Lecture 1: Introduction

Jun-Yan Zhu

16-726, Spring 2024

Teaching Staff

Instructors

Jun-Yan Zhu

junyanz at
cs.cmu.edu

Teaching Assistants

Jaskaran Singh Sodhi

jsodhi at
andrew.cmu.edu

Hariharan Ravichandran

hravicha at
andrew.cmu.edu

Jun-Yan Zhu

- Computer Vision, Computer Graphics, Machine Learning, Computational Photography
- Love pets (cat & dog)
- Swimming, tennis, gaming.

Cat Paper Collection

As reported by Cisco, 90% of net traffic will be visual, and indeed, most of the visual data are cat photos and videos. Thus, understanding, modeling, and synthesizing our feline friends becomes a more and more critical research problem these days, especially for our cat lovers.

Cat Paper Collection is an academic paper collection that includes computer graphics, computer vision, and machine learning papers that produce experimental results related to **cats**. If you would like to add/remove an article, please send an email to **Jun-Yan Zhu** (junyanz at cs dot cmu dot edu). We thank all the authors for their contribution and support.

See also **GitHub** | **CSV file**

Efficient Halftoning via Deep Reinforcement Learning
Haitian Jiang, Dongliang Xiong, Xiaowen Jiang, Li Ding, Liang Chen, Kai Huang
In IEEE TIP 2023
[Paper]

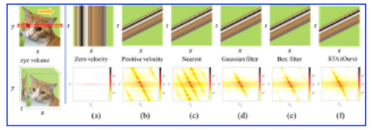

Learning Spatio-Temporal Downsampling for Effective Video Upscaling
Xiaoyu Xiang, Yapeng Tian, Vijay Rengarajan, Lucas Young, Bo Zhu, Rakesh Ranjan
In ECCV 2022
[Paper]

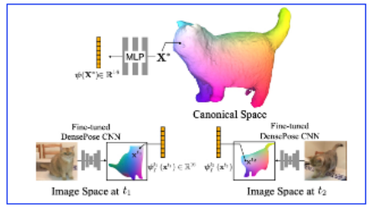

BANMo: Building Animatable 3D Neural Models from Many Casual Videos
Gengshan Yang, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, Hanbyul Joo
In CVPR 2022
[Paper] [Project]

Jaskaran Singh Sodhi

- MRSD student at RI
- Working on Subcanopy Wildfire Monitoring with AirLab
- Interested in applications of deep learning to perception and motion planning

Hariharan Ravichandran

- 2nd year MRSD student in RI
- Project: “Autonomous Construction on Lunar-like Terrains” under Prof. Red Whittaker
- Currently working on learning-based planning methods for off-road driving

Visual Content Creation

Cave art

Time

32,000 BC

Visual Content Creation

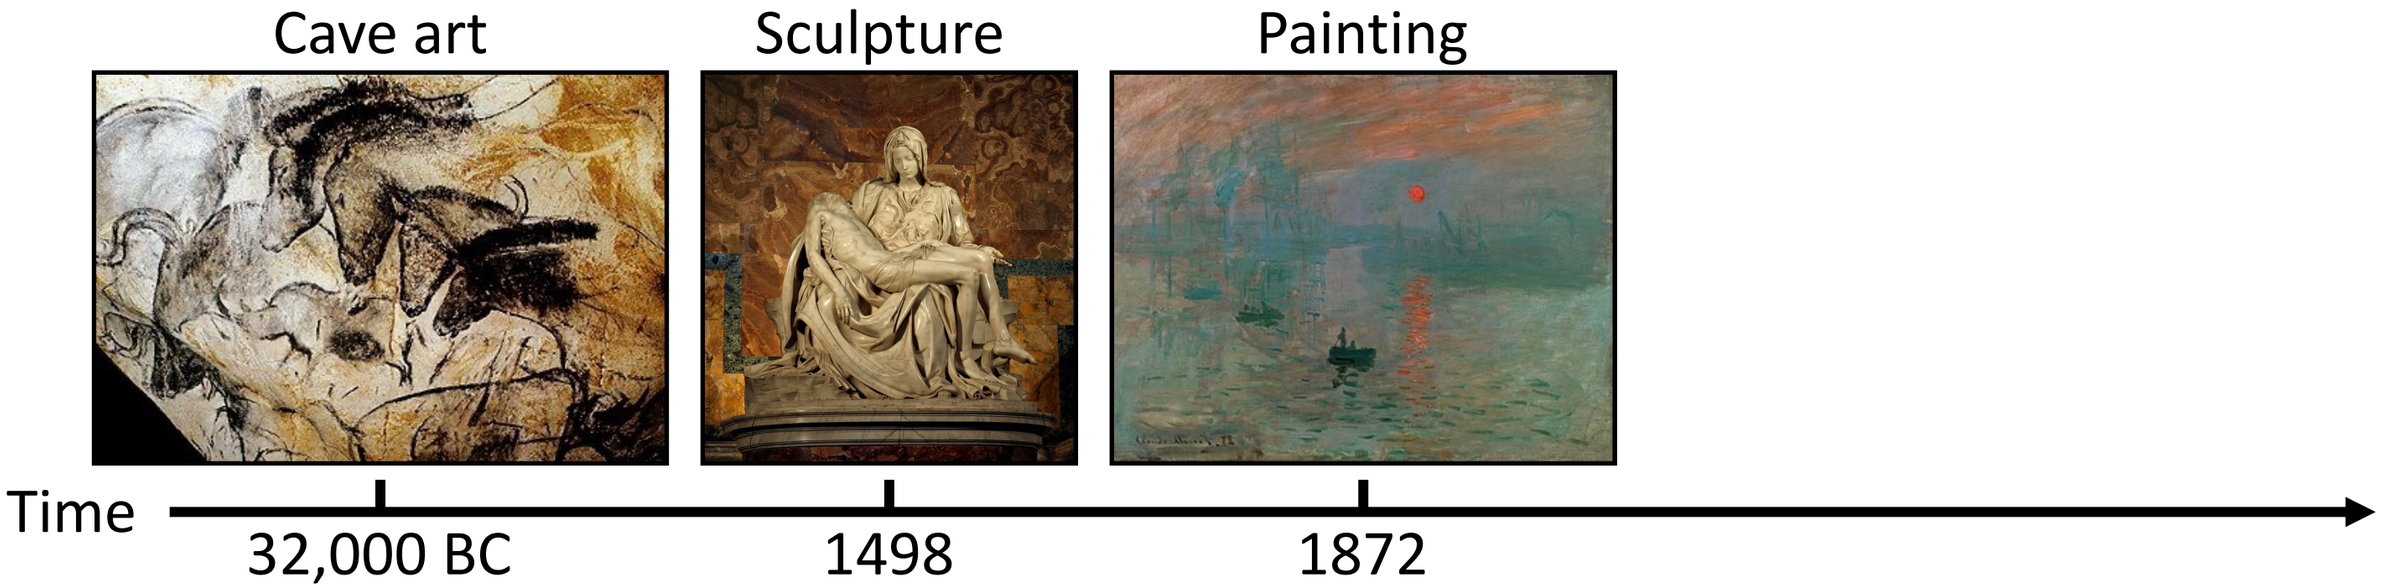

Visual Content Creation

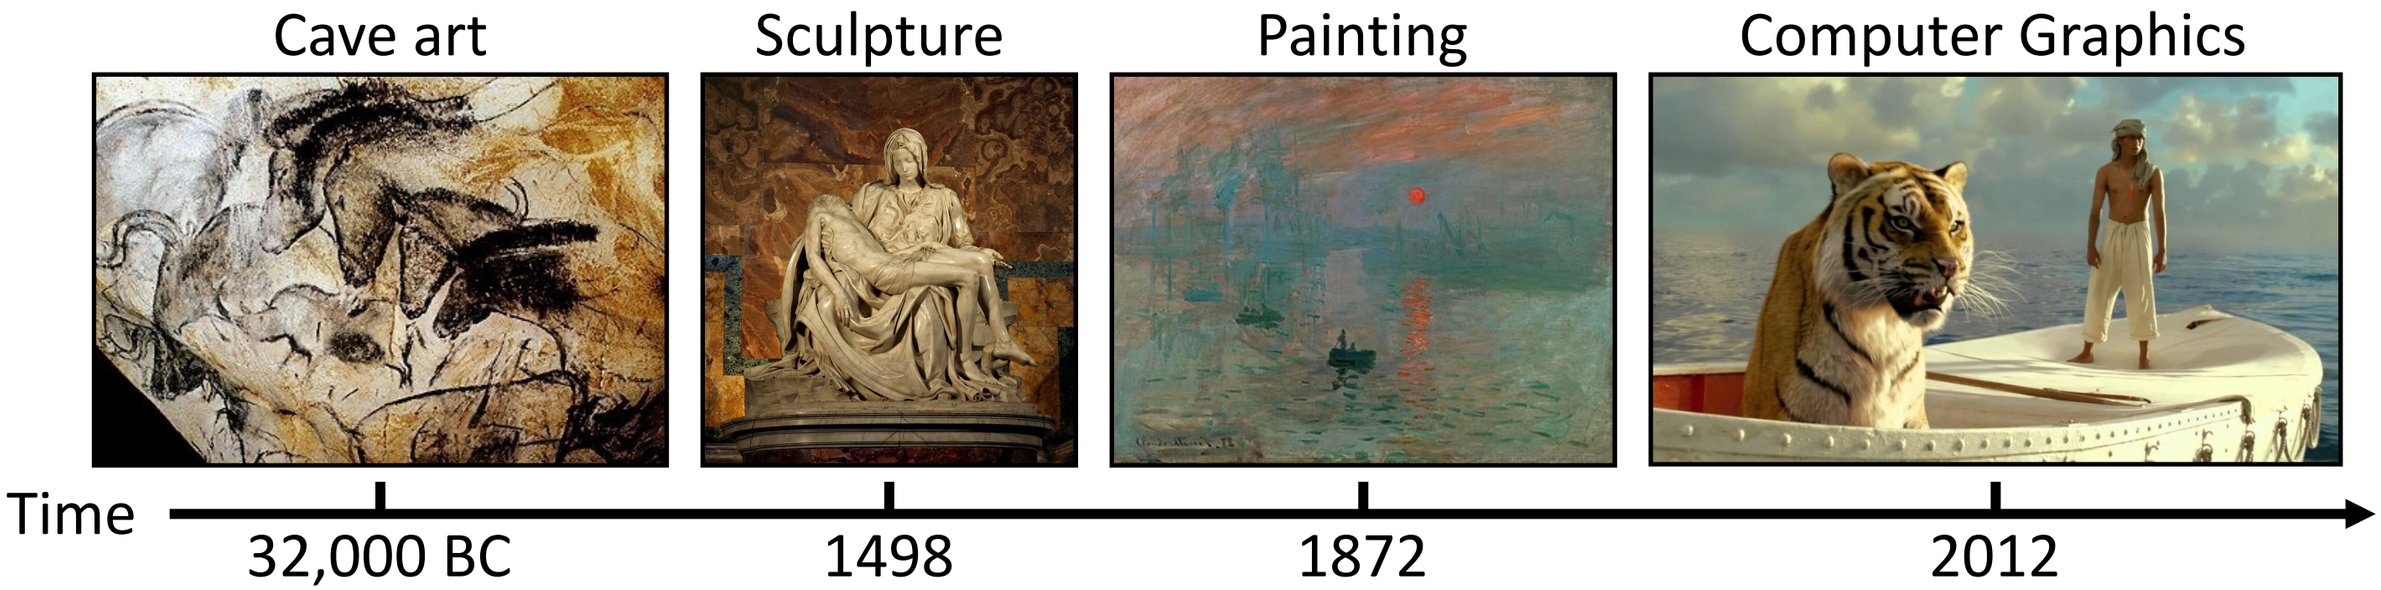

Who is creating visual content?

Cave art

Sculpture

Painting

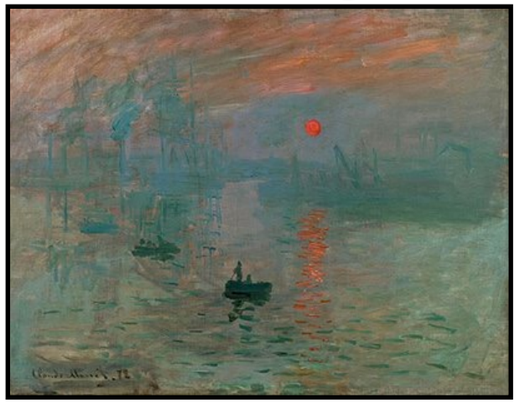

Computer Graphics

Time

32,000 BC — 1498 — 1872 — 2012

Michelangelo

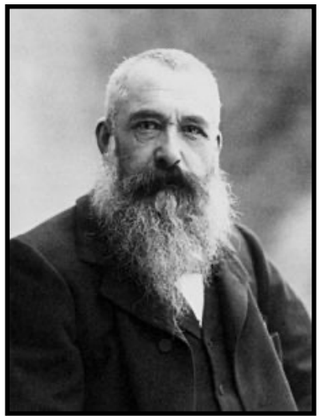

Claude Monet

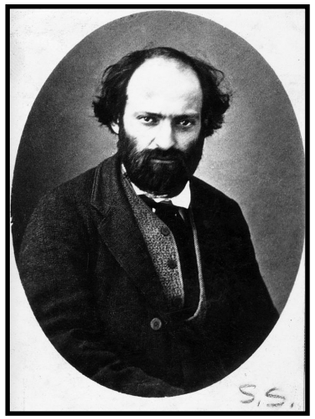

Paul Cezanne

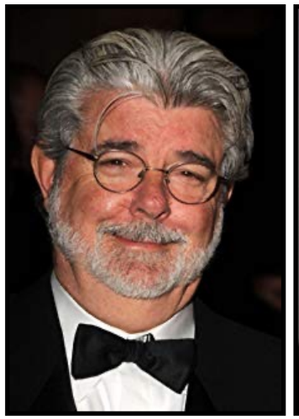

George Lucas

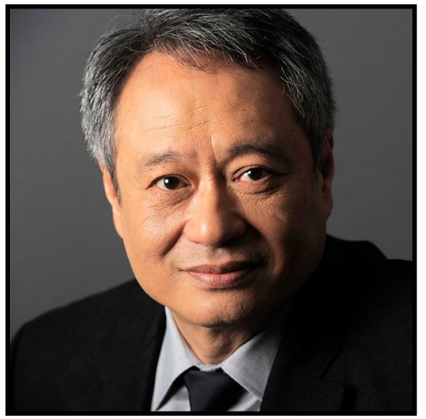

Ang Lee

Who is creating visual content?

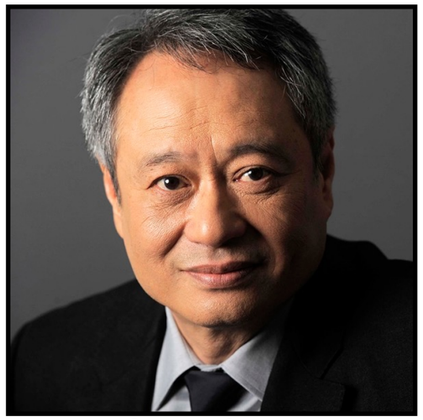

Ang Lee

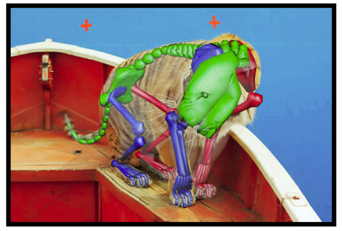

Skeleton

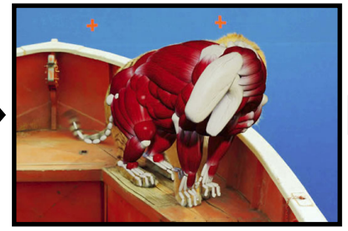

Geometry

Texture

Details

Image

Idea

Visual Content

Who is creating visual content?

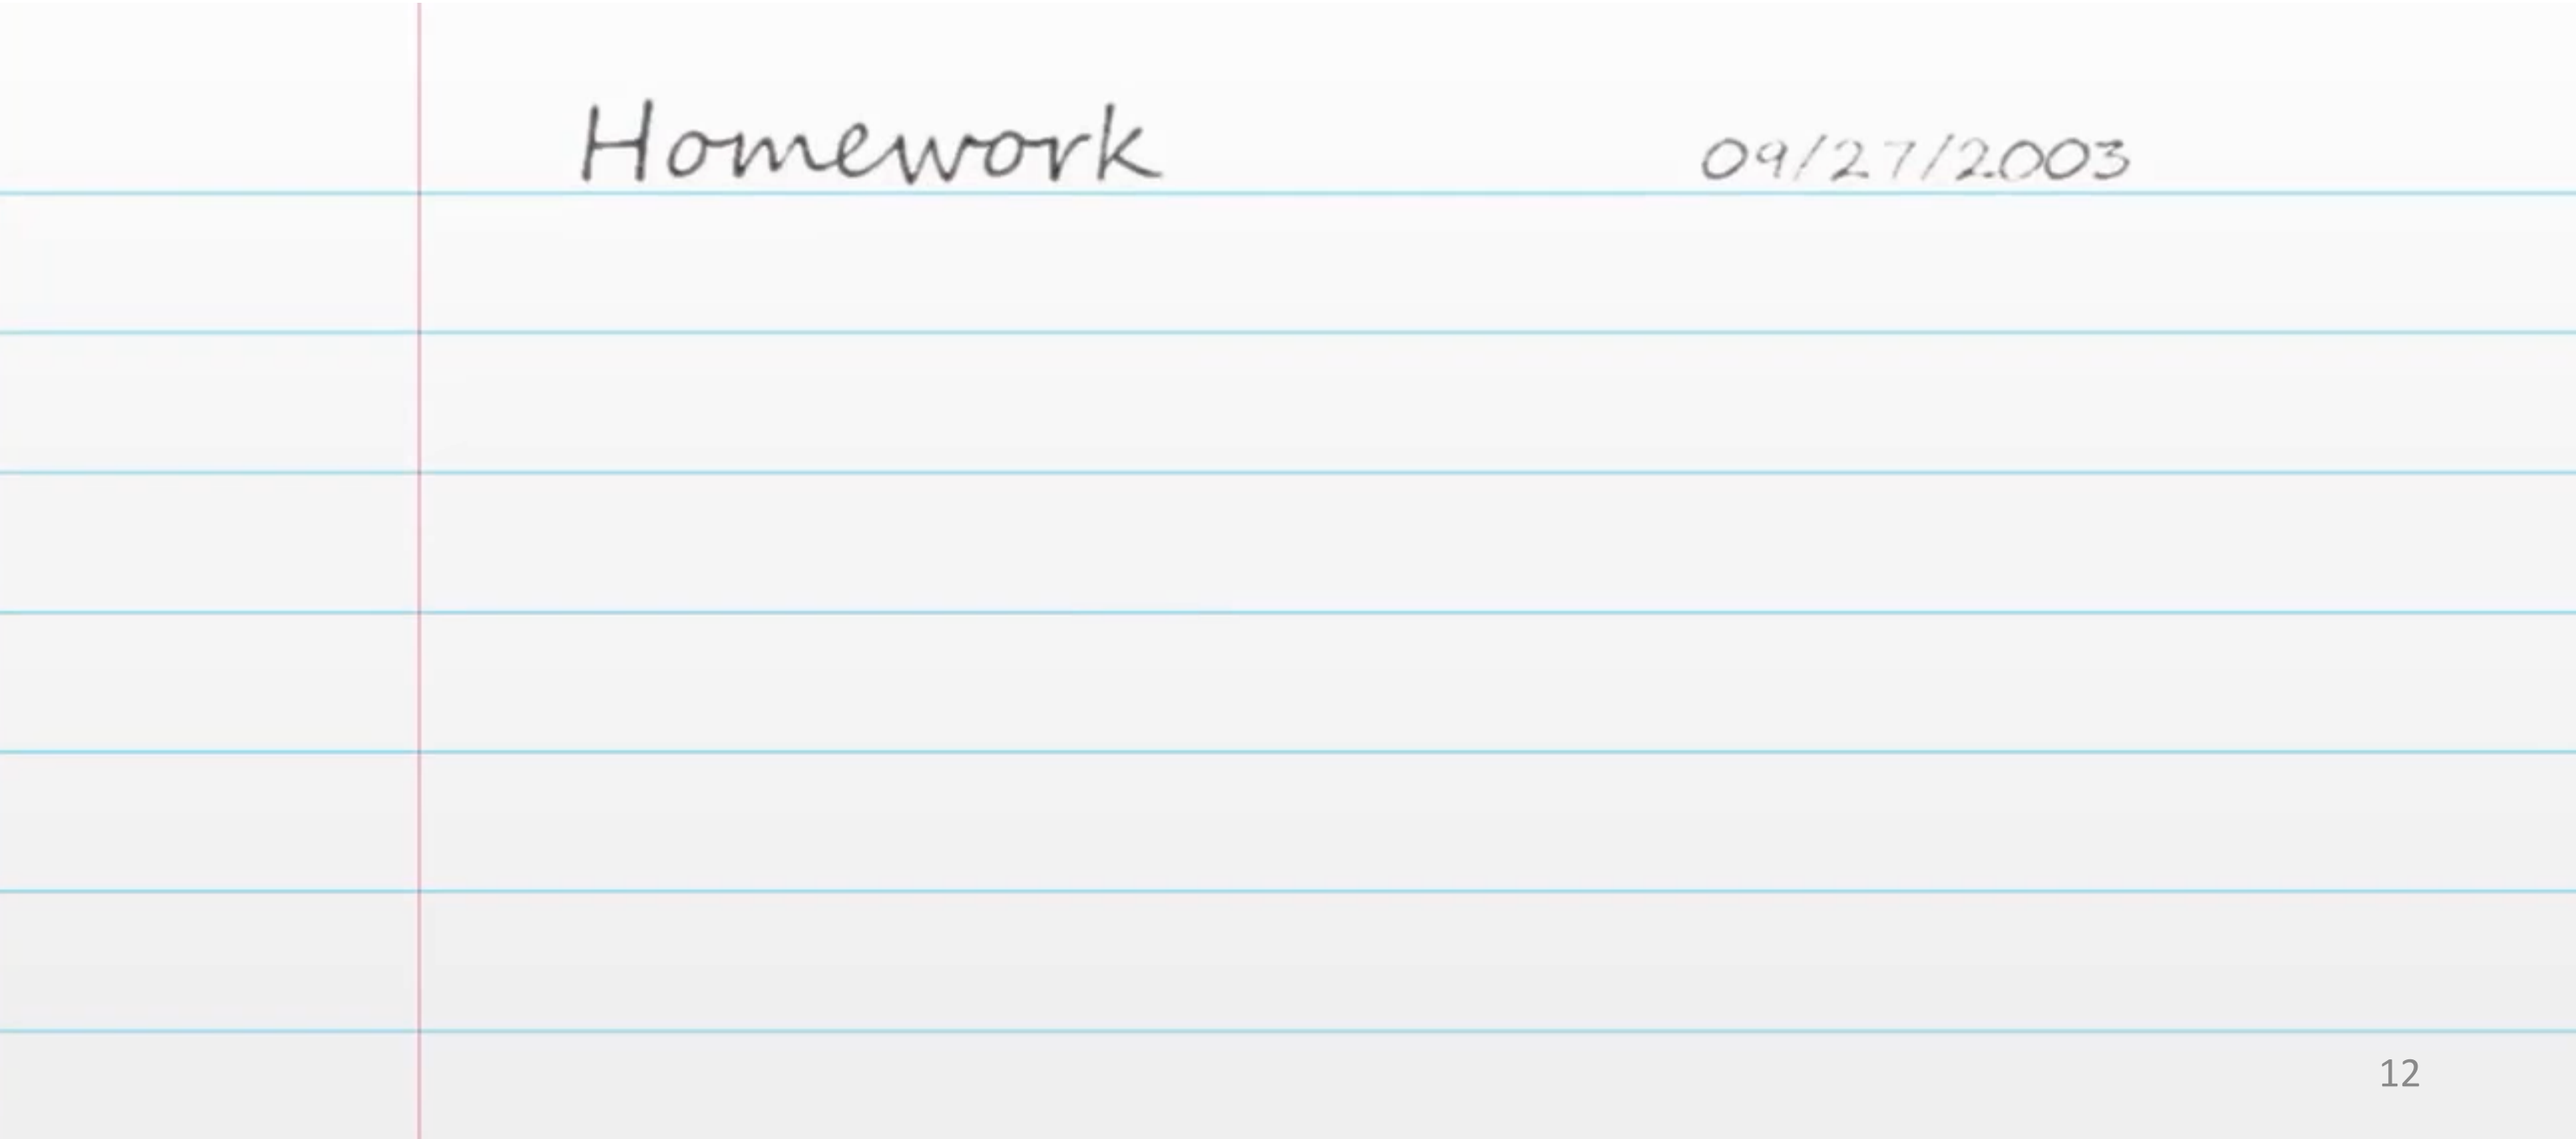

Who is creating visual content?

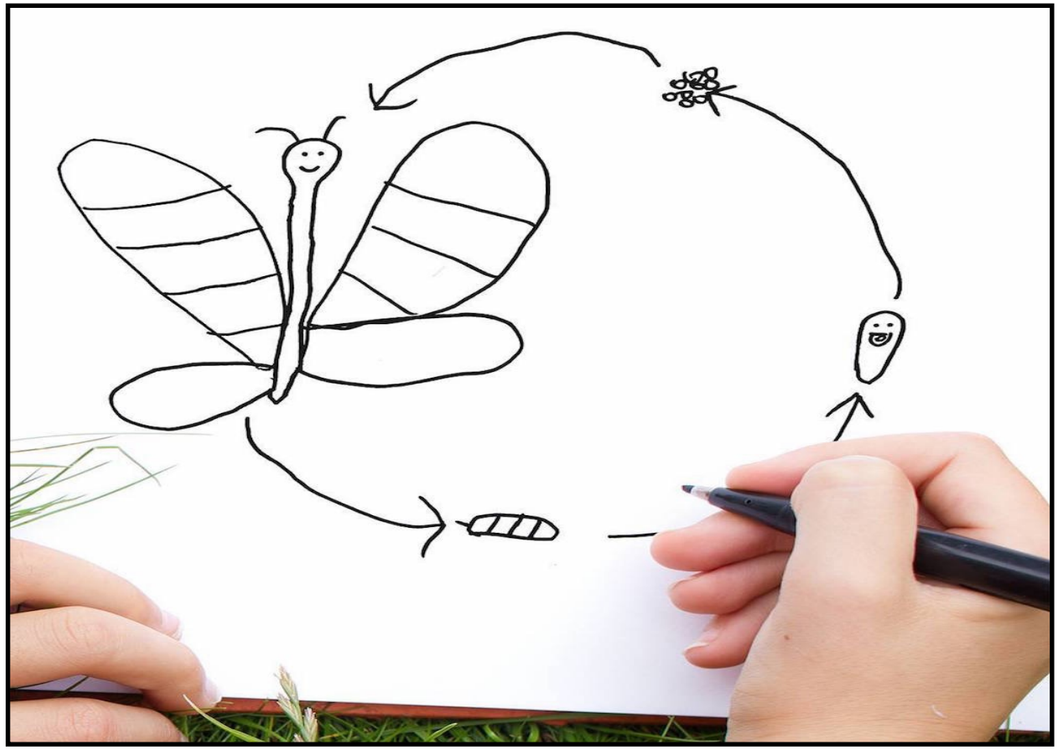

Kid’s drawing

Photoshop result by his father

Creating Visual Realism Manually

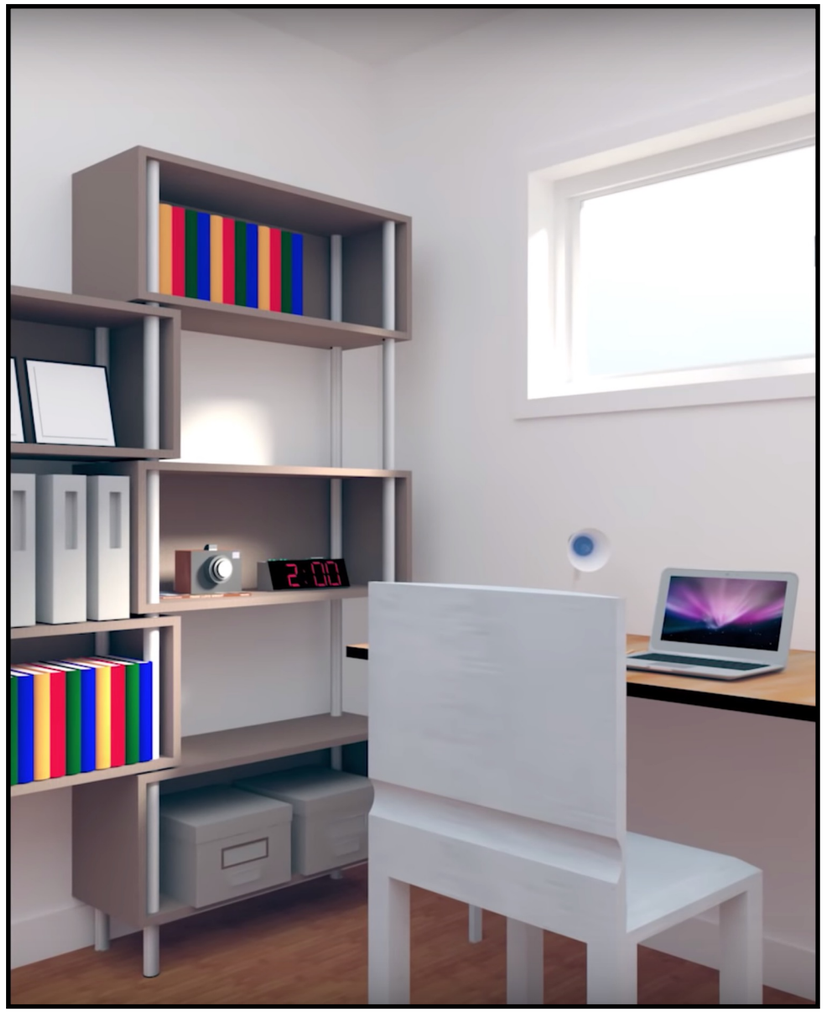

CG office

CG office (more details)

My advisor's office

Data-Driven Graphics (2000s)

Graphics → Image Retrieval

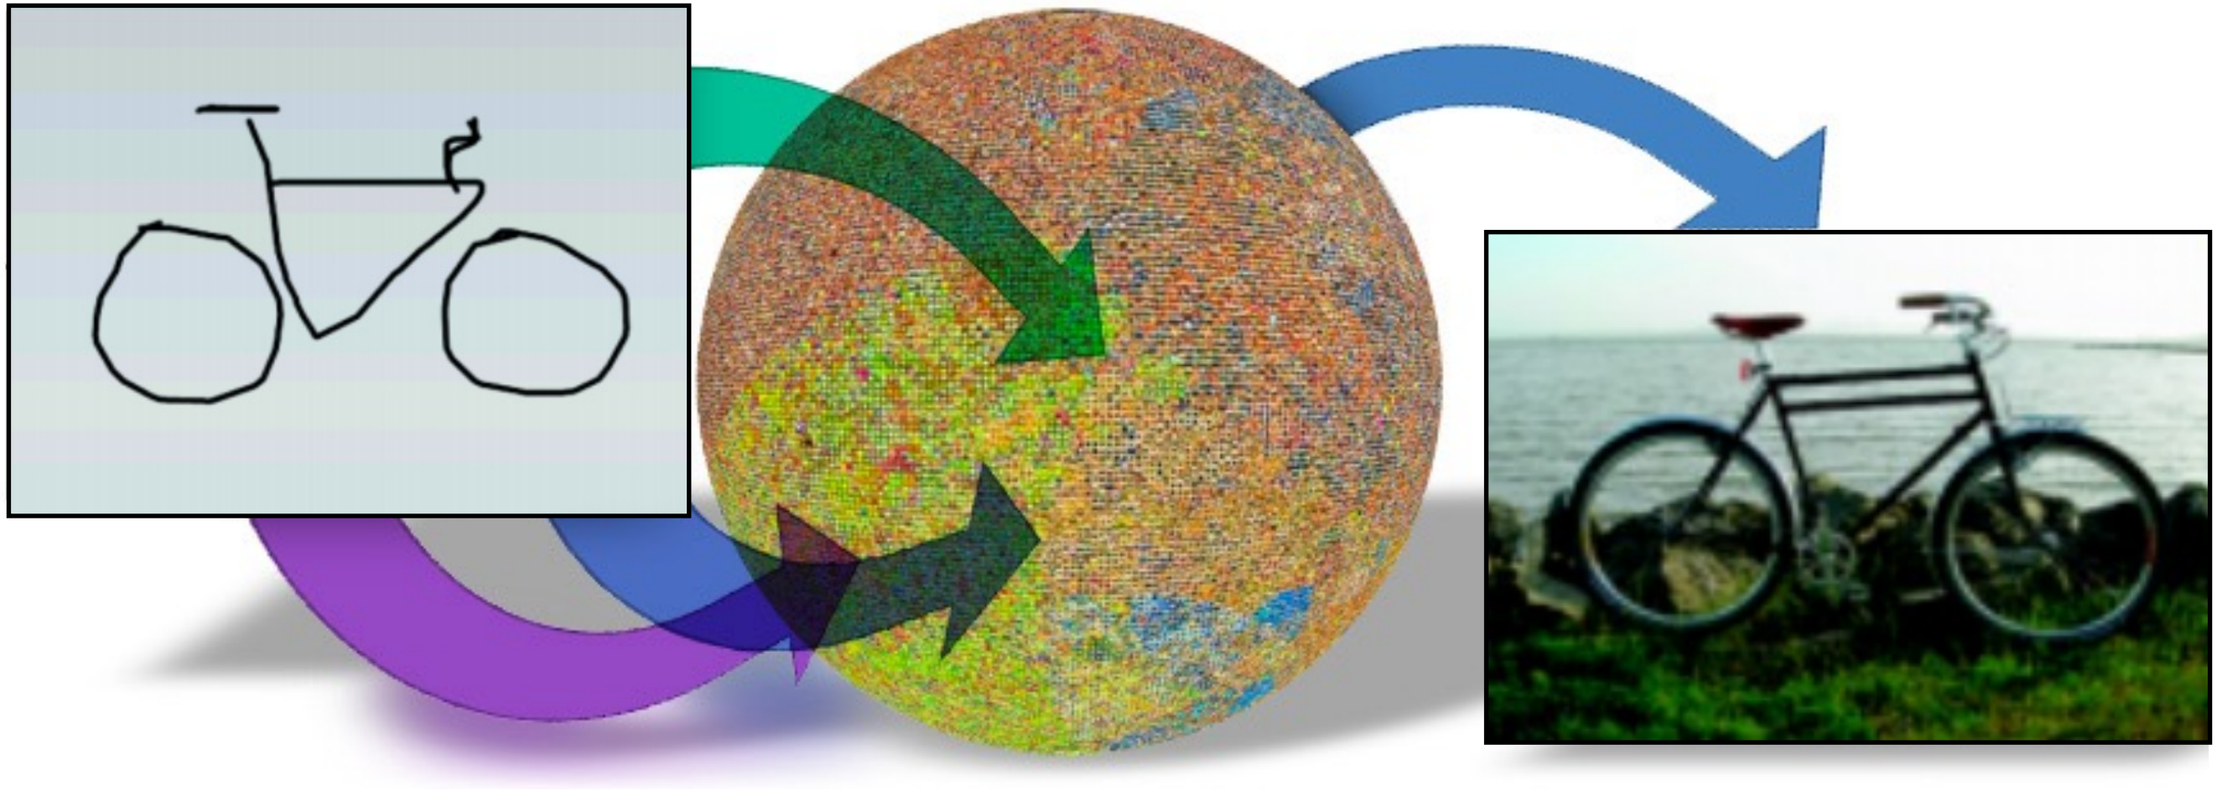

Picture from James Hays

Data-Driven Graphics (2000s)

Compositing multiple parts

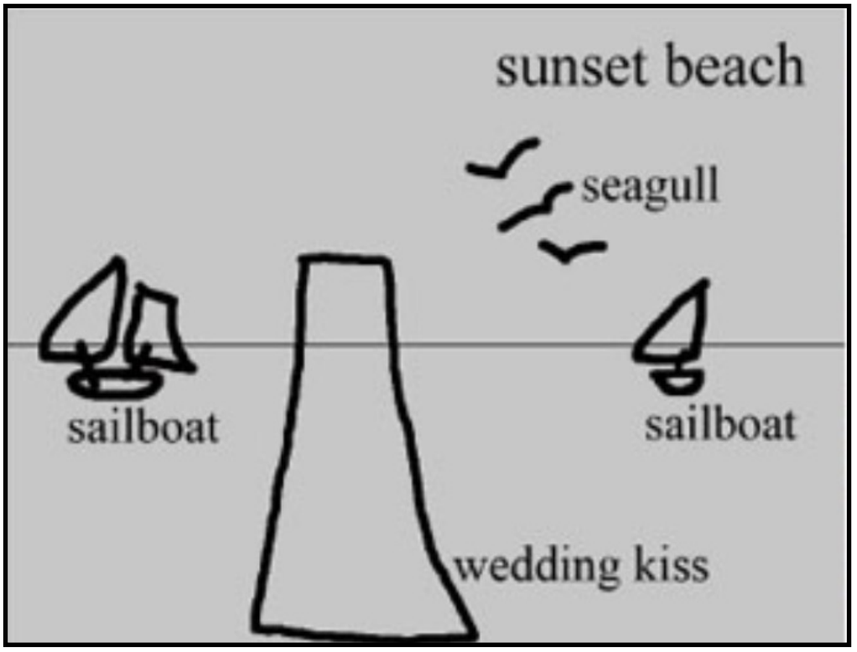

User Input

Database images

Output

Sketch2Photo [Tao et al. SIGGRAPH Asia 2009]

Data-Driven Graphics (2000s)

- Hard to combine pieces
- No understanding of visual realism

Help everyone
easily create visual content

Teach machines
how to create realistic content

Image-to-Image Translation with **pix2pix**

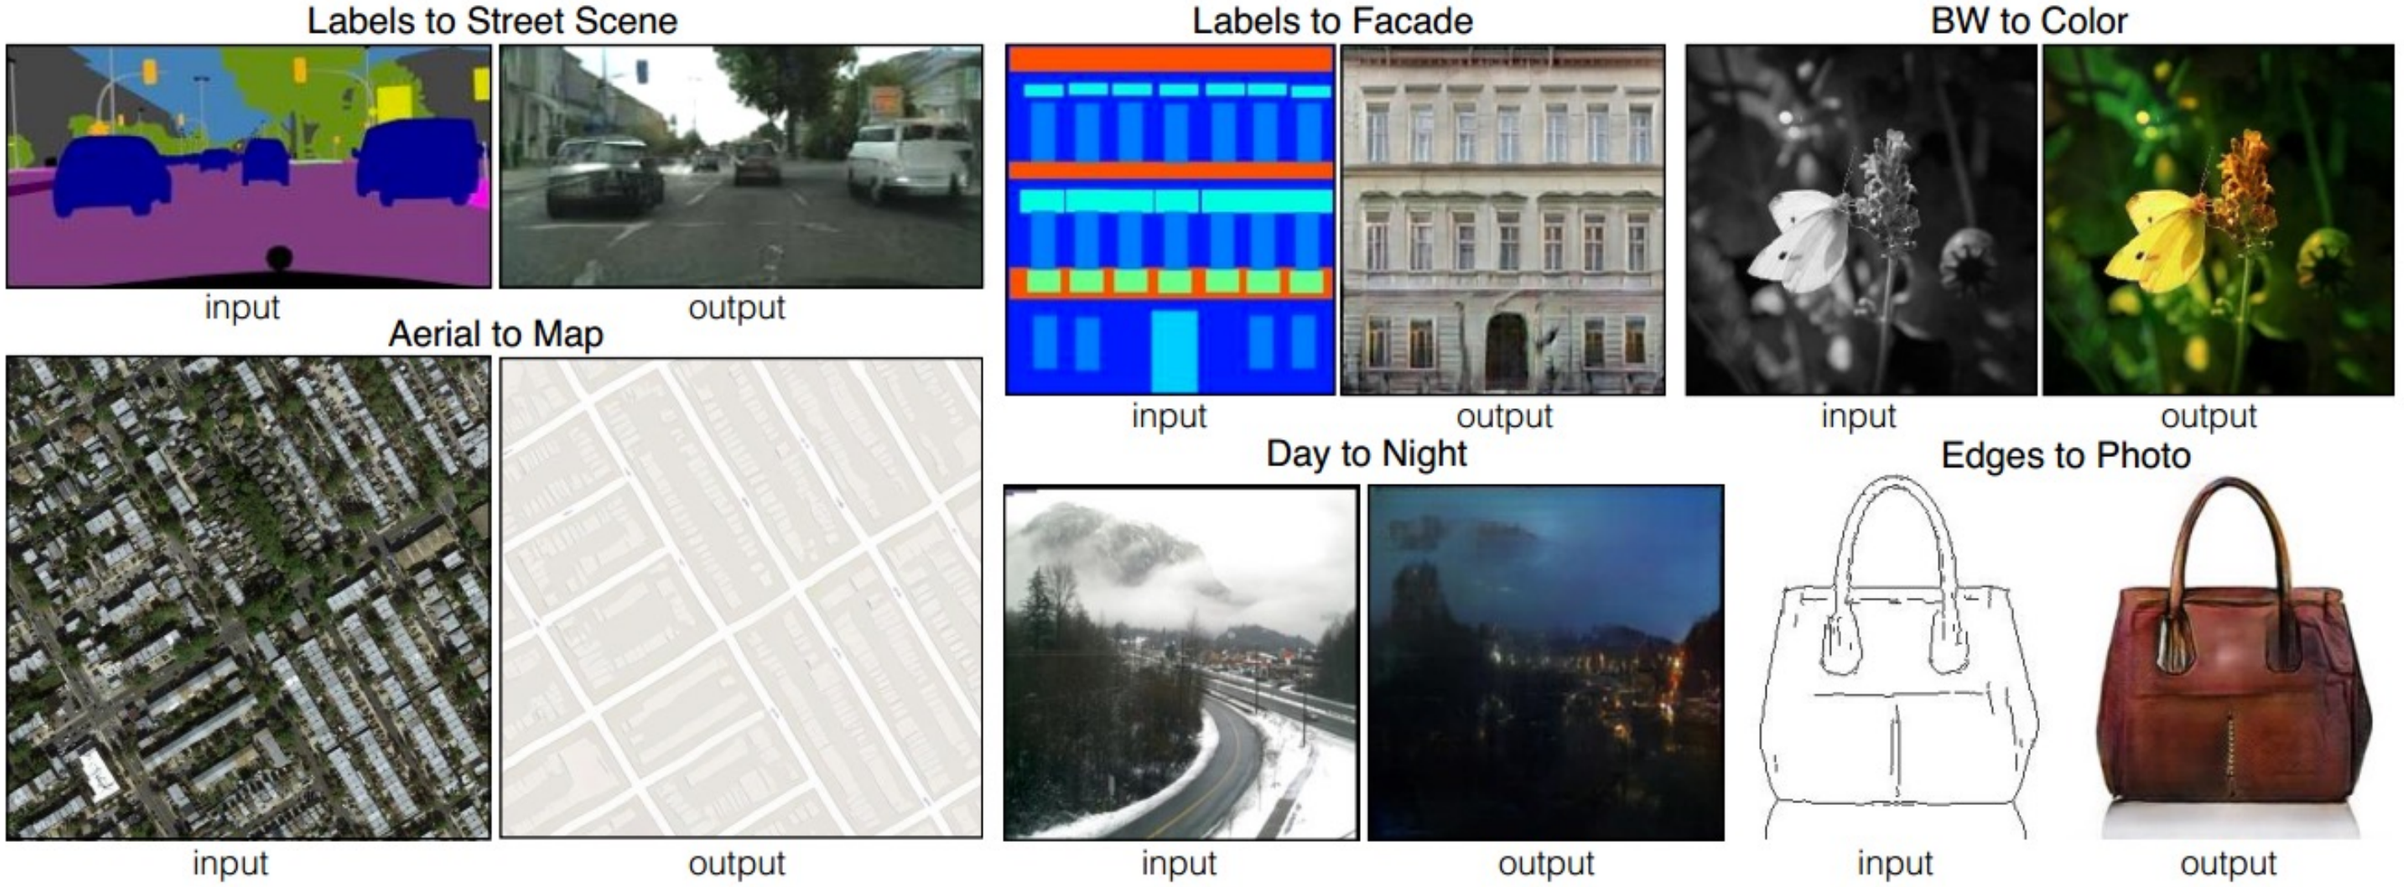

pix2pix [Isola, Zhu, Zhou, Efros. CVPR 2017]

#edges2cats with pix2pix

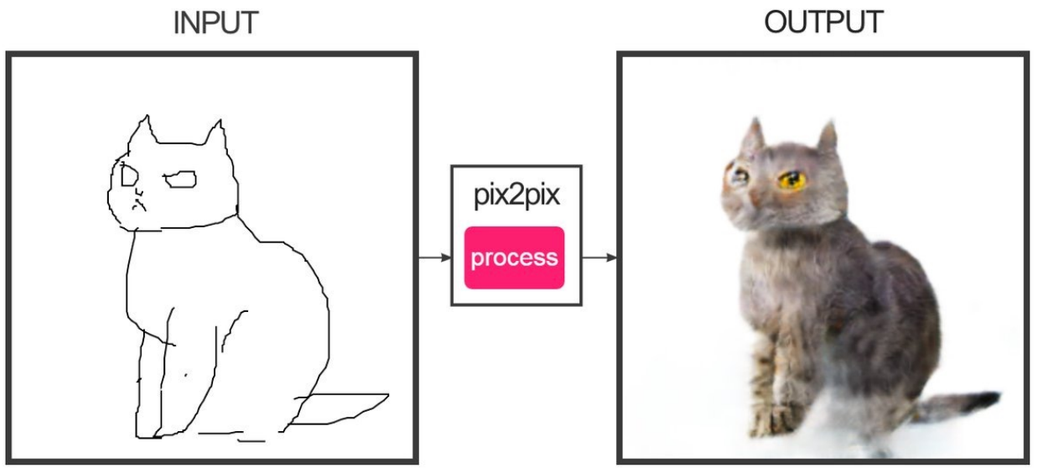

@gods_tail

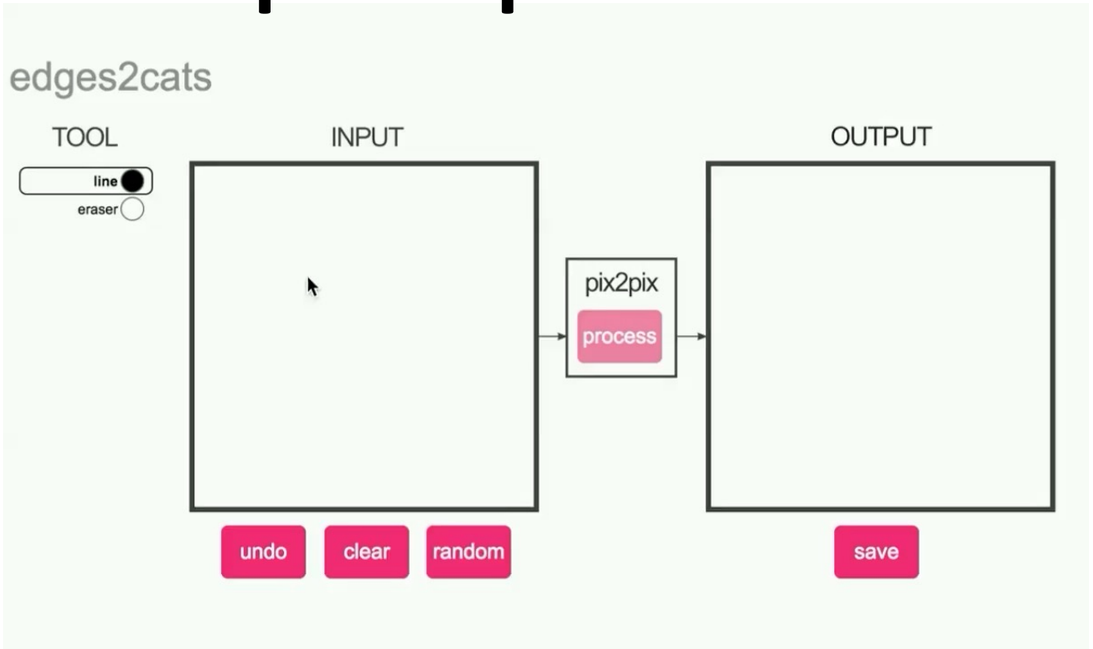

@matthematician

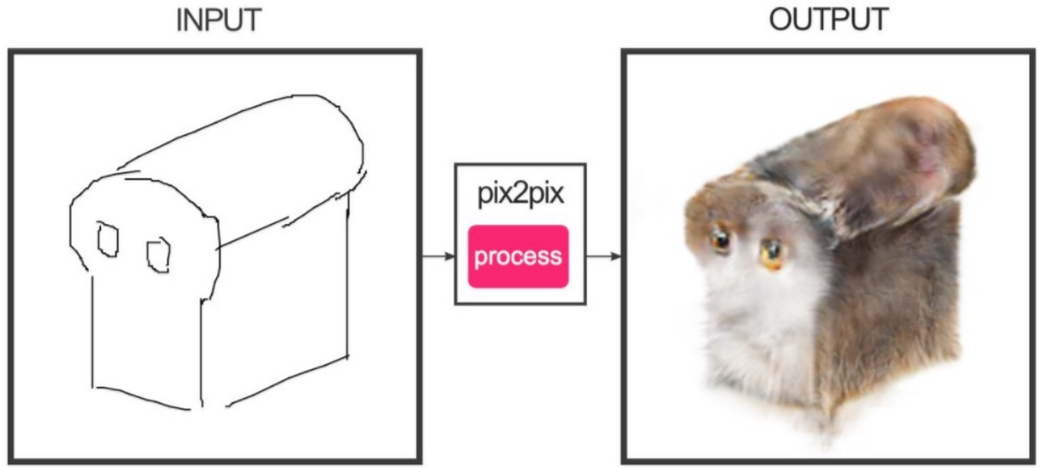

Ivy Tasi @ivymyt

Vitaly Vidmirov @vvid

By Christopher Hesse https://affinelayer.com/pixsrv/

GauGAN [Park, Liu, Wang, Zhu. 2019]

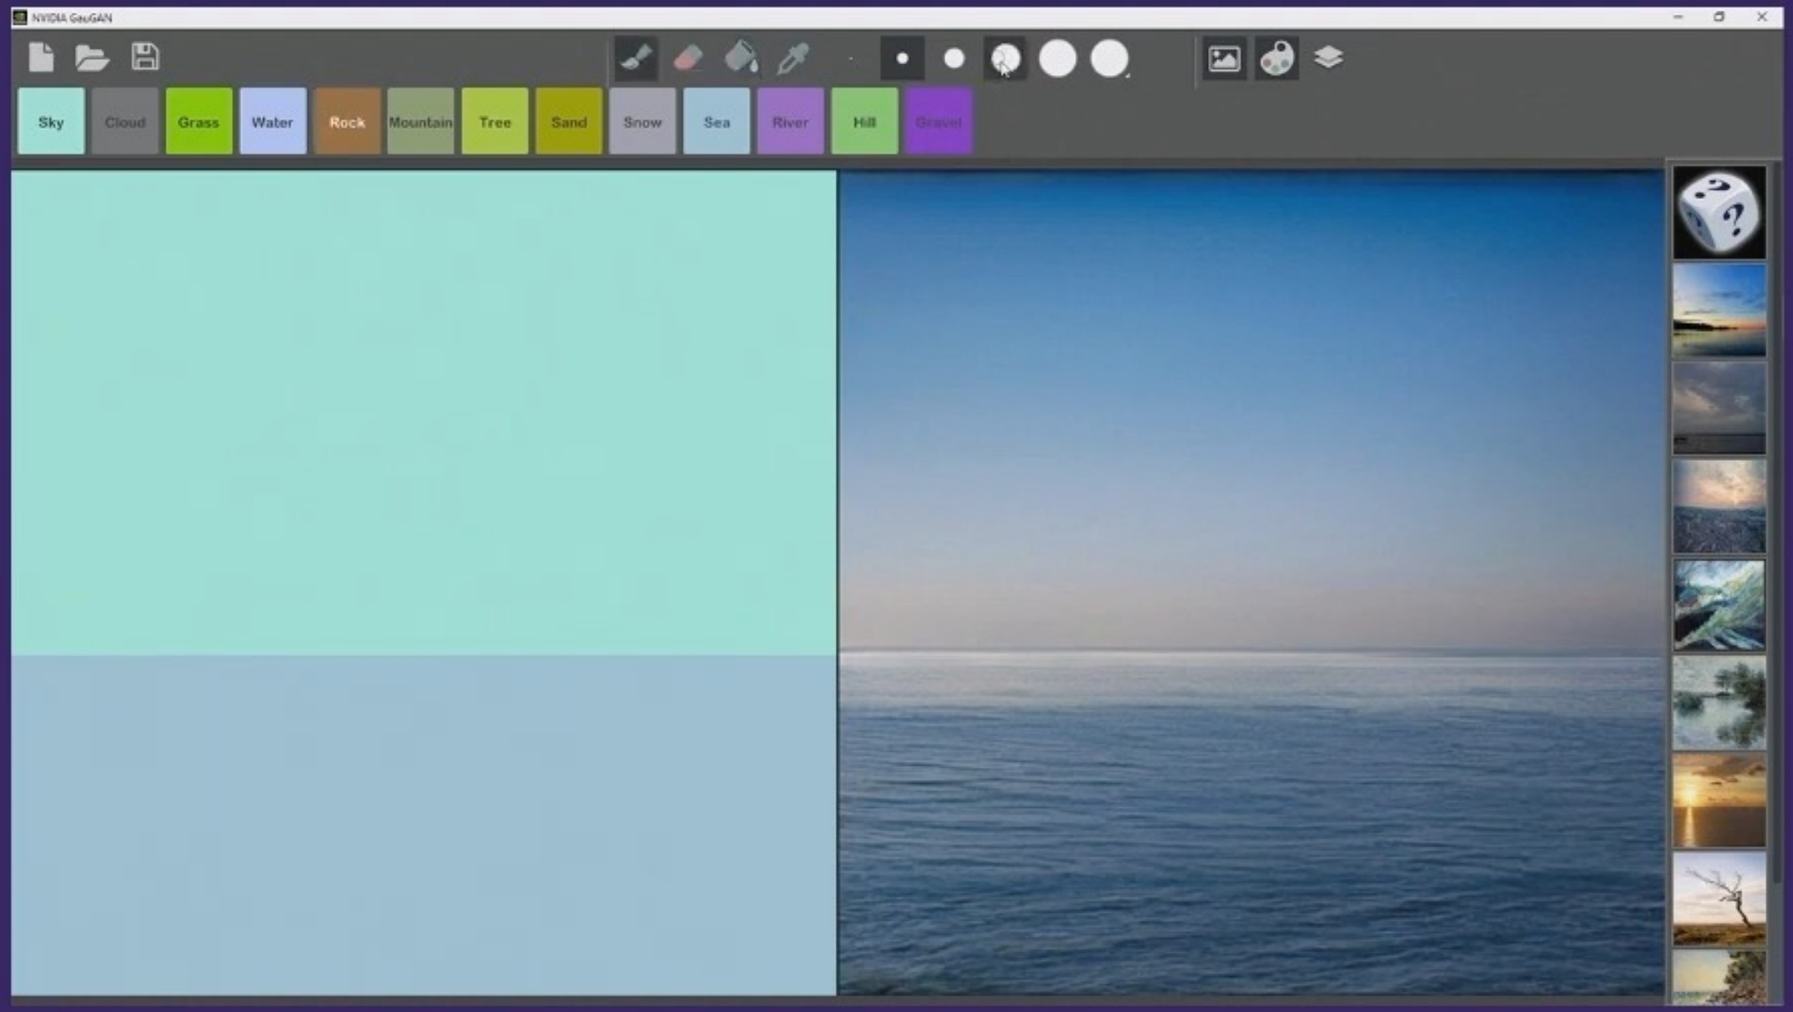

SIGGRAPH 2019 Real-time Live! “Best of Show Award" and "Audience Choice Award"

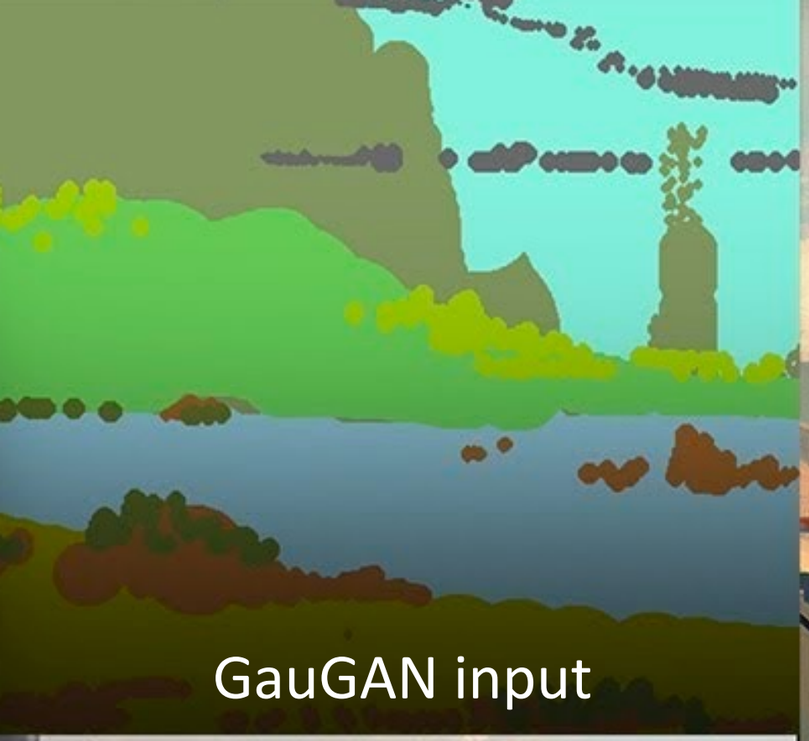

GauGAN input

GauGAN result

By Darek Zabrocki, Concept Designer and Illustrator

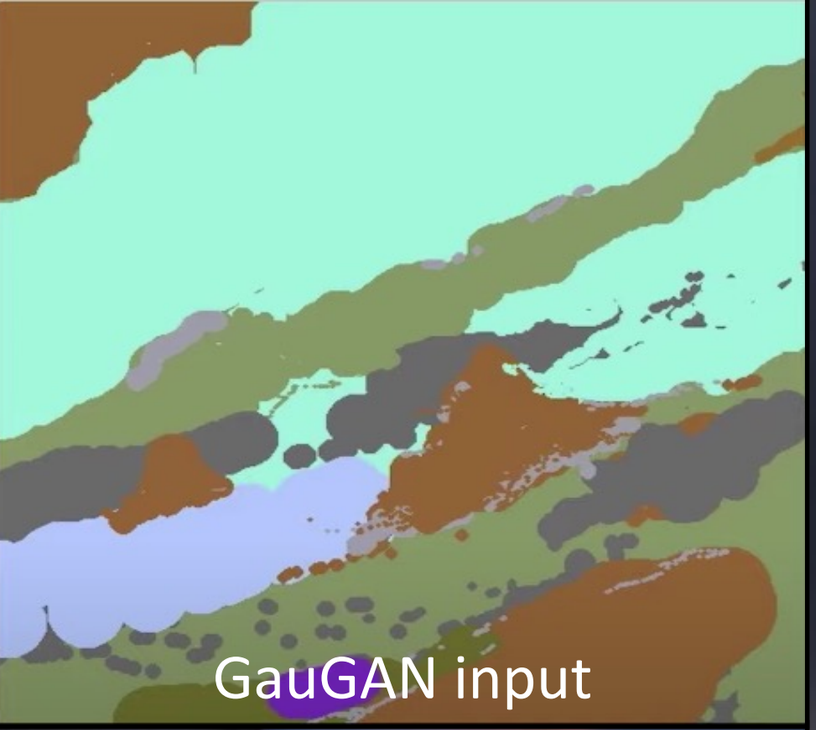

GauGAN input

GauGAN result

By Darek Zabrocki, Concept Designer and Illustrator

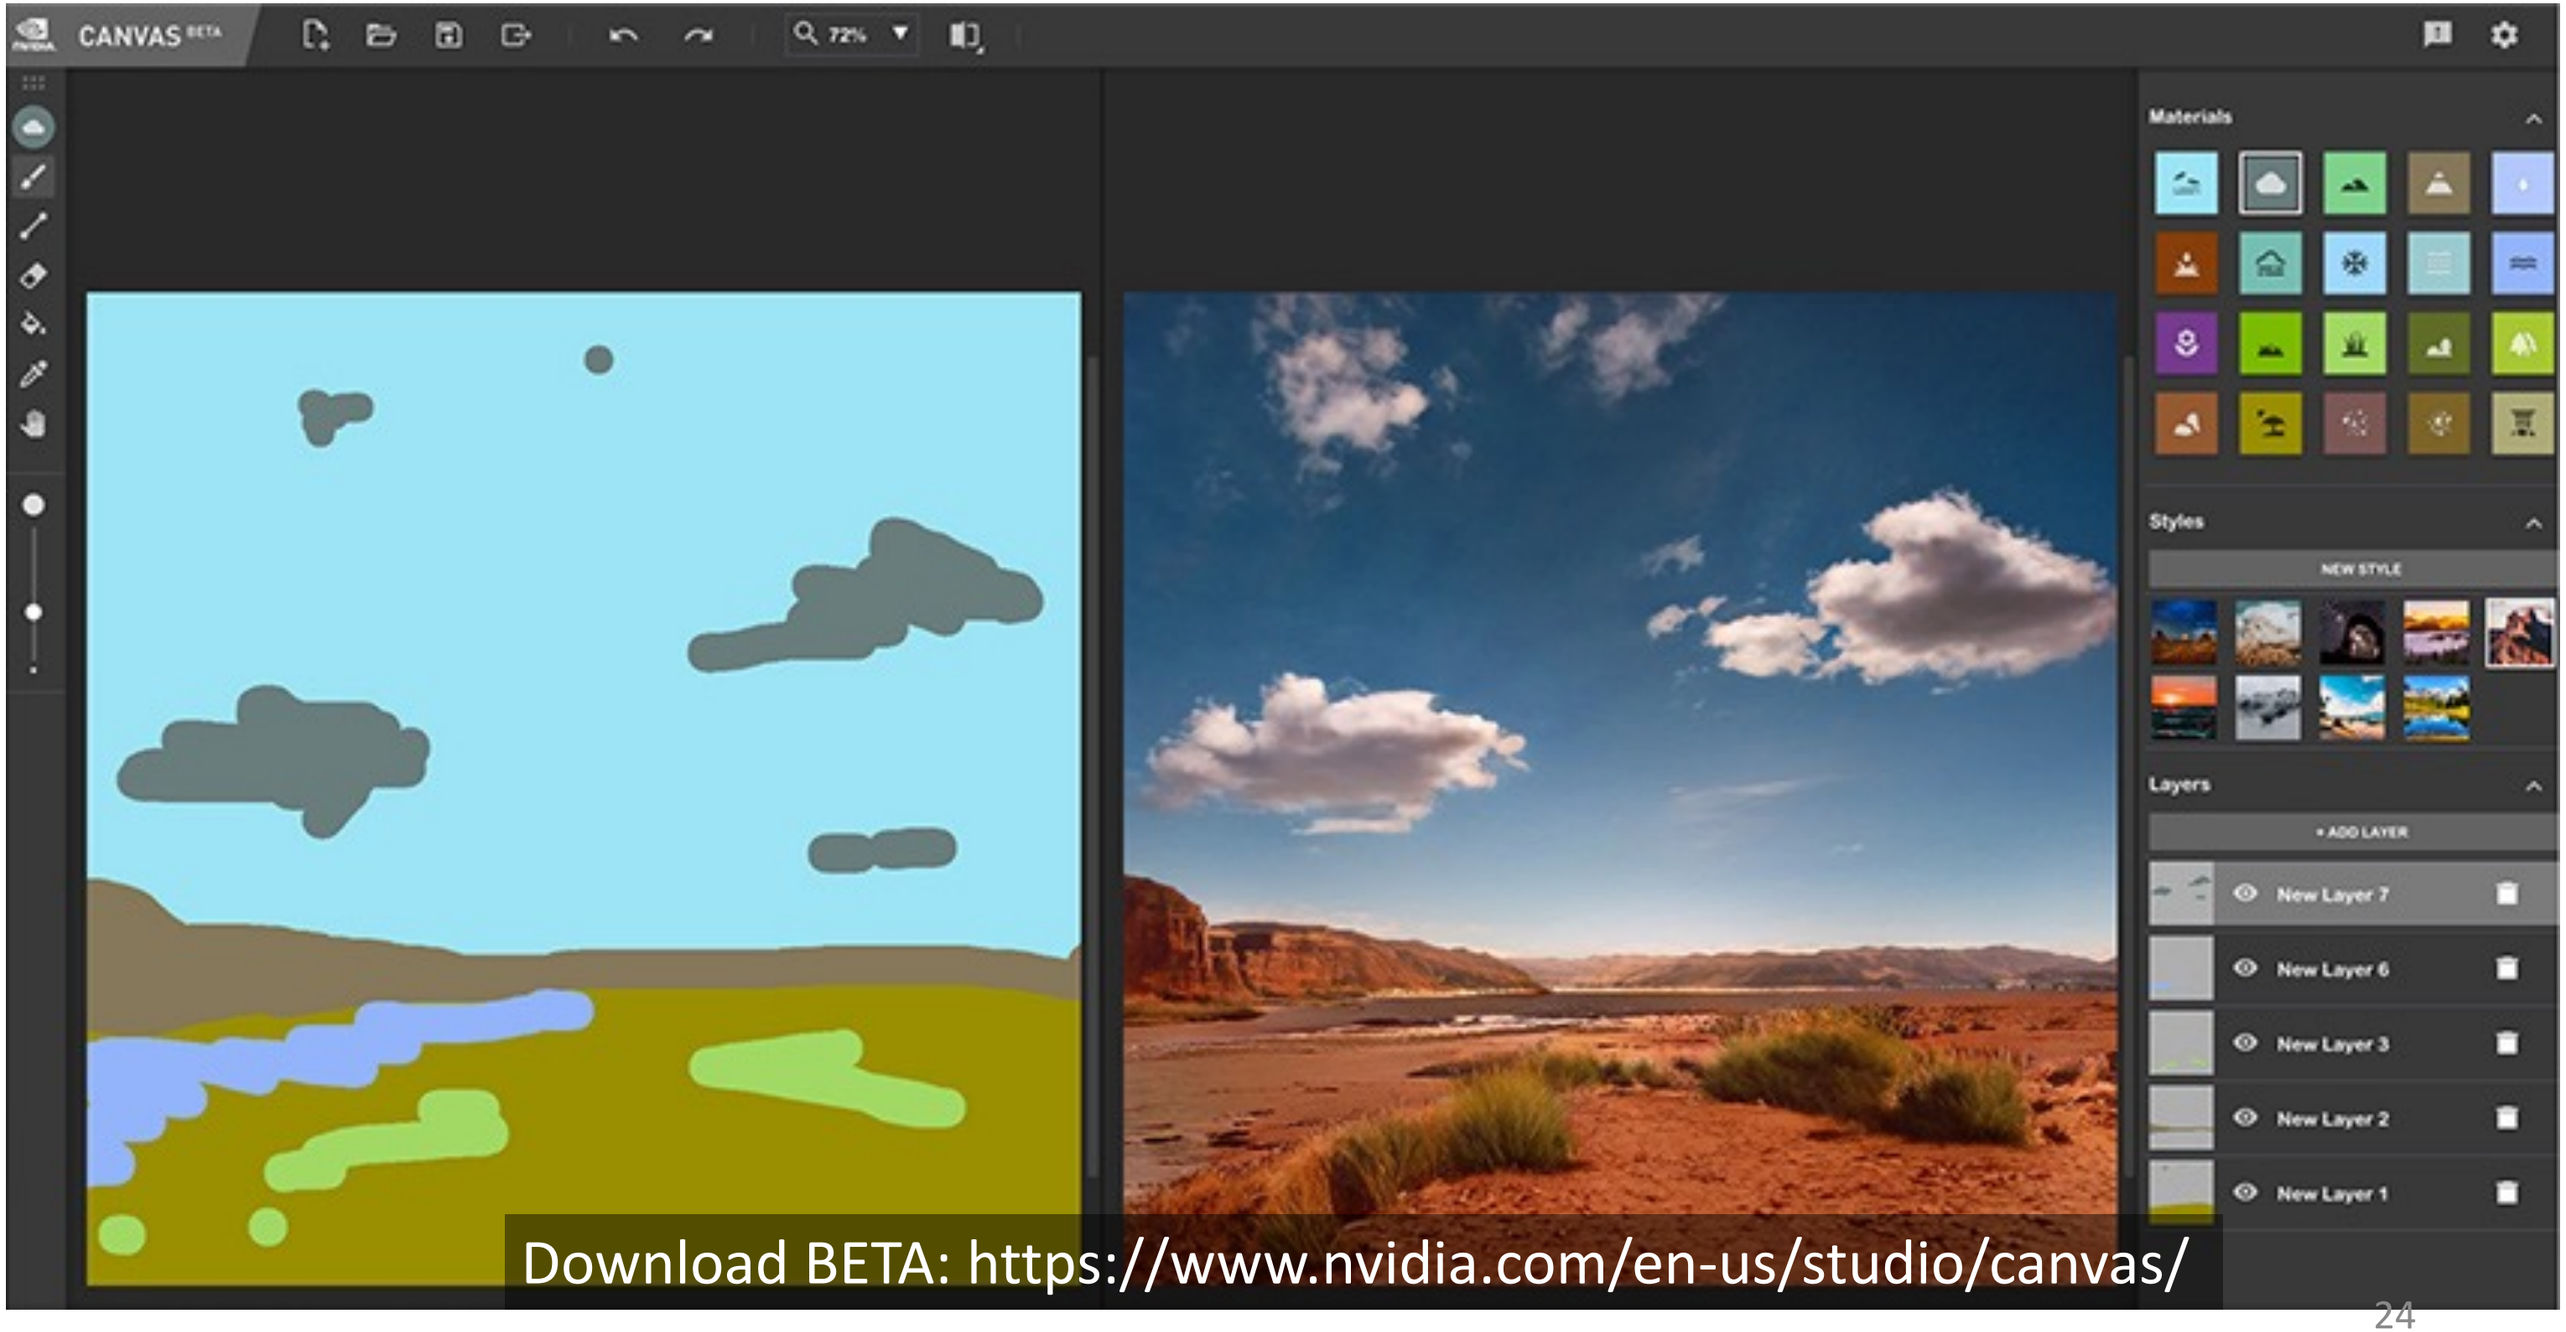
CANVAS BETA
72%
Materials
Styles
NEW STYLE
Layers
+ ADD LAYER
New Layer 7
New Layer 6
New Layer 3
New Layer 2
New Layer 1
Download BETA: https://www.nvidia.com/en-us/studio/canvas/

Collection Style Transfer

Photograph ©Alexei Efros

Monet

Van Gogh

Cezanne

Ukiyo-e

CycleGAN [Zhu, Park, Isola, Efros. 2017]

Monet’s paintings → photographic style

CycleGAN [Zhu, Park, Isola, Efros. 2017]

Horse → Zebra

CycleGAN [Zhu, Park, Isola, Efros. 2017]

Failure case

Failure case

Swapping Autoencoder [Park et al. NeurIPS 2020]

<not_ads>


Ps
Photoshop 22 Landscape Mixer

Photoshop 2021 Neural Filters

</not_ads>

Custom Stable Diffusion

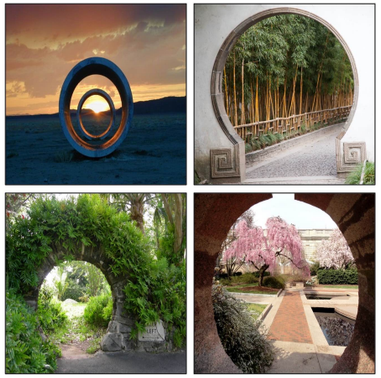

A photo of a **moongate**

A **moongate** in the snowy ice

A squirrel in front of **moongate**

A photo of a **V* dog**

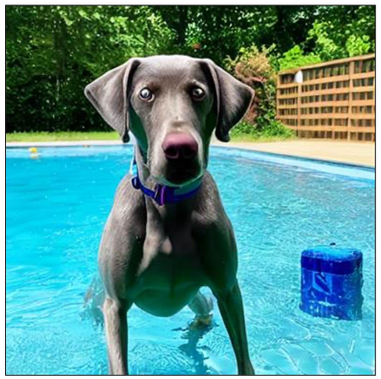

A **V* dog** in a swimming pool

A **V* dog** wearing sunglasses

User input images

Single-concept generation

A digital illustration of a **V* dog** in front of a **moongate**

V* dog wearing sunglasses in front of a **moongate**

Multi-concept composition

custom-diffusion [Nupur Kumari et al., 2022]

Research Highlights

from other universities & industry labs

Synthesizing High-res Portraits

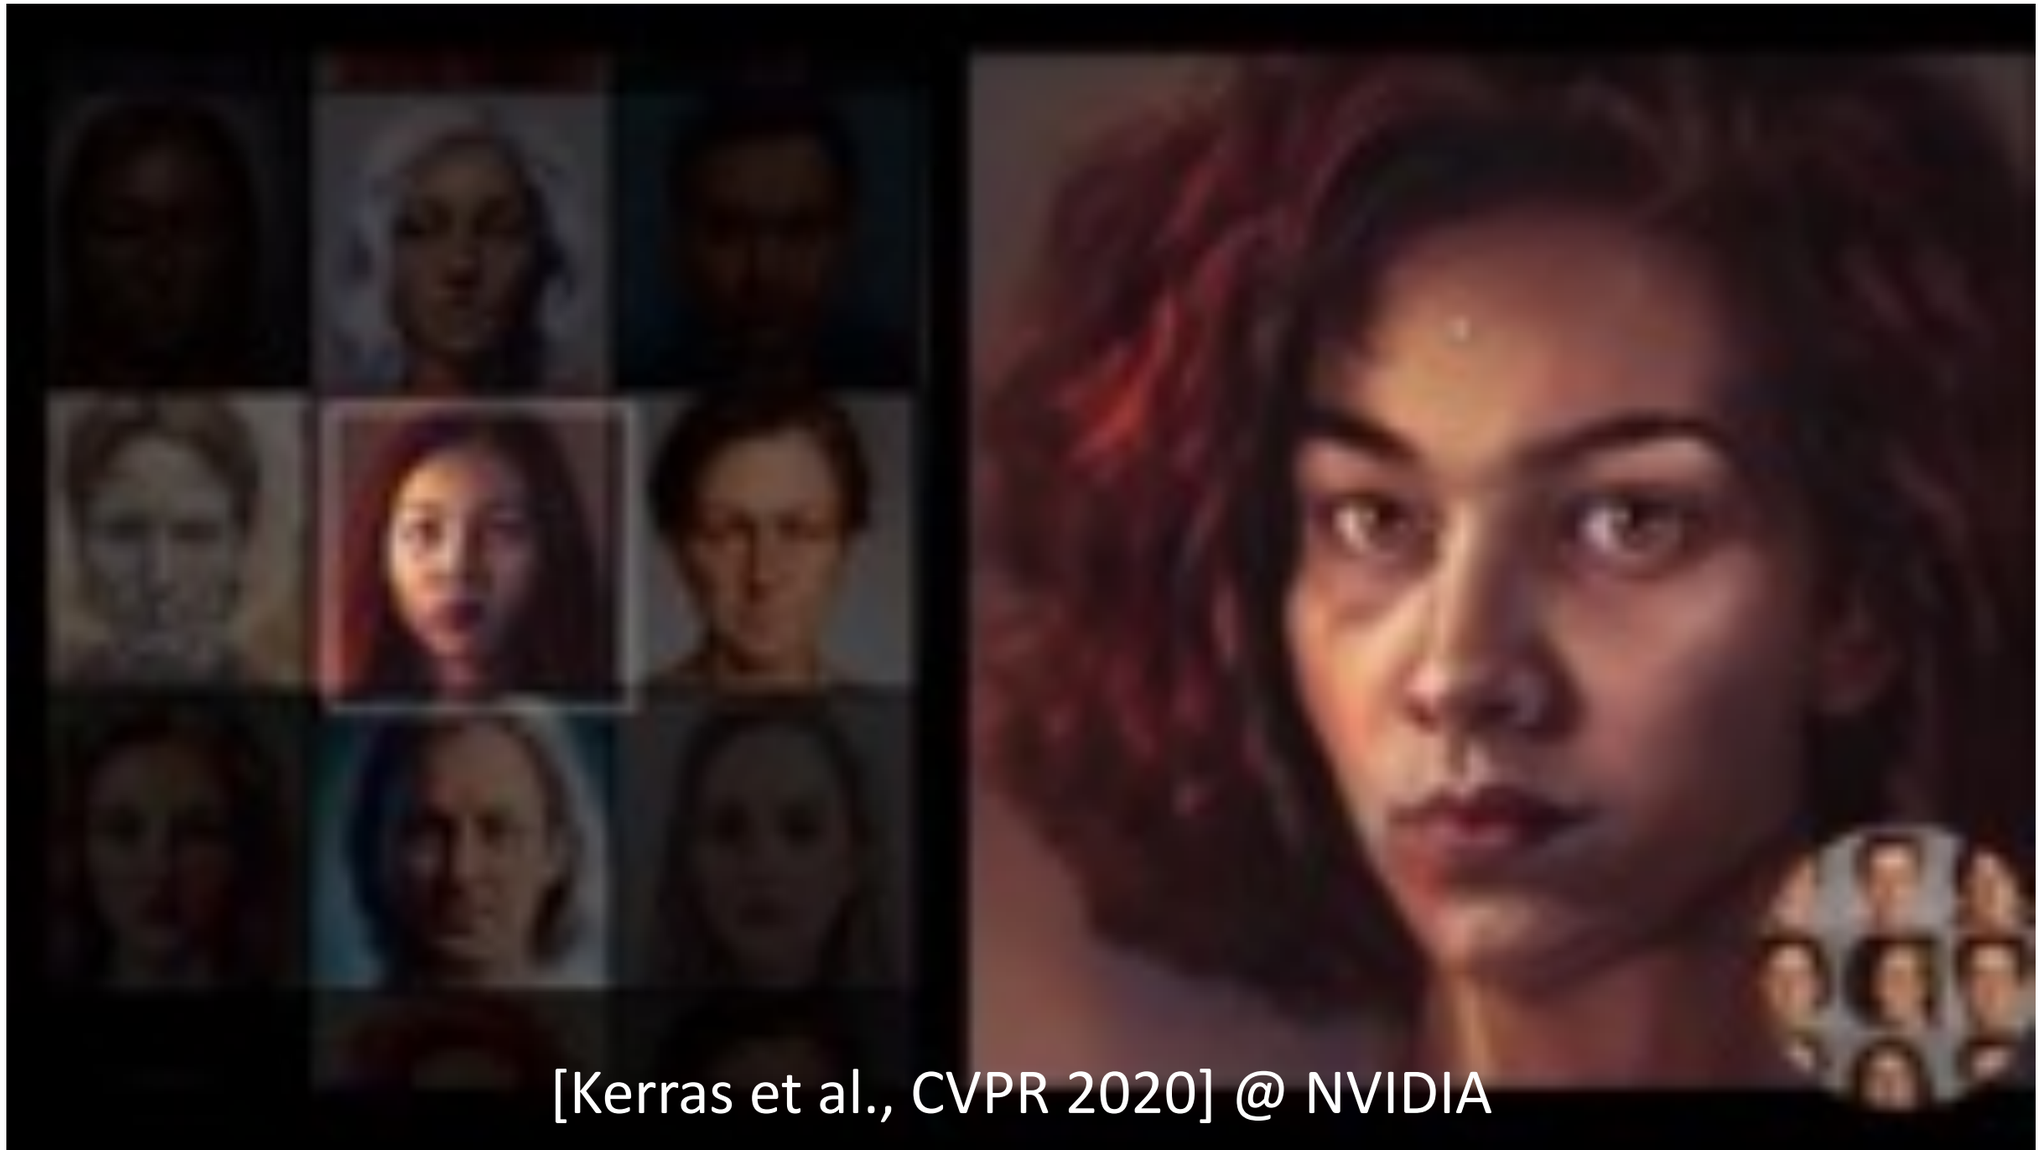

[Kerras et al., CVPR 2020] @ NVIDIA

Everybody Dances Now

Caroline Chan , Shiry Ginosar, Tinghui Zhou, Alexei A. Efros, ICCV 2019. @ UC Berkeley

Neural Talking-Head Synthesis

face-vid2vid: One-Shot Free-View Neural Talking-Head Synthesis for Video Conferencing
Ting-Chun Wang, Arun Mallya, Ming-Yu Liu. CVPR 2021 @ NVIDIA

NeRF in the Wild

[Martin-Brualla et al., CVPR 2021] @ Google Research

Text-based Image Editing

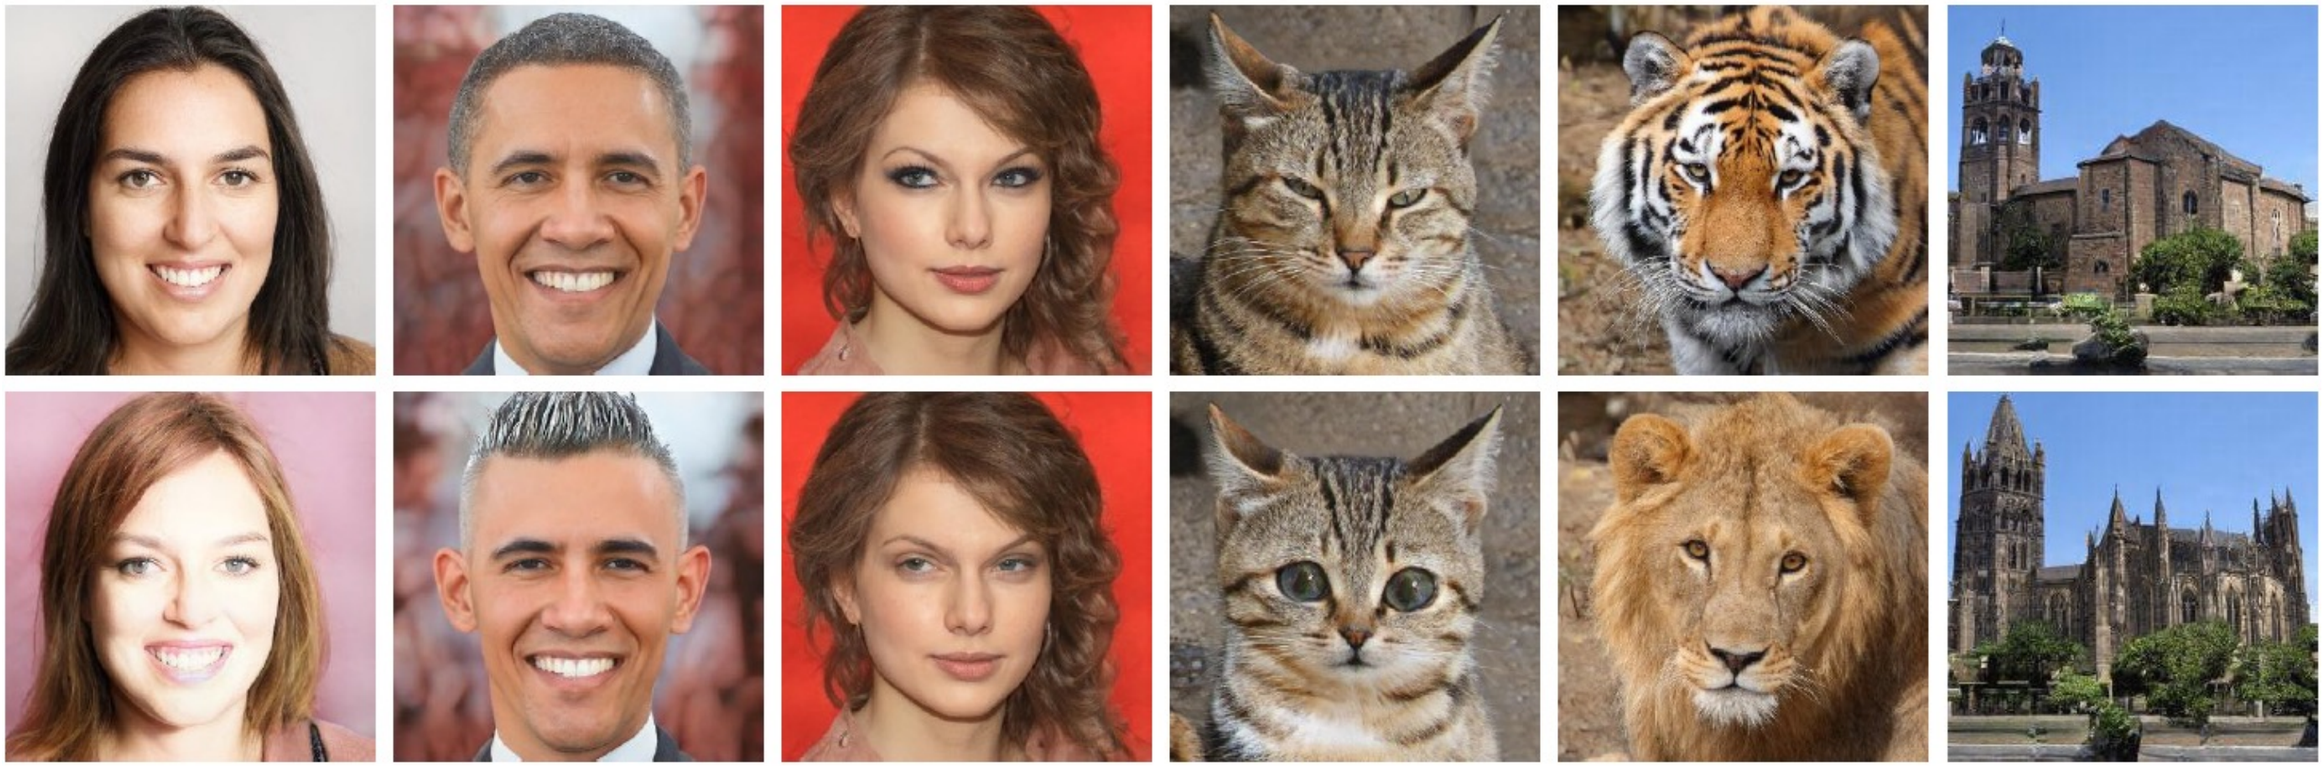

StyleCLIP [Or Patashnik*, Zongze Wu*, et al., ICCV 2021]

Generative Models (2016)

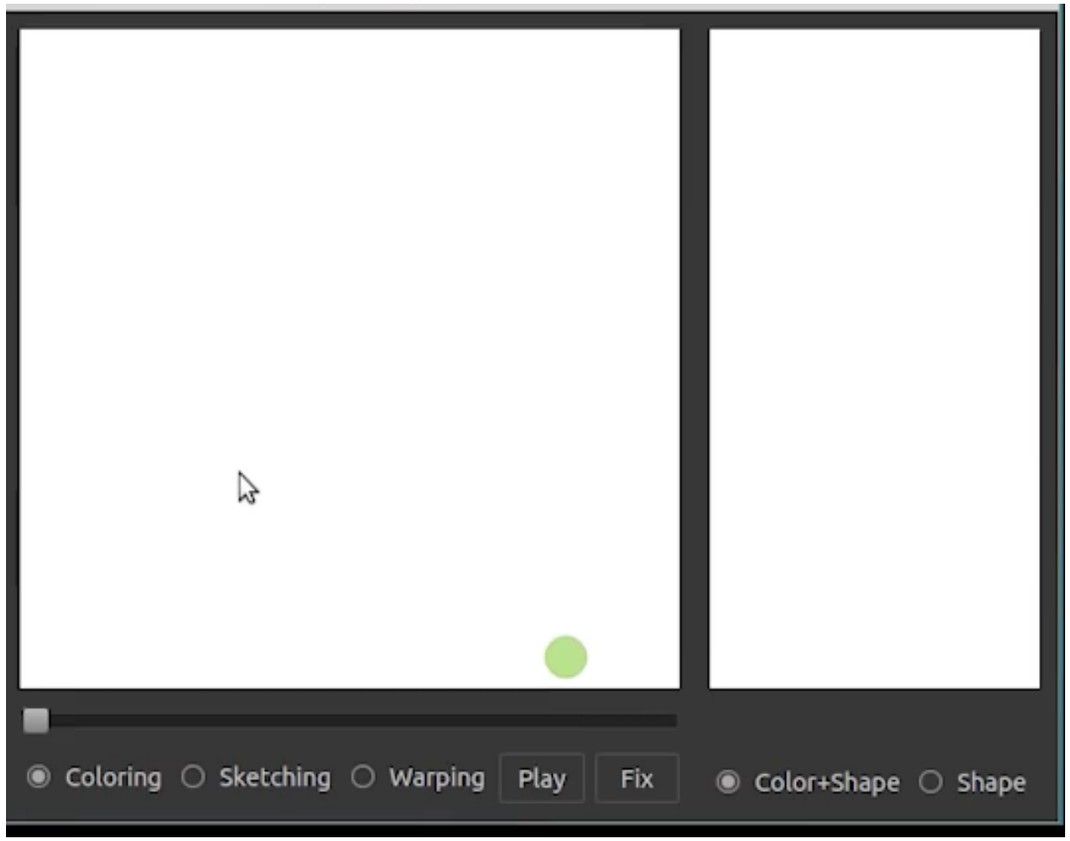

iGAN [Zhu et al., ECCV 2016]

A toilet seat sits open in the grass field.

Text2image [Mansimov et al., ICLR 2016]

Generative Models (2023)

Diffusion models
(DALL-E 2, Imagen, SD)

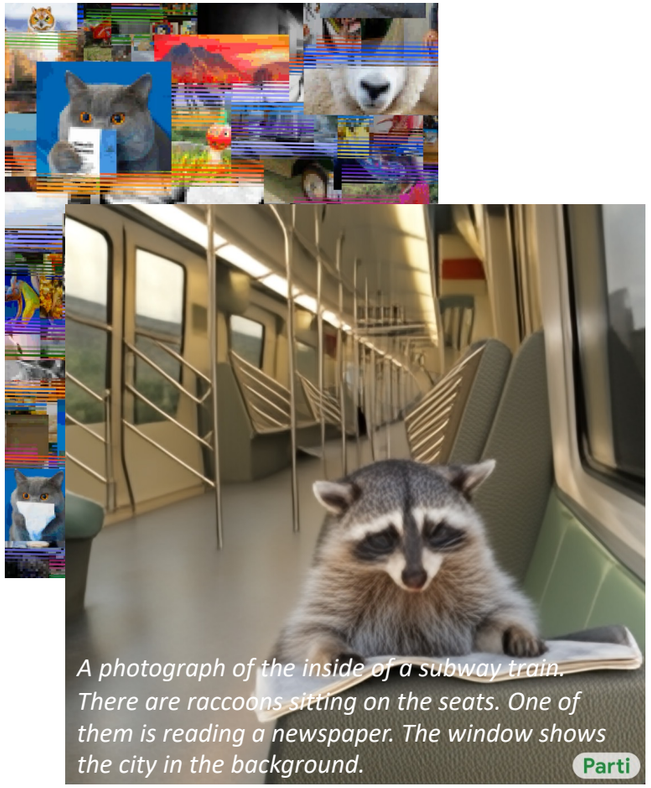

Autoregressive models
(Image GPT, Parti)

GANs, Masked GIT
(GigaGAN, MUSE)

Generative Models (2023)

By DALL·E 3

Generative Models (2023)

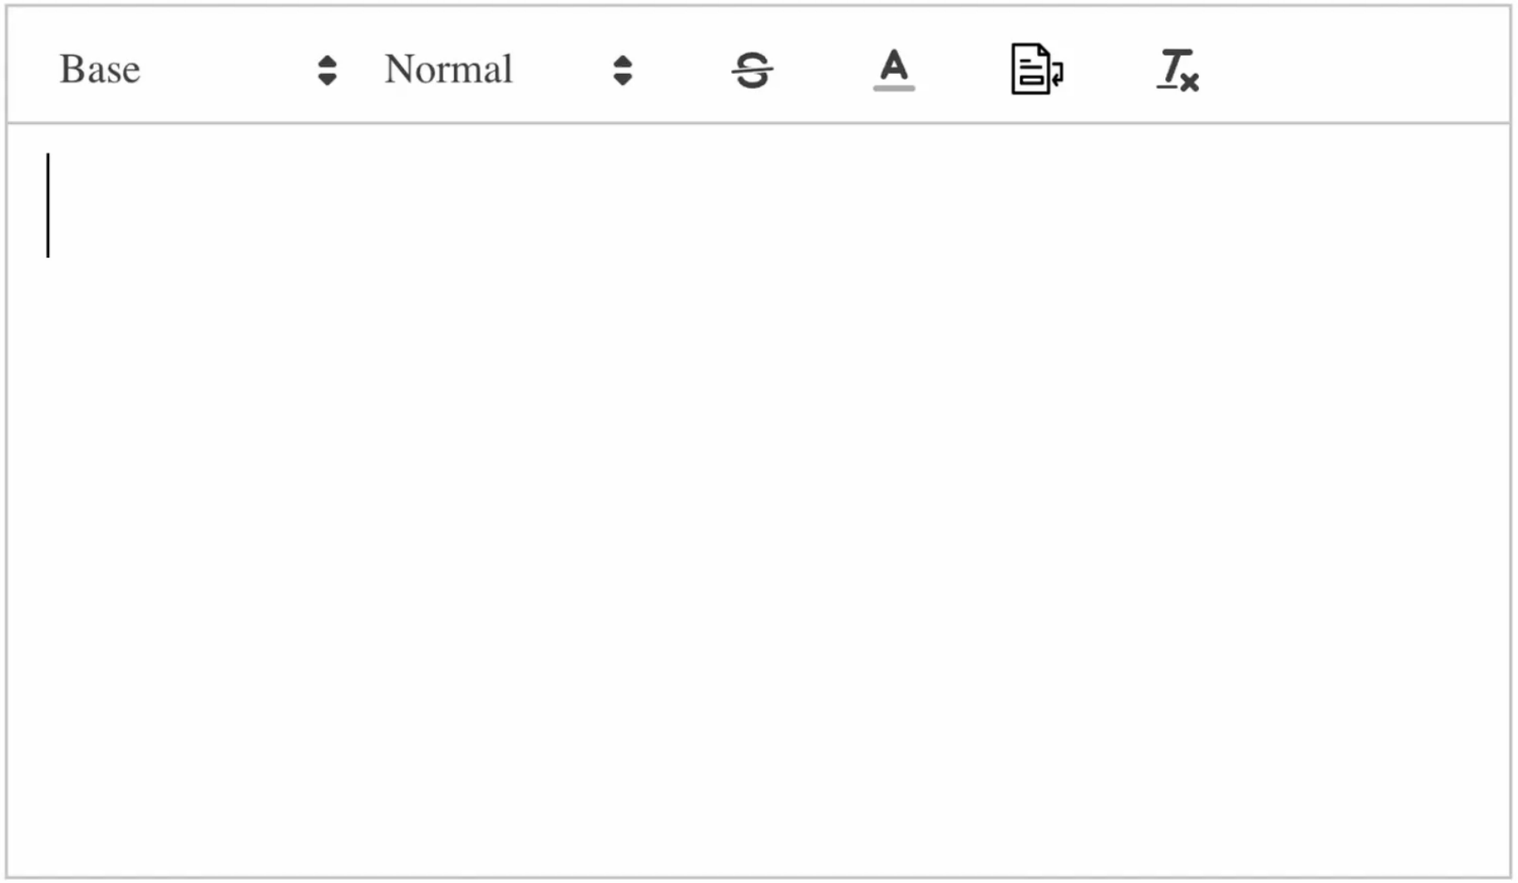

*The computation time is skipped in this video.

Text-Guided Synthesis of Eulerian Cinemagraphs [Mahapatra et al., SIGGRAPH Asia 2023]

Generative Models (2023)

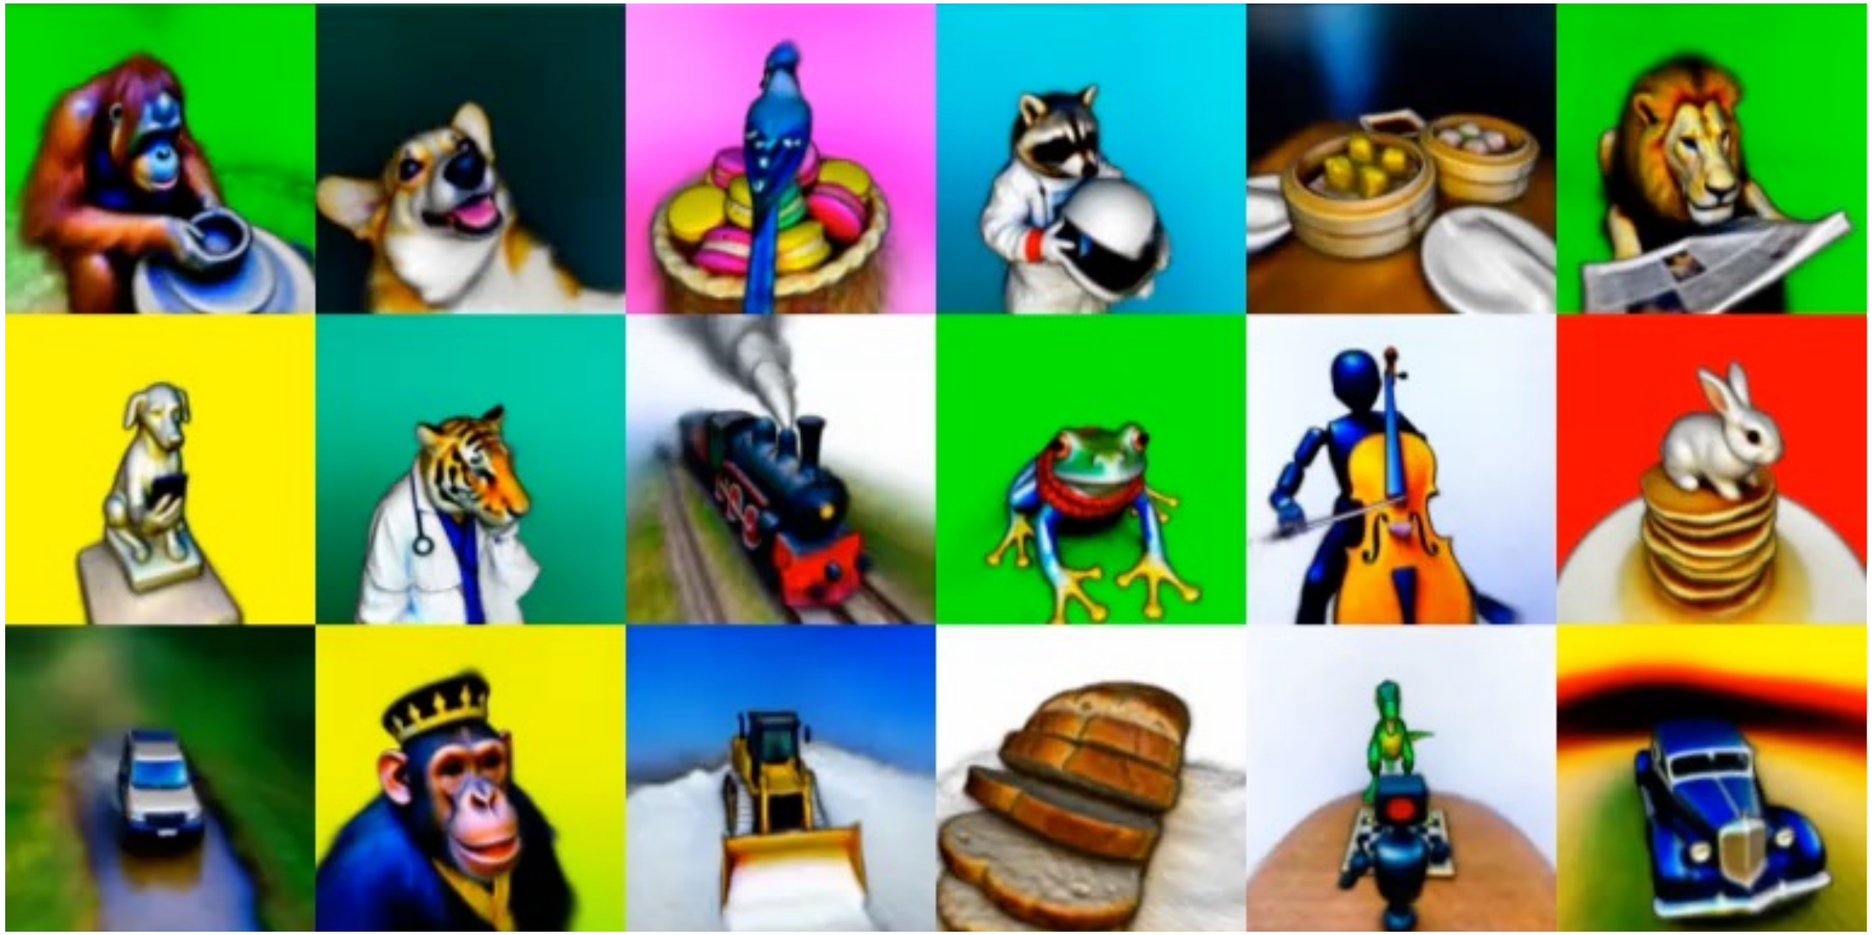

DreamFusion: Text-to-3D using 2D Diffusion [Poole et al., ICLR 2023] from UC Berkeley/Google

Cave art

32,000 BC

Sculpture

1498

Painting

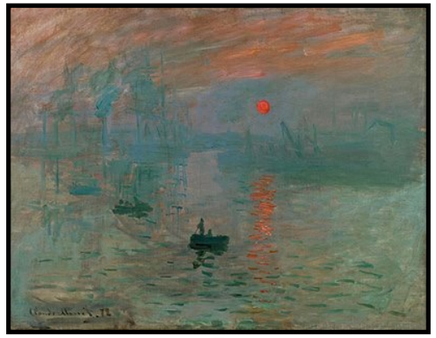

1872

Computer Graphics

2012

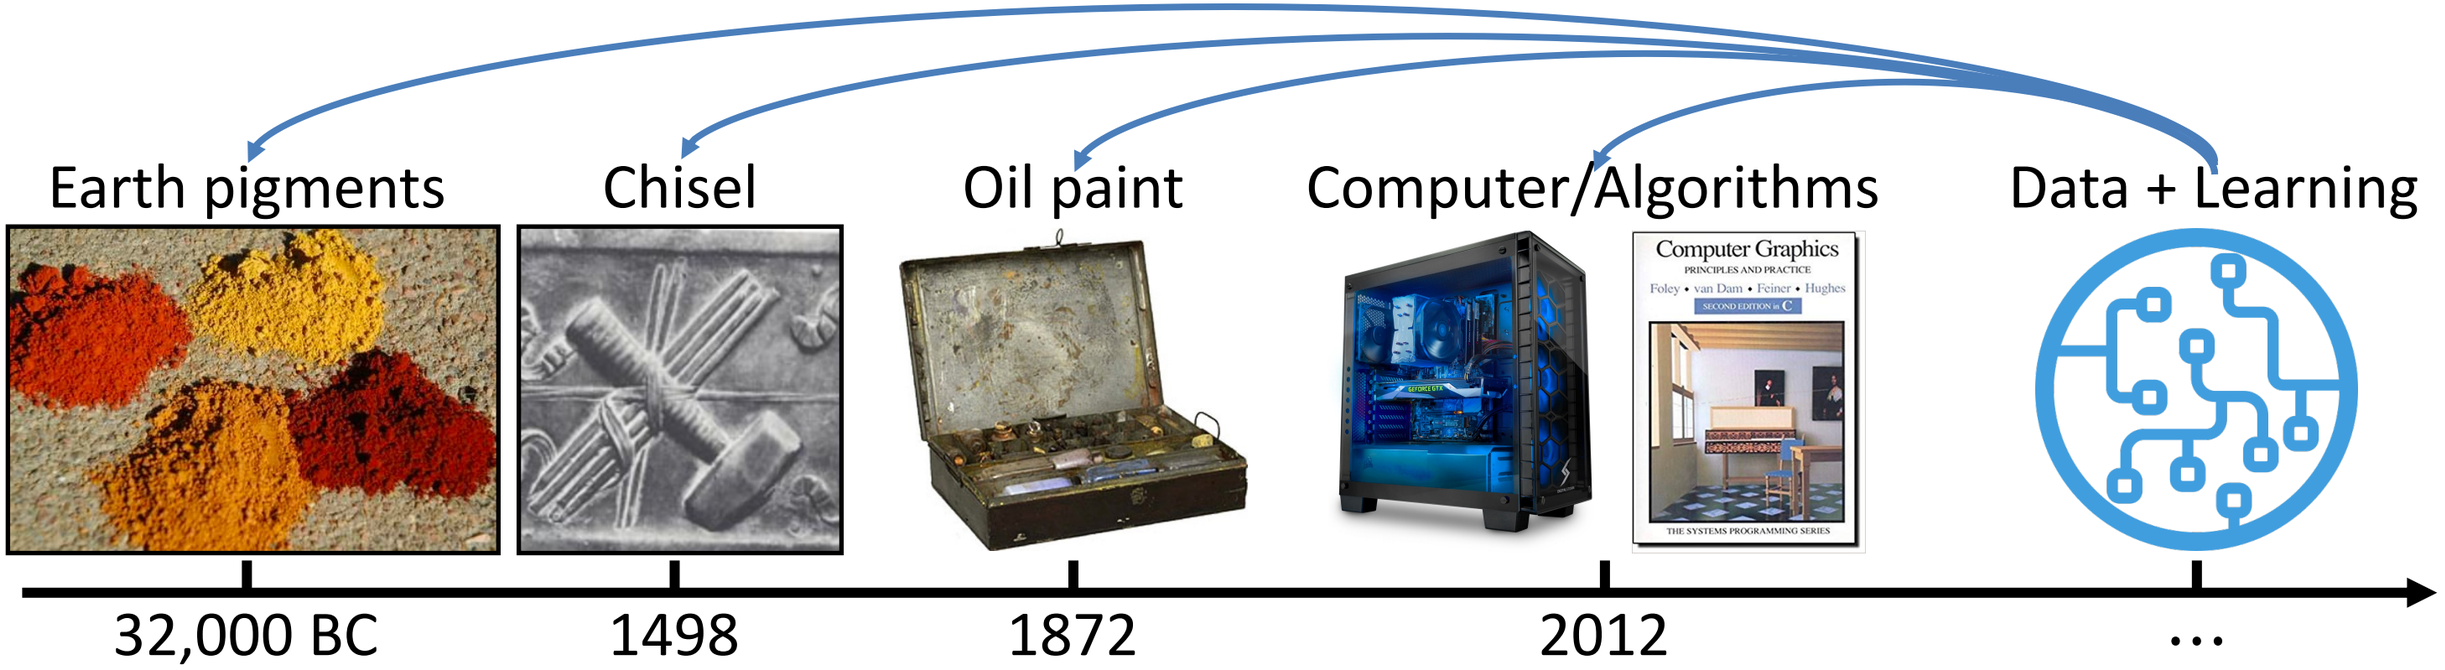

Course preview

- A modern machine learning perspective
- Widely-used learning algorithms
- Interactive content creation tools

Logistics

Course objectives

1. You will get a foundation in image editing and synthesis.
   - Texture synthesis and style transfer.
   - Face modeling and synthesis.
   - Image colorization and inpainting.
   - Video generation and editing.
   - Image-to-image translation.
   - Image and video editing. (warping, morphing, compositing)
   - Societal Implications: forensics, copyrights, biases.

Course objectives

2. You will get a foundation of machine learning concepts
   - (fast) Nearest neighbor search.
   - Principal component analysis, Gaussian Mixture model. Markov Random Field (MRF)
   - Convolutional neural networks.
   - Deep generative models: Auto-encoder, Generative Adversarial Networks, Flow-based models, Variational Auto-encoder, Autoregressive Models, Diffusion Models.
   - Conditional generative models.
   - Neural Radiance Fields (NeRF)

Course objectives

3. You will have some cool results with your own photos

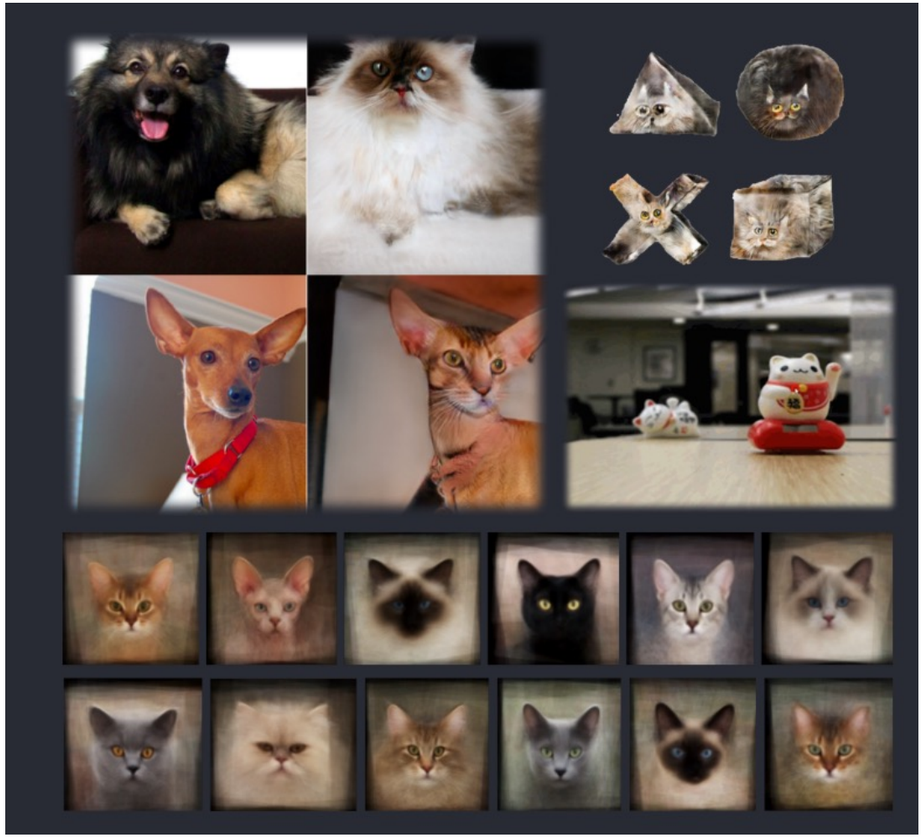

Textbook

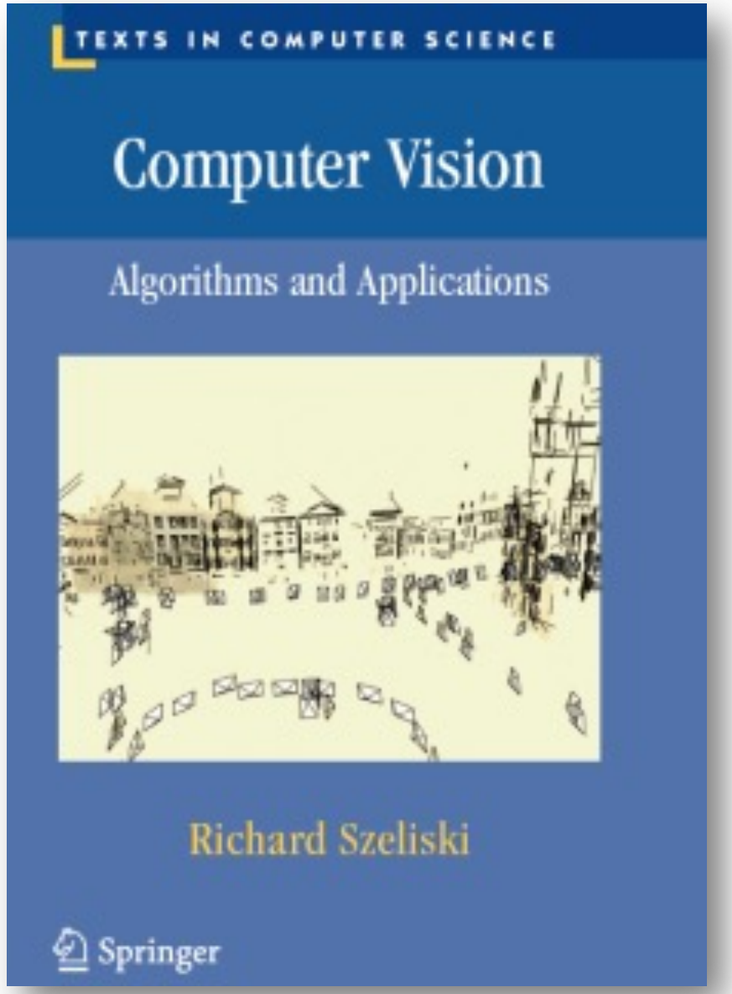

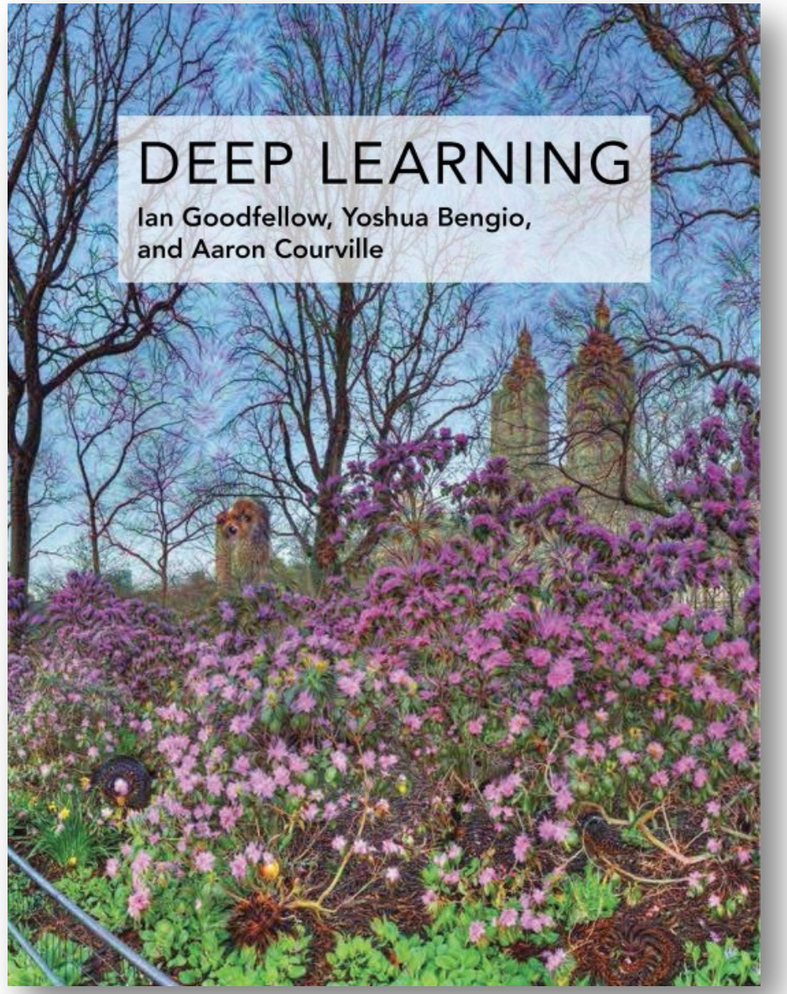

https://szeliski.org/Book/

https://www.deeplearningbook.org/

Grading

- Emphasis on programming projects (**60%**).
  - Classic: 1. image alignment. 2. image blending
  - Deep learning: 3. neural style transfer. 4. GANs and conditional GANs.
    5. reconstructing and editing an image with GANs.
- Late Policy for programming assignments.
  - Five (5) emergency late days for semester, to be spent wisely
  - 10% of penalty per 24 hours afterwards
- Class attendance (**5%**)
- One paper presentation (**10%**):
  - 10-20 min, 1-2 people in a group.
  - Need to answer questions about this paper from now on.
- Final Project (**25%**)
  - A webpage-based report + a presentation.
  - No late day.
  - 2-3 people per group.

Assignments

Assignment #0 - How to submit assignments?

Assignment #1 - Colorizing the Prokudin-Gorskii Photo Collection
Winner: [Konwoo Kim]
Honorable Mentions: [Juyong Kim] [Zihang Lai] [Manuel Rodriguez]

Assignment #2 - Gradient Domain Fusion
Winner: [Manuel Rodriguez]
Honorable Mentions: [George Cazanavette]

Assignment #3 - When Cats meet GANs
Winner: [Jun Luo]
Honorable Mentions: [George Cazanavette]

Assignment #4 - Neural Style Transfer
Winner: [Zihang Lai]
Honorable Mentions: [Zijie Li] [Tarang Shah]

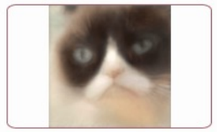

Assignment #5 - GAN Photo Editing
Winner: [George Cazenavette]
Honorable Mentions: [Manuel Guevara] [Zijie Li] [Zhe Huang]

+ more Diffusion Model modules

For each assignment

- Derive the math, implement stuff from scratch (+ starter code), and apply it to your own photos
- Every person does their own project (except final)
- Reporting via web page (+ submit code to Canvas)
- Afterwards, vote for class **favorite**(s)! **Gift**!
- Programming Language:
  - Python and PyTorch
  - you can use other languages, but you are on your own

Academic Integrity

- Can discuss projects, but don’t share code
- Don’t look up code or copy from a friend
- If you’re not sure if it’s allowed, ask
- Acknowledge any inspirations
- If you get stuck, come talk to us

Getting help outside of class

- Course Web Page
  - https://16726-image-synthesis.github.io/sp24/
- Discussion board:
  - Piazza.com
- Assignment submission
  - Canvas
- Office hours
  - See piazza posts

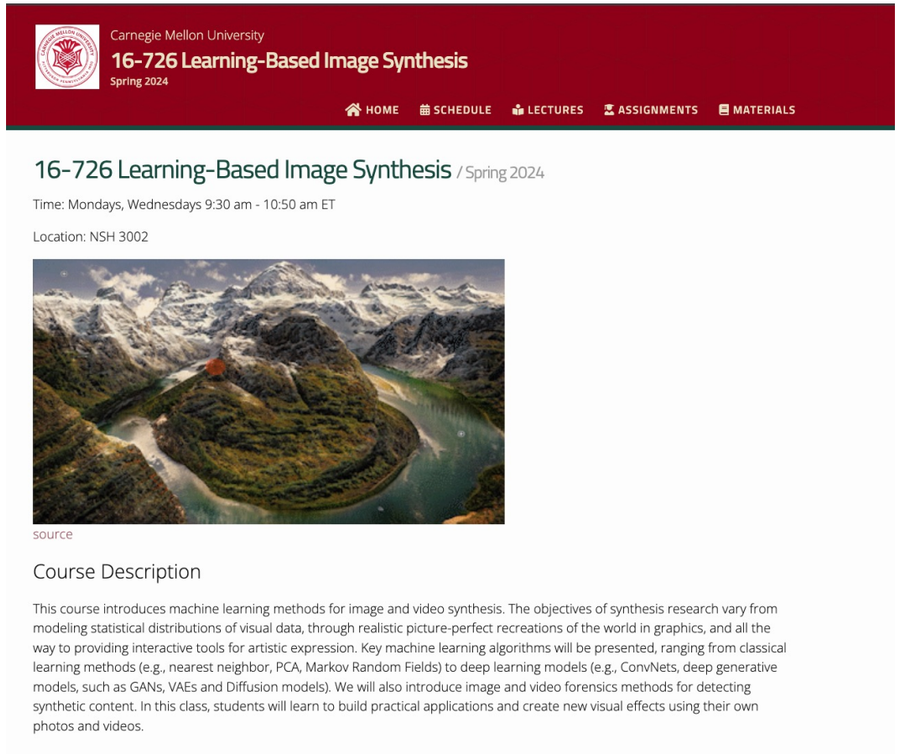

Why you should NOT take this class

- Project-based class
  - No canned problem sets.
  - Not theory-heavy.
  - will read many research papers.
  - Open-ended by design.
- Need time to think, not just hack
  - **Creativity** is a class requirement.
- Not worth it if you don’t enjoy it.

Now… reasons TO take this class

- Not too many similar courses at other places.
- You get to create pictures and unleash your creative potential.
- Interested in grad school and research? ☺
- Interested in industry jobs? ☺

Become a friend with every pixel!

- Jitendra Malik, UC Berkeley

Thank You!

16-726, Spring 2024

https://learning-image-synthesis.github.io/sp24/