The Human Touch

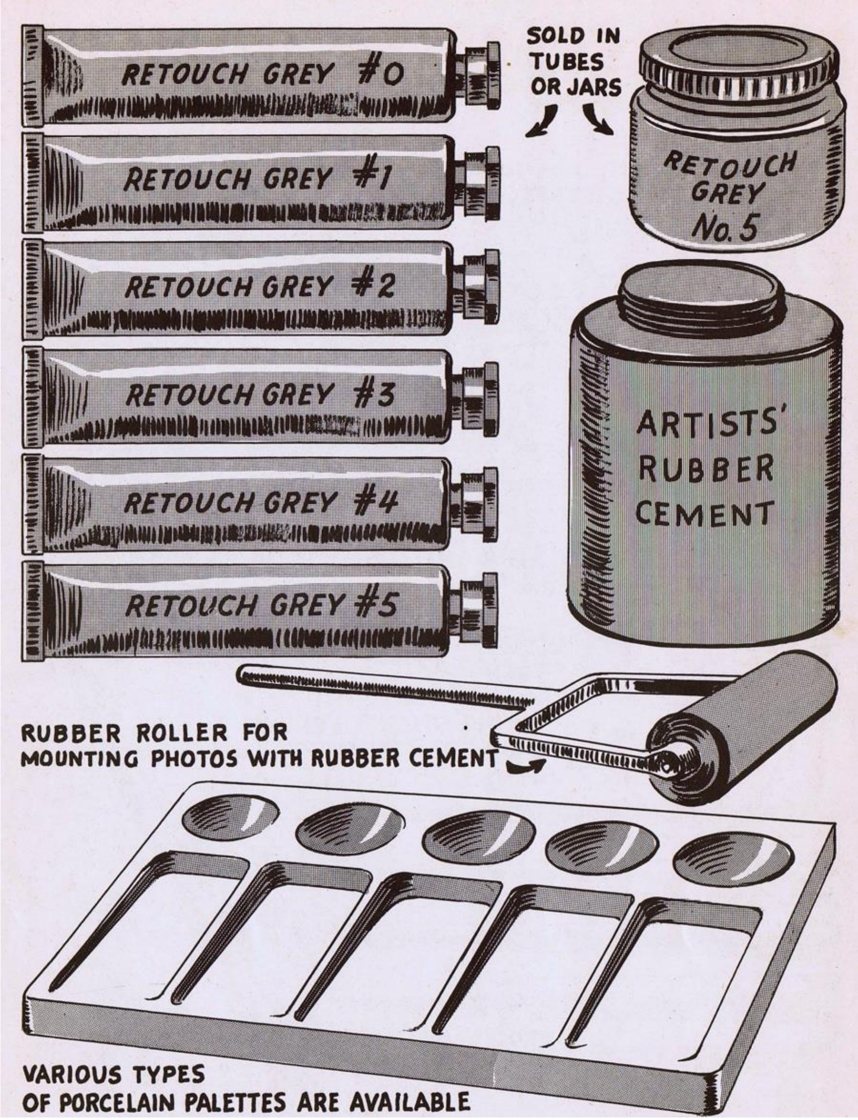

RETOUCH EQUIPMENT

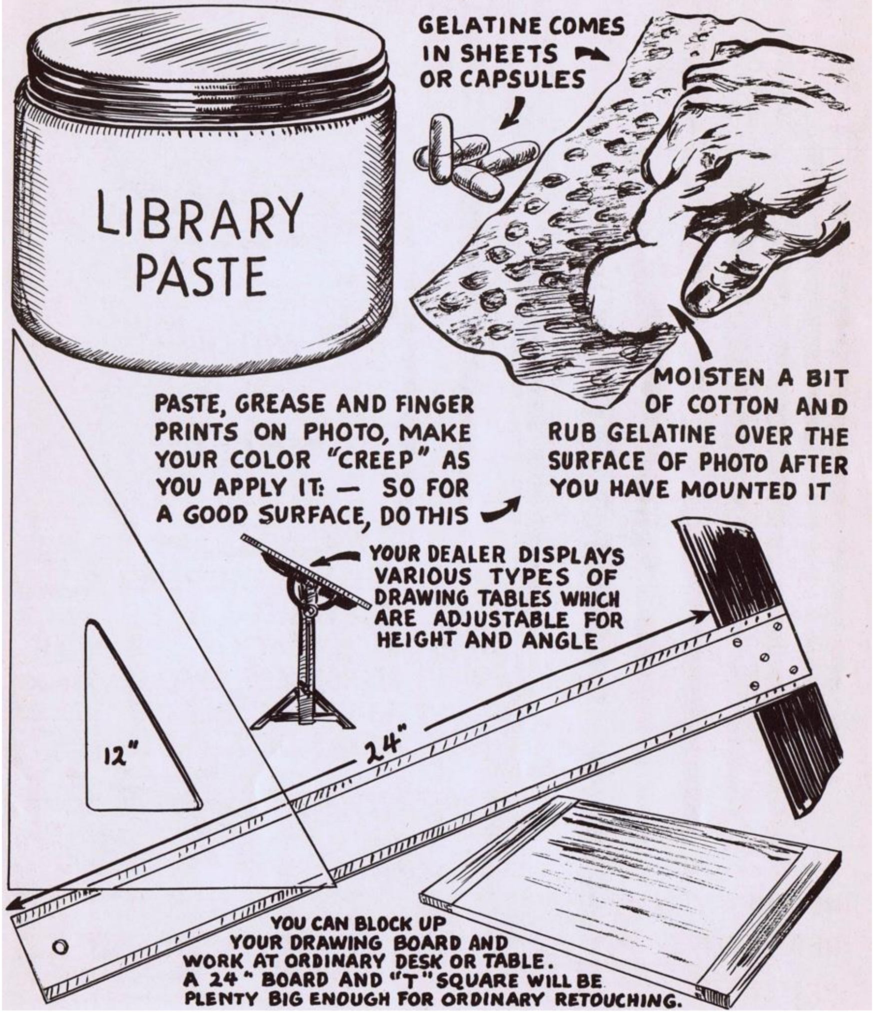

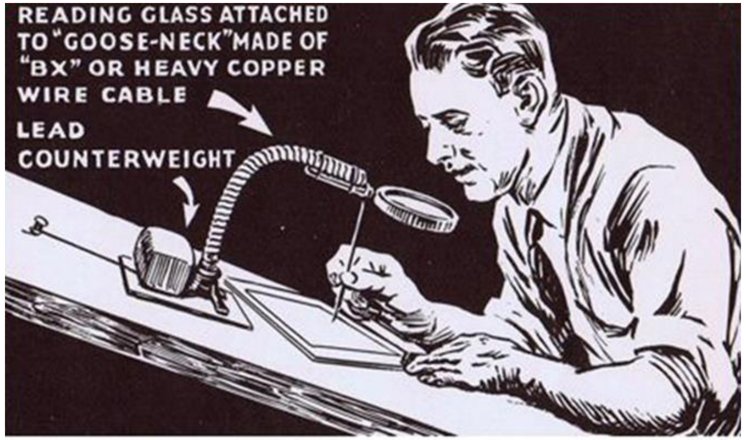

HERE'S WHAT YOU SEE

PART OF PHOTO OF WATCH DIAL MAGNIFIED AND TIP OF YOUR BRUSH SNAPPING UP THE HAIRLINE CALIBRATIONS AND NUMERALS

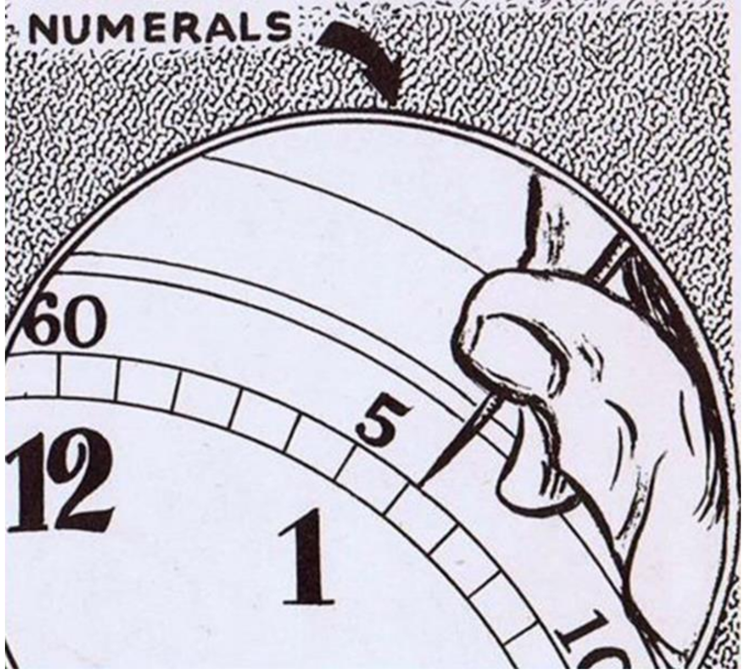


universal
Enlarger Height 65:45
20,000 LEAGUES UNDER THE SEA
CINEMASCOPE
APPLIANCES


20,000 LEAGUES
UNDER THE SEA
CINEMASCOPE
APPLIANCES

From Analog to AI


OM
39439 - UD

OM
39439 - UD

GenAI


Ps
AUTO
RETOUCH

Turning Pixels to Pictures

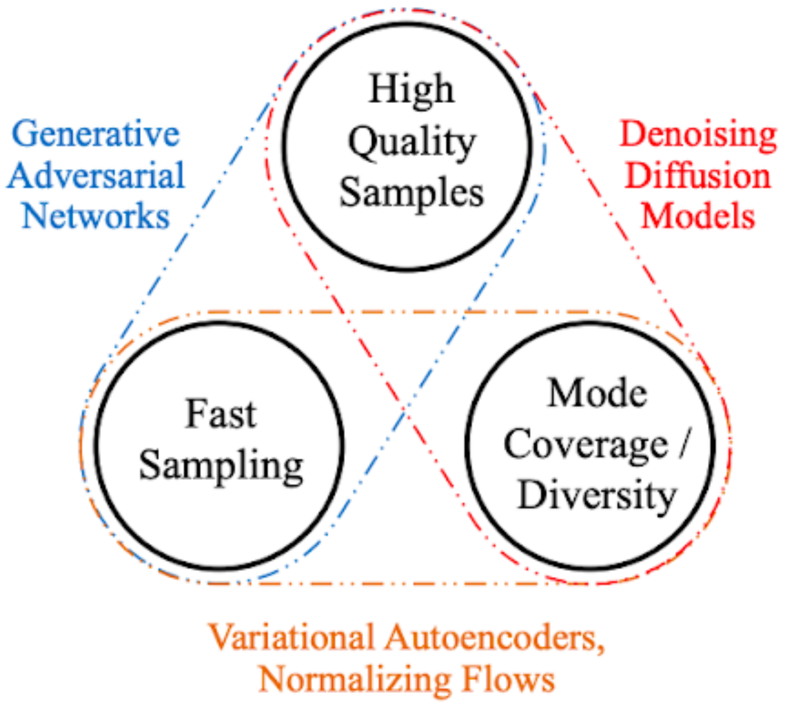

Generative Adversarial Networks
Denoising Diffusion Models
High Quality Samples
Fast Sampling
Mode Coverage / Diversity
Variational Autoencoders, Normalizing Flows

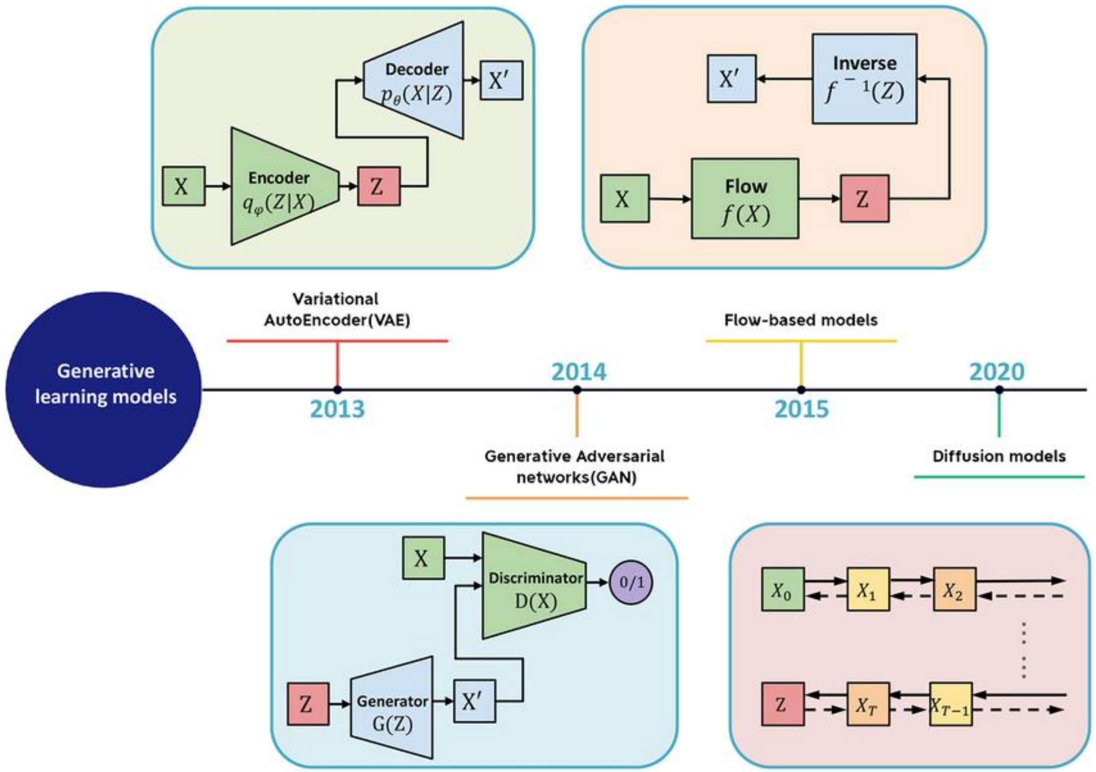

Decoder
X'
X
Encoder
Z
Inverse
Flow
Variational AutoEncoder(VAE)
Flow-based models
Generative learning models
2013
2014
2015
2020
Generative Adversarial networks(GAN)
Diffusion models
Discriminator
D(X)
0/1
Generator
G(Z)

Deep Learning requires large amounts of carefully **labeled data** which is **difficult** to acquire and **expensive** to annotate

(a) Texture image

81.4%	**Indian elephant**
10.3%	indri
8.2%	black swan

(b) Content image

71.1%	**tabby cat**
17.3%	grey fox
3.3%	Siamese cat

(c) Texture-shape cue conflict

63.9%	**Indian elephant**
26.4%	indri
9.6%	black swan

Classification predictions of a ResNet-50 trained on ImageNet

Valuable (natural) image **features** should not be specialized for solving **a particular supervised task**, but rather encapsulate richer chracteristics exploitable for **various downstream tasks**

How Much Information is the Machine Given during Learning?

- **“Pure” Reinforcement Learning (cherry)**
  - The machine predicts a scalar reward given once in a while.
  - **A few bits for some samples**
- **Supervised Learning (icing)**
  - The machine predicts a category or a few numbers for each input
  - Predicting human-supplied data
  - **10→10,000 bits per sample**
- **Self-Supervised Learning (cake génoise)**
  - The machine predicts any part of its input for any observed part.
  - Predicts future frames in videos
  - **Millions of bits per sample**

Input: `The man went to the [MASK]`$_1$ `. He bought a [MASK]`$_2$ `of milk .`
Labels: `[MASK]`$_1$ `= store; [MASK]`$_2$ `= gallon`

Missing word prediction task.

Sentence A = `The man went to the store.`
Sentence B = `He bought a gallon of milk.`
Label = `IsNextSentence`

Sentence A = `The man went to the store.`
Sentence B = `Penguins are flightless.`
Label = `NotNextSentence`

So What?

Let there be Color!

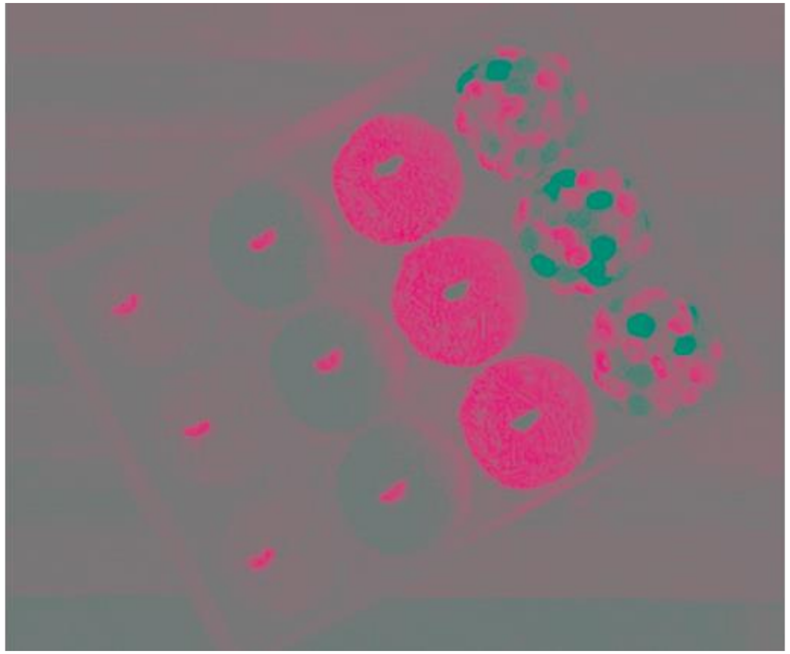

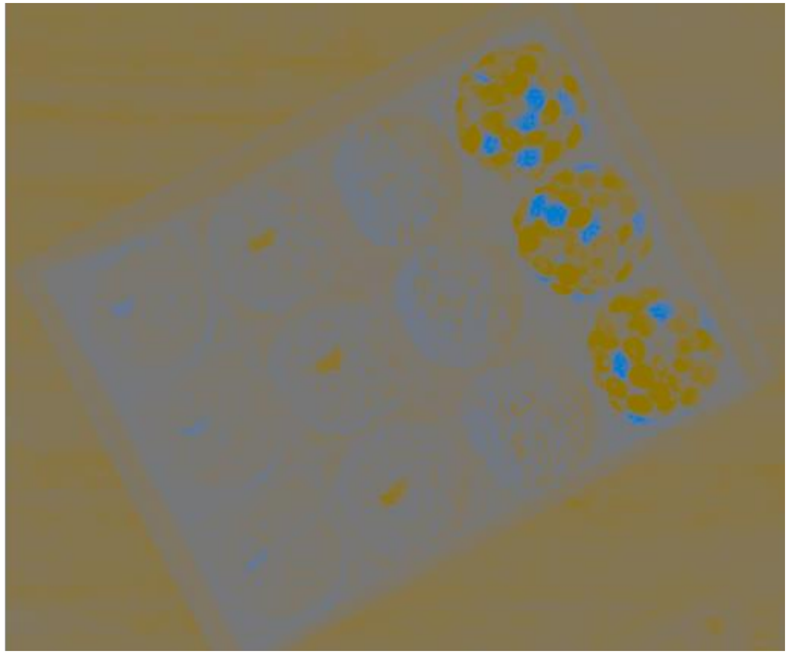

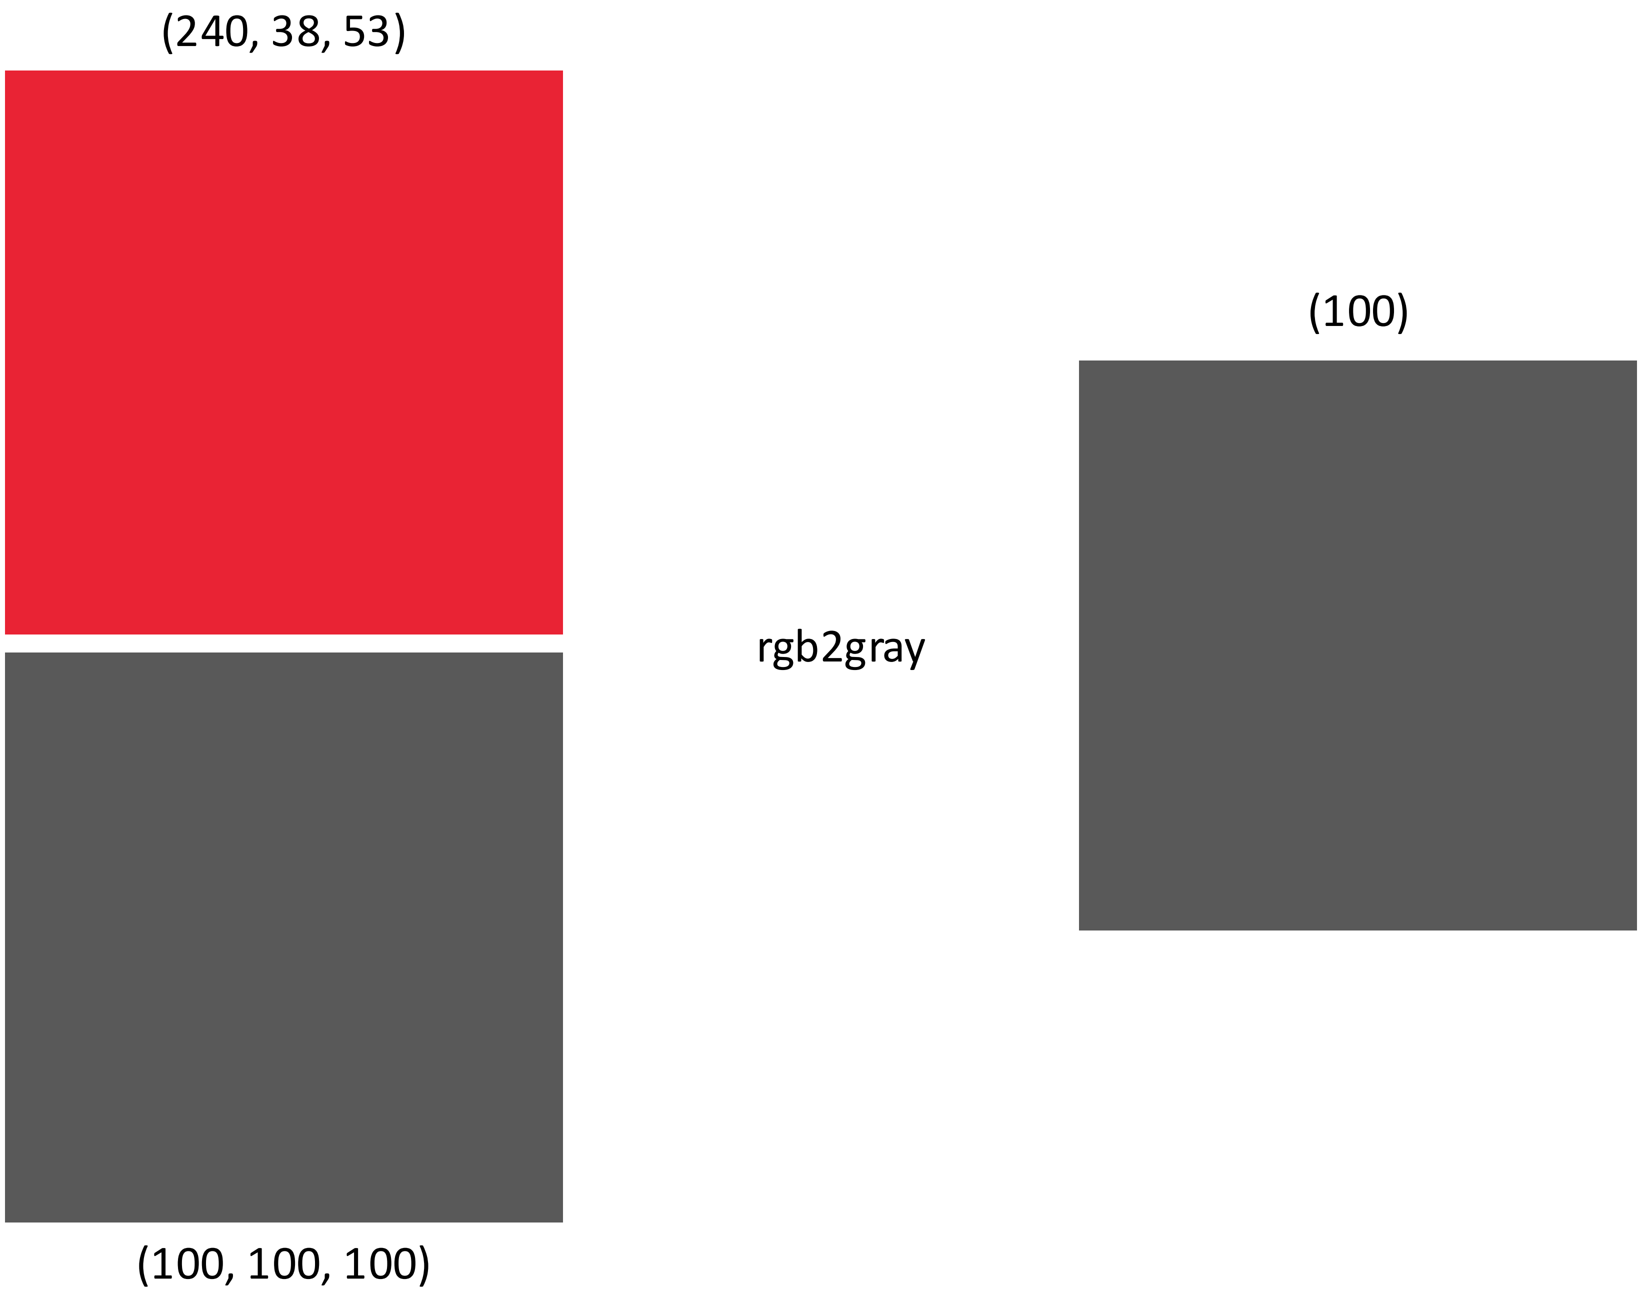
(240, 38, 53)
(100, 100, 100)
rgb2gray
(100)

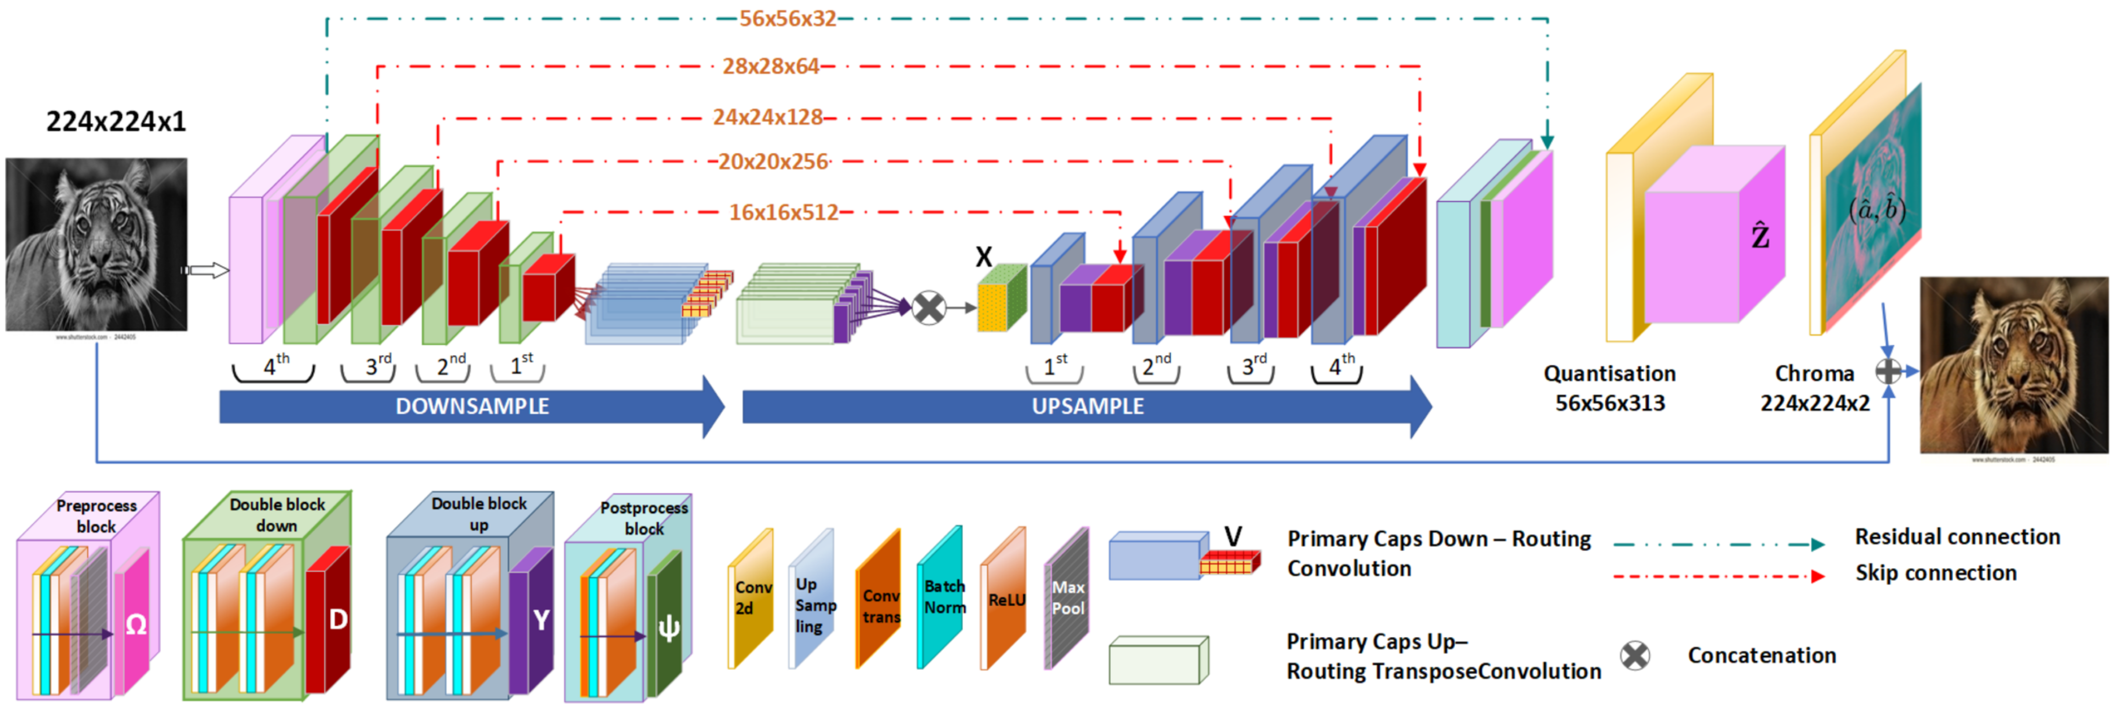

R. Pucci, C. Micheloni, G. L. Foresti, N. Martinel – *Is It a Plausible Colour? UCapsNet for Image Colourisation.* NeurIPS Workshop: Self-Supervised Learning -- Theory and Practice (2021)
R. Pucci, C. Micheloni, N. Martinel – *Collaborative image and object level features for image colourisation.* CVPR (2021)
R. Pucci, C. Micheloni, G. L. Foresti, N. Martinel – *ProCCaps: Progressively Teaching Colourisation to Capsules*. WACV (2022)

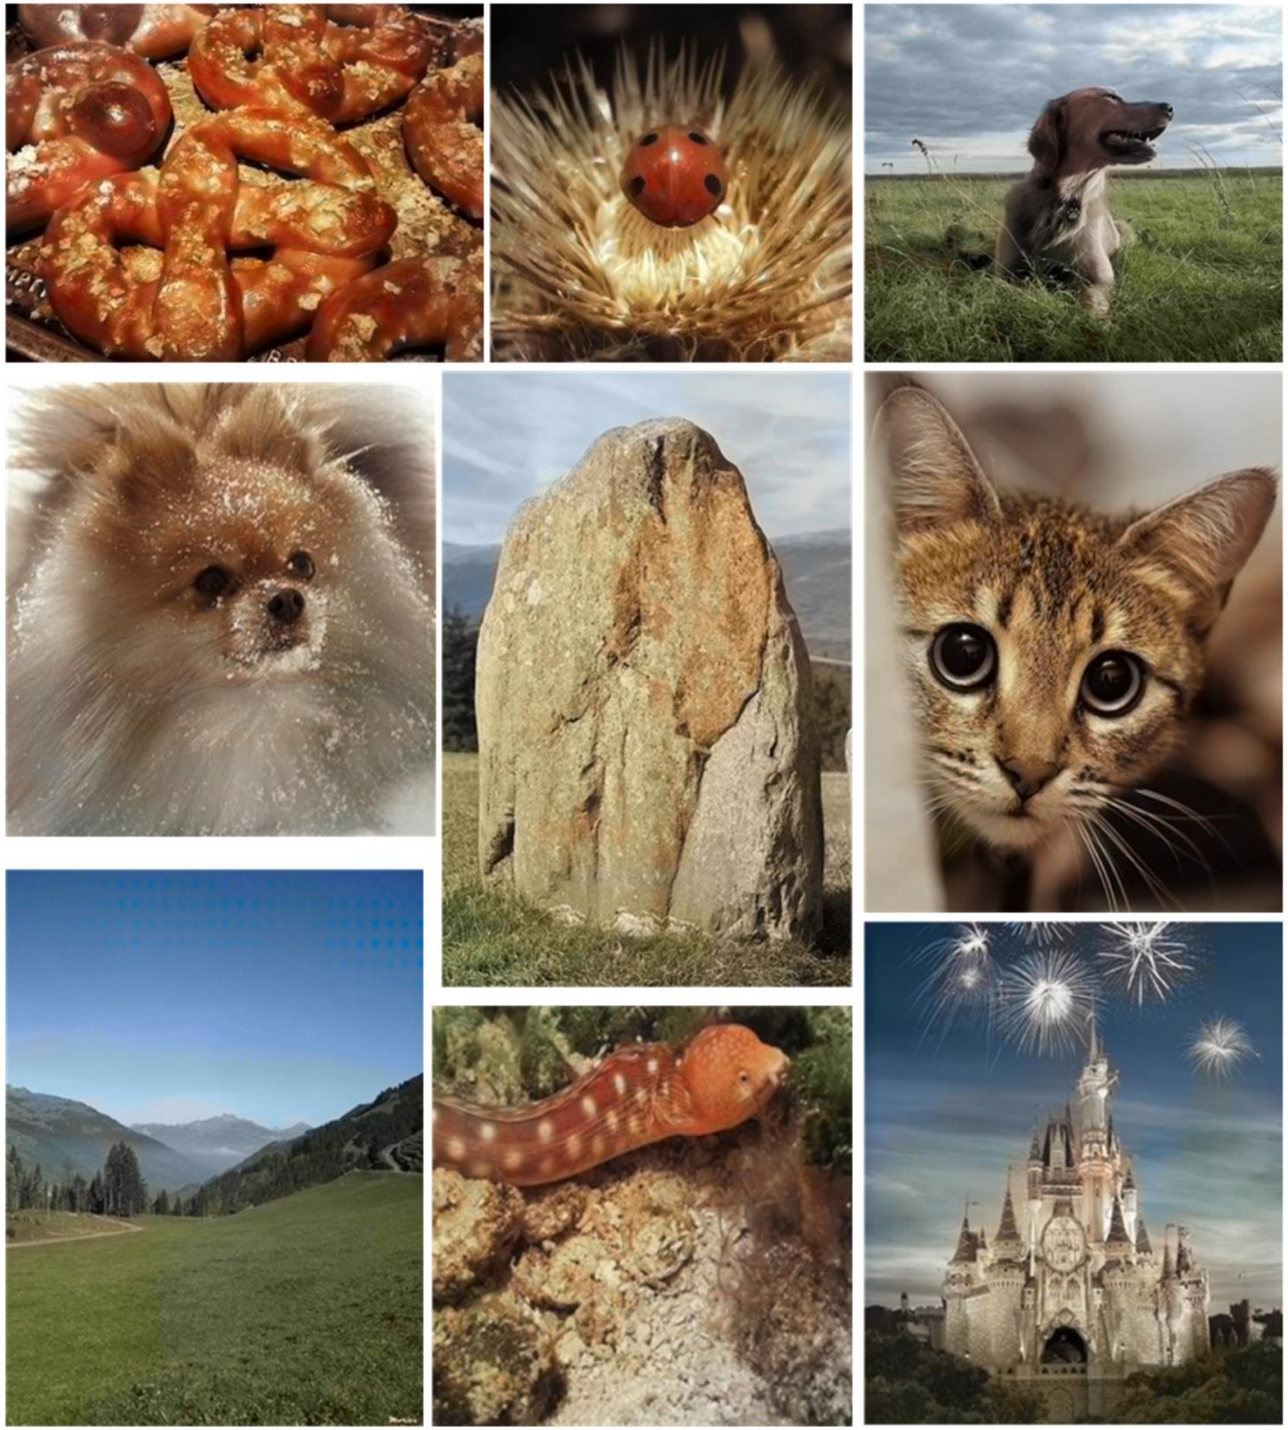

Henri Cartier Bresson, Ansel Adams, Alinari. Historical archives.

Where are the details?

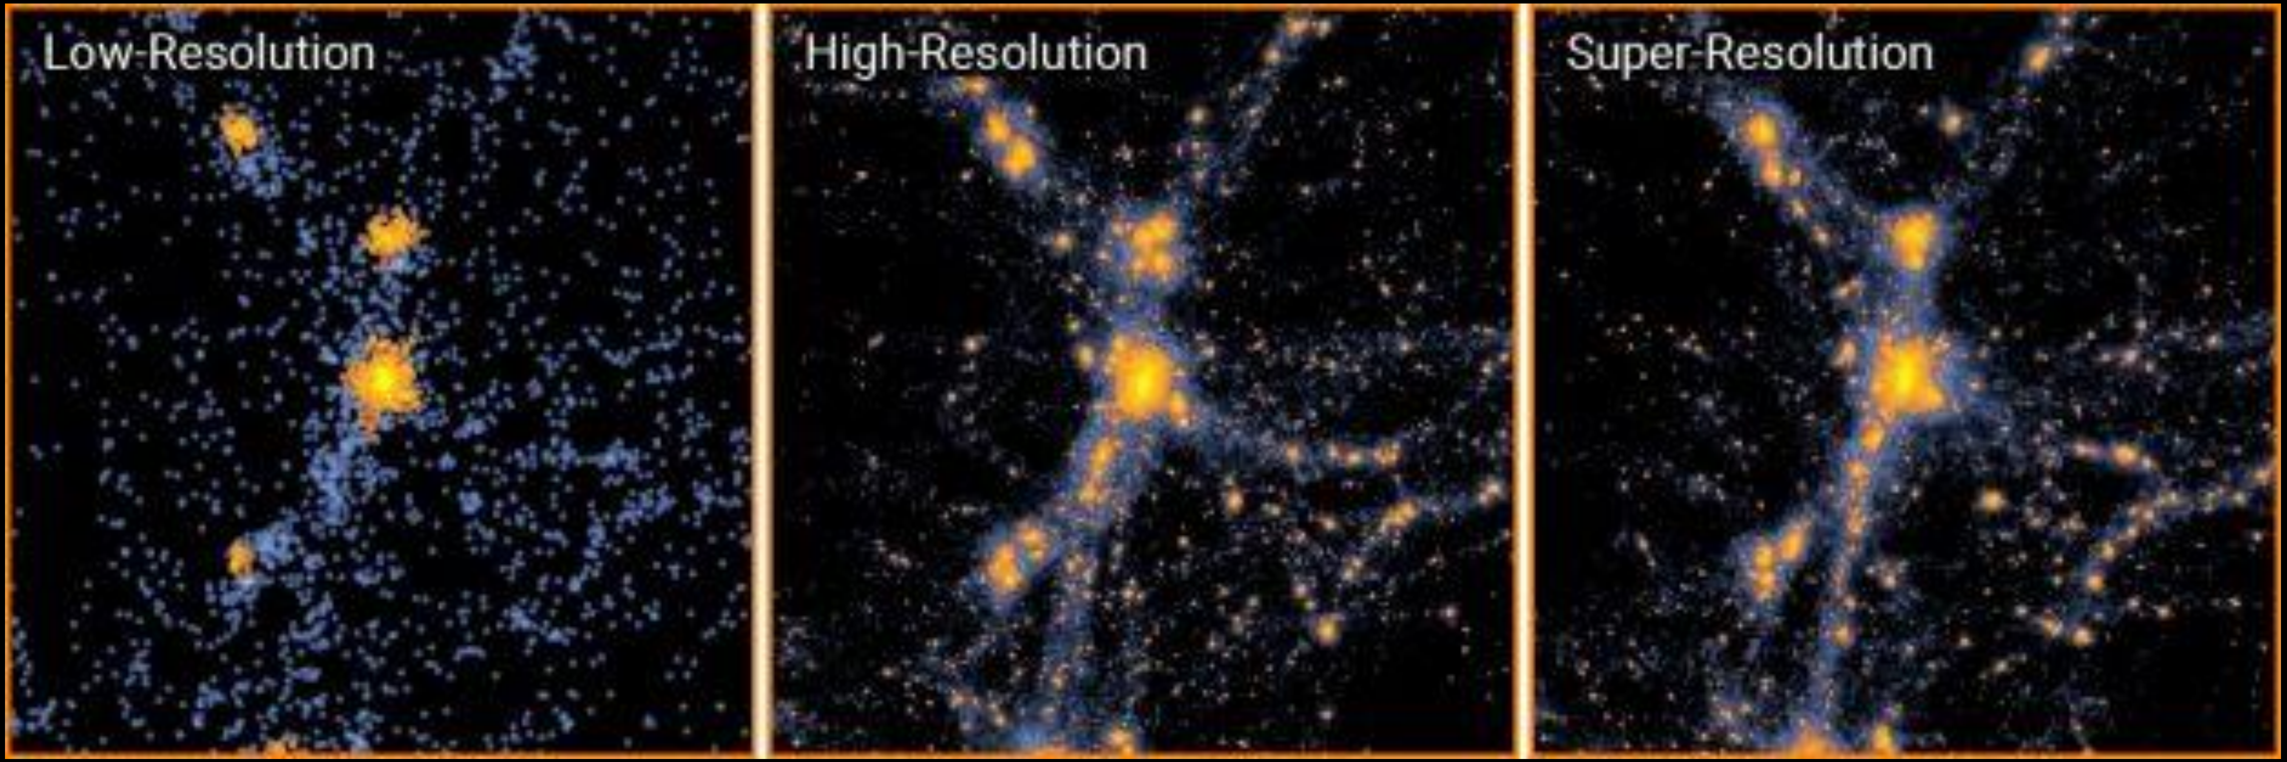

Yin Li, Yueying Ni, Rupert A. C. Croft, Tiziana Di Matteo, Simeon Bird, and Yu Feng, *AI-assisted superresolution cosmological simulations*, PNAS 2021

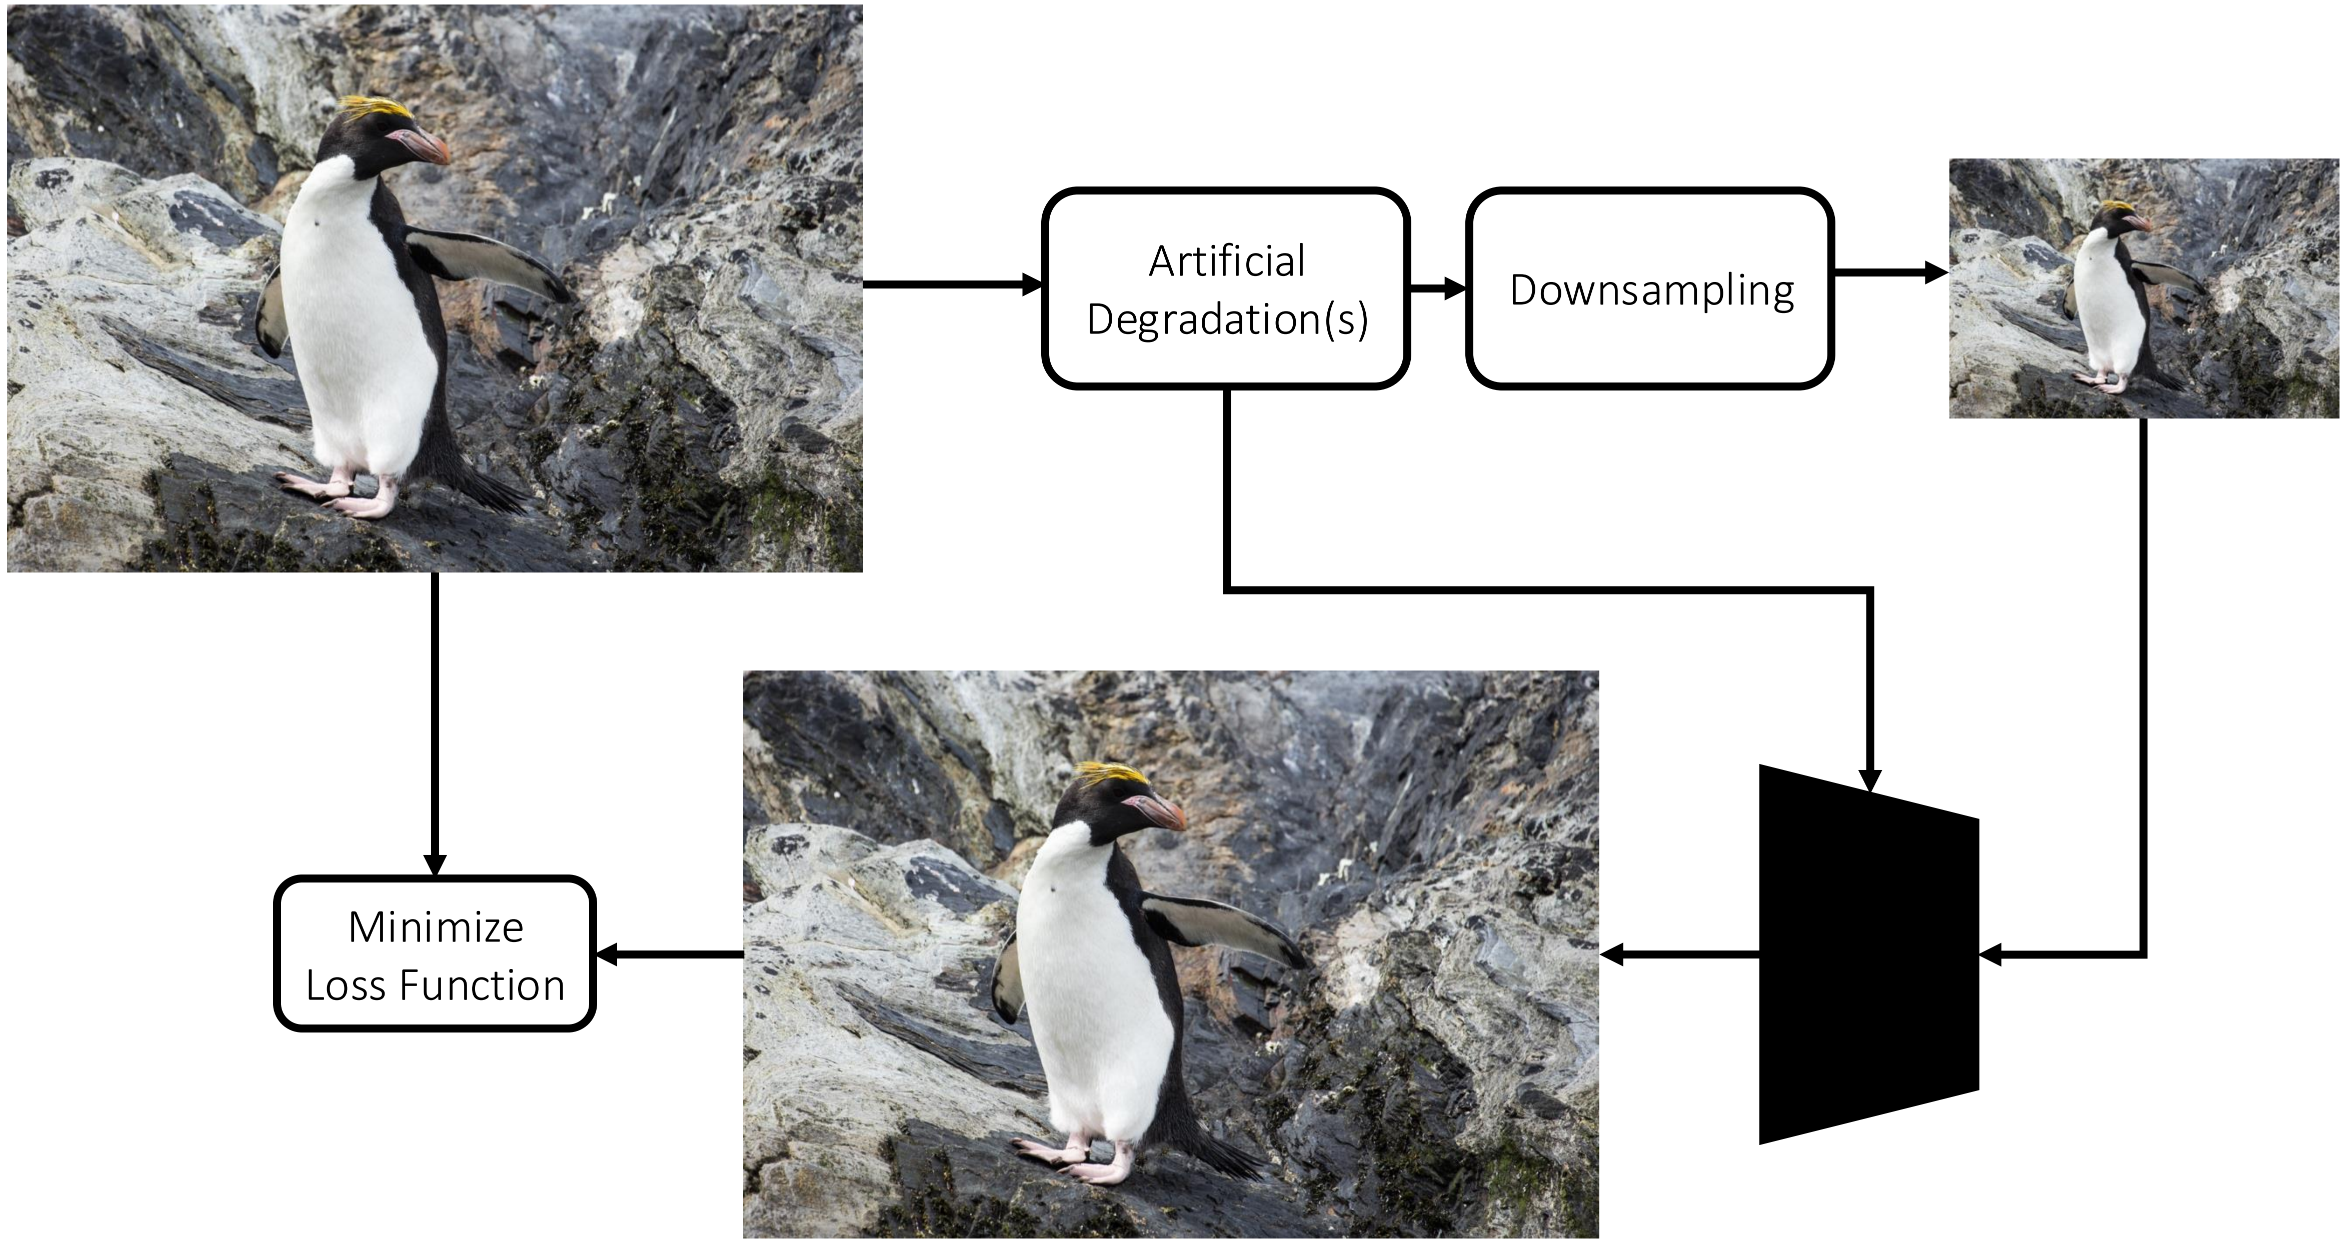
Artificial
Degradation(s)
Downsampling
Minimize
Loss Function

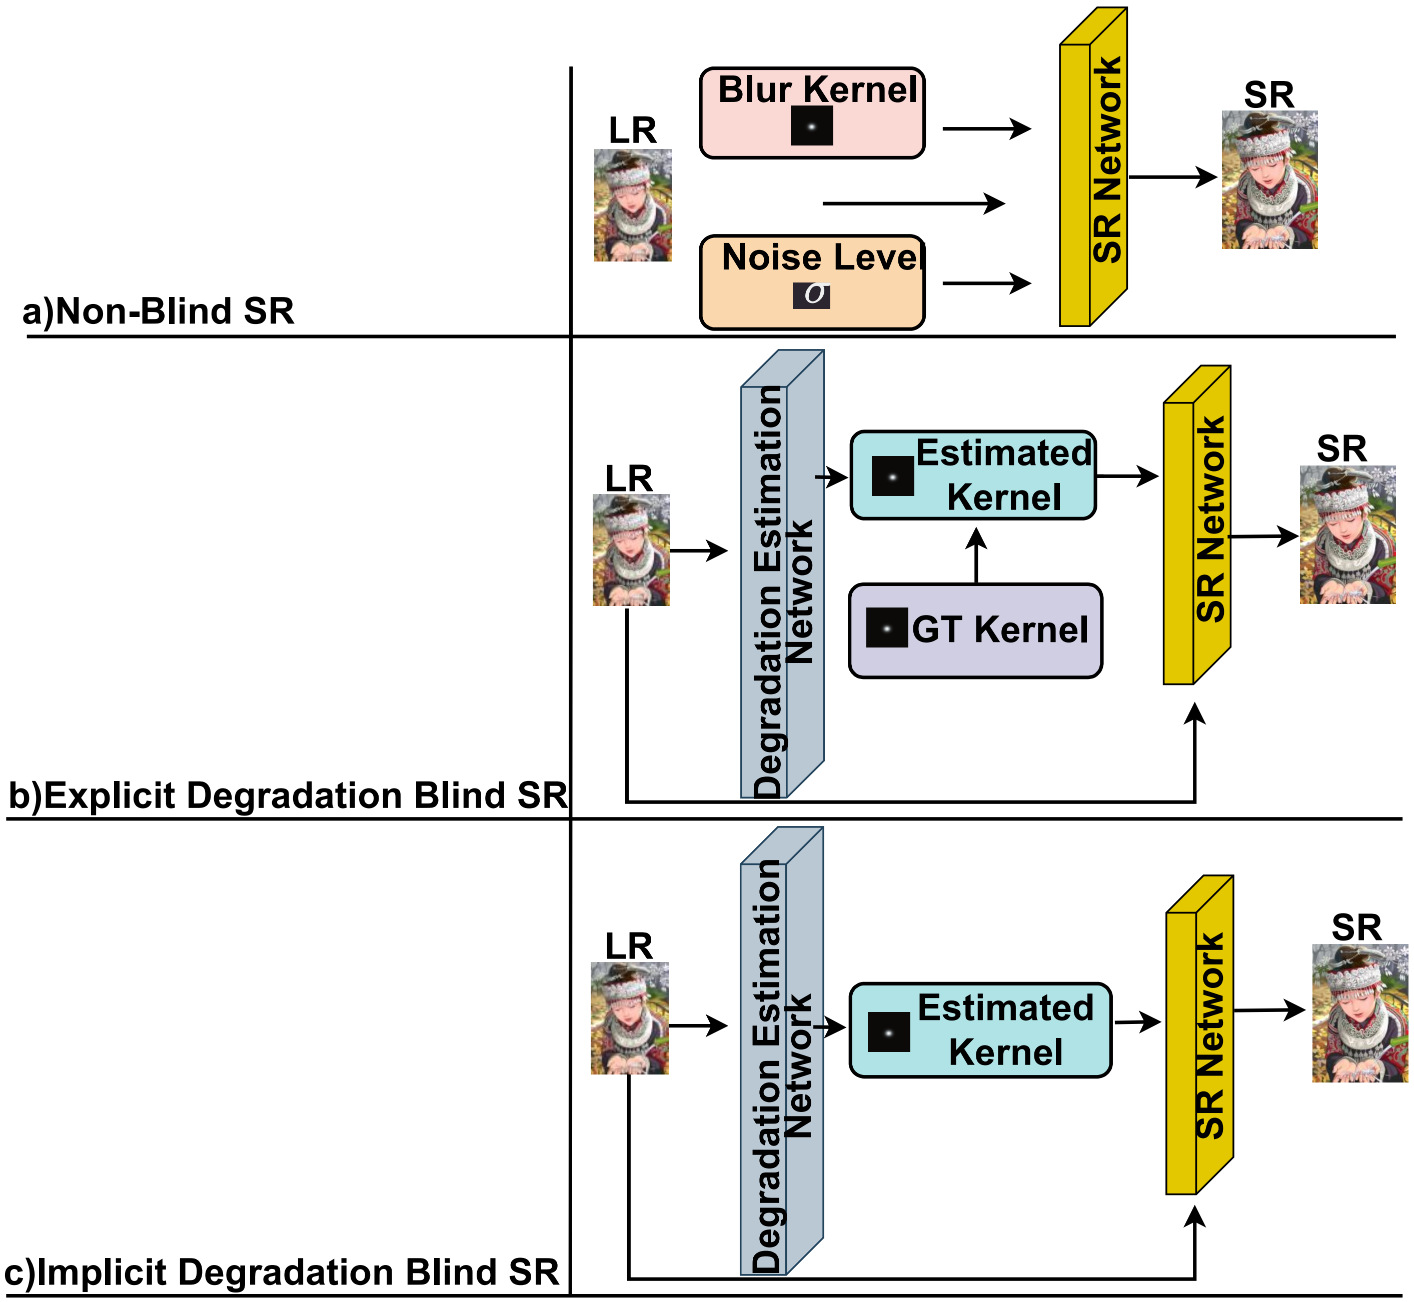

a)Non-Blind SR
LR
Blur Kernel
Noise Level
σ
SR Network
SR
b)Explicit Degradation Blind SR
LR
Degradation Estimation Network
Estimated Kernel
GT Kernel
SR Network
SR
c)Implicit Degradation Blind SR
LR
Degradation Estimation Network
Estimated Kernel
SR Network
SR

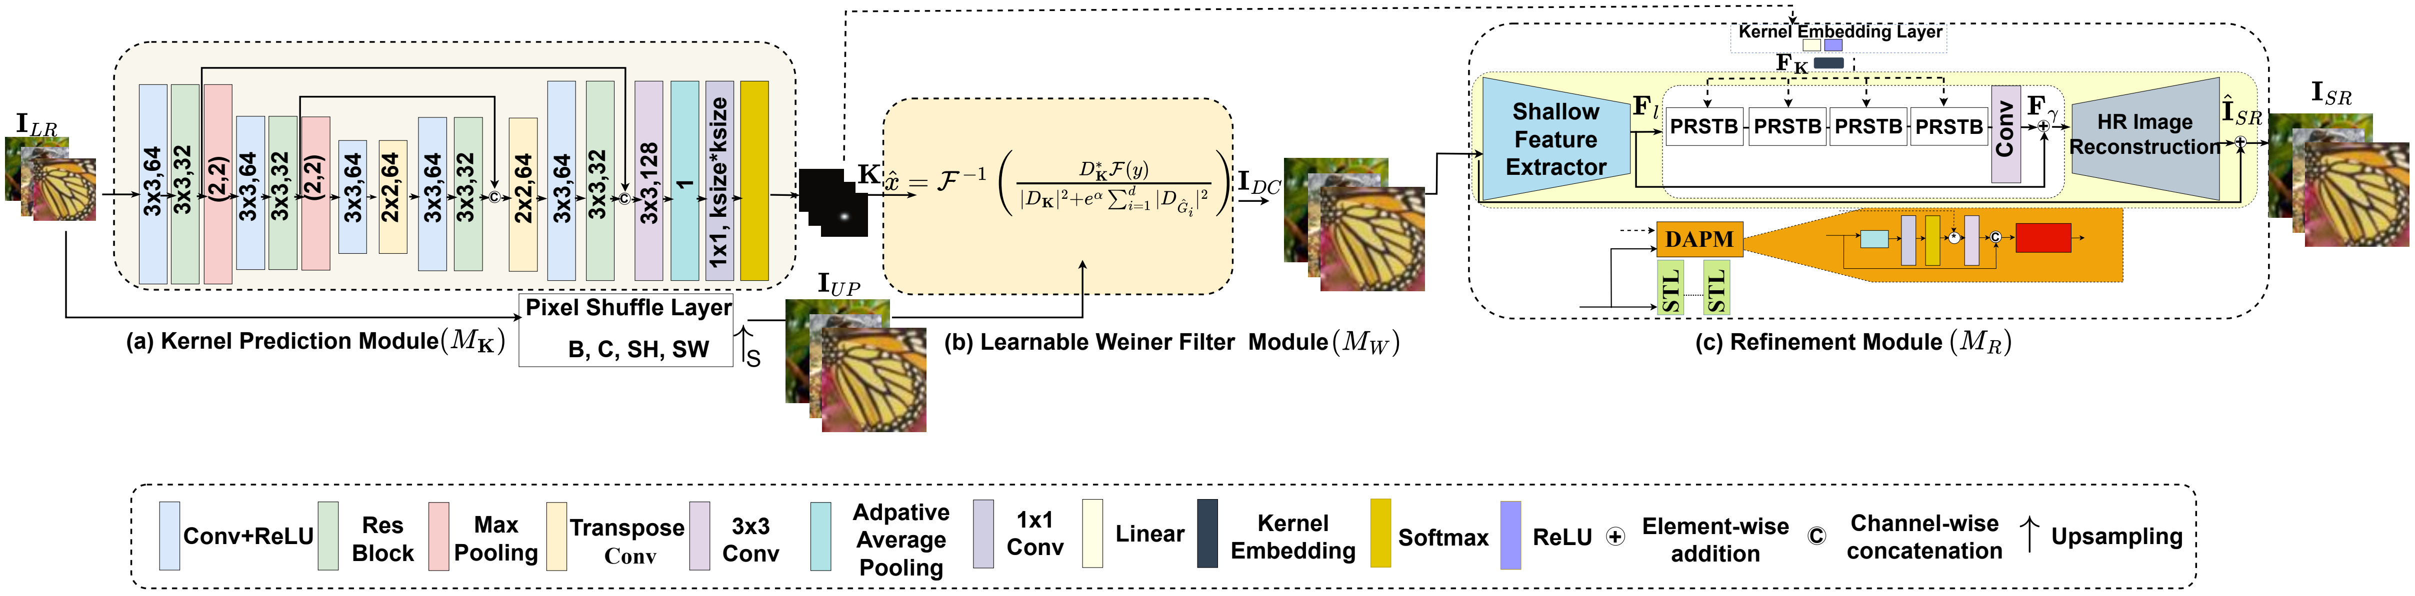

Khan, Asif Hussain, Christian Micheloni, and Niki Martinel. IDENet: Implicit Degradation Estimation Network for Efficient Blind Super Resolution. IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (2024)

Khan, Asif Hussain, Christian Micheloni, and Niki Martinel. Lightweight Prompt Learning Implicit Degradation Estimation Network for Blind Super Resolution. IEEE Transactions on Image Processing (2024)

Now What?

Thanks!

niki.martinel@uniud.it

   

Machine Learning and Perception Lab
Università degli Studi di Udine