancement of compressed video](#), ICASSP 2022
- [Katakol et al., Distributed Learning and Inference with Compressed Images](#), IEEE Trans. Image Processing 2021
- Liu et al., Slimmable Video Codec, 2022 (under review)

THANK YOU!

lherranz@cvc.uab.es

www.lherranz.org

