


Differentiable Visual Computing

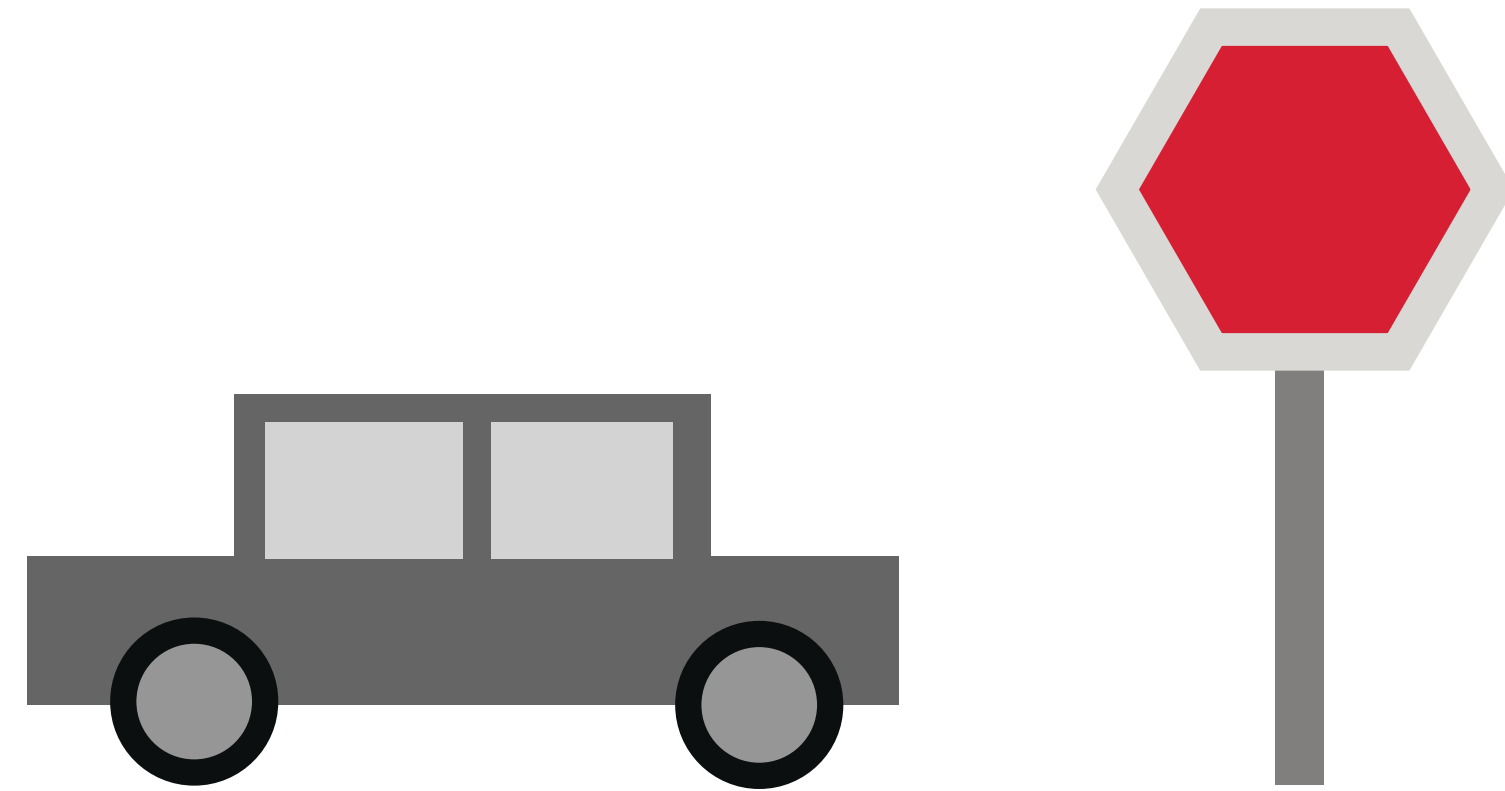
Tzu-Mao Li (李子懋)

UCSD


The world is visual




robots




autonomous driving




camera



virtual reality


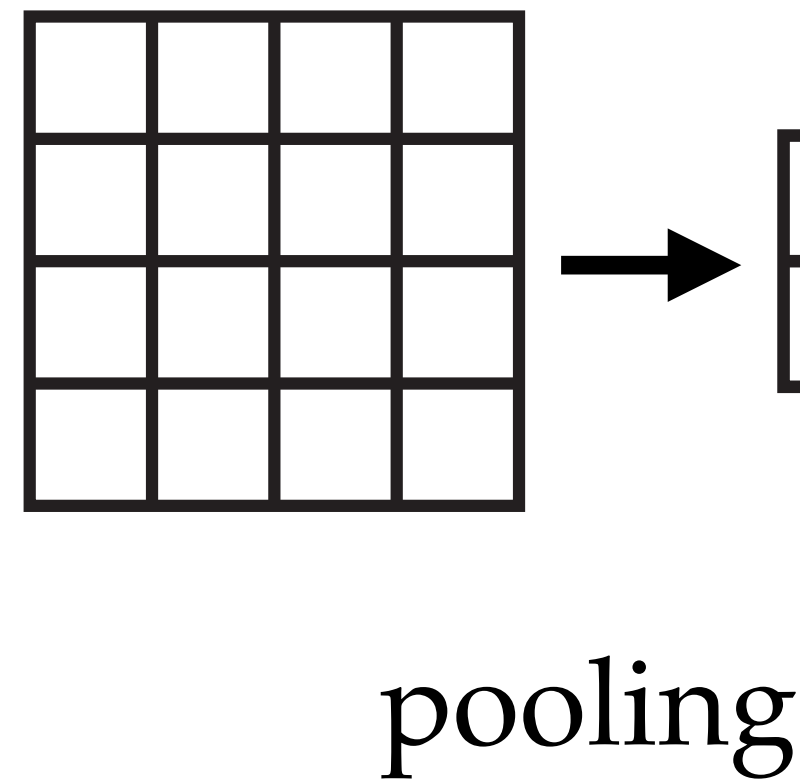
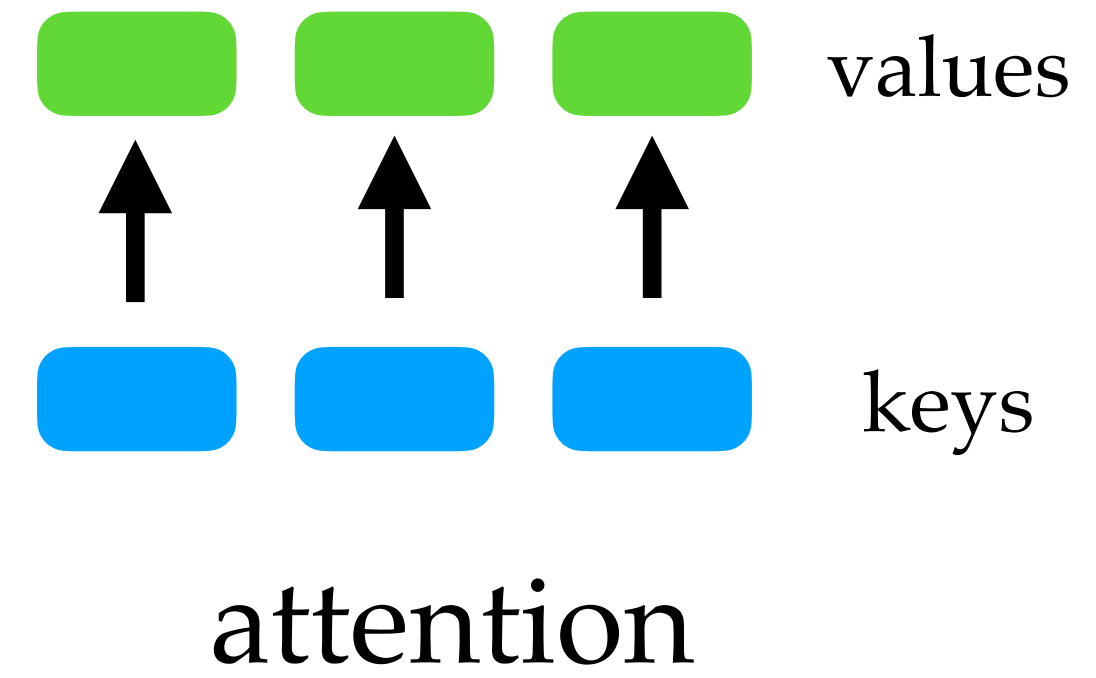
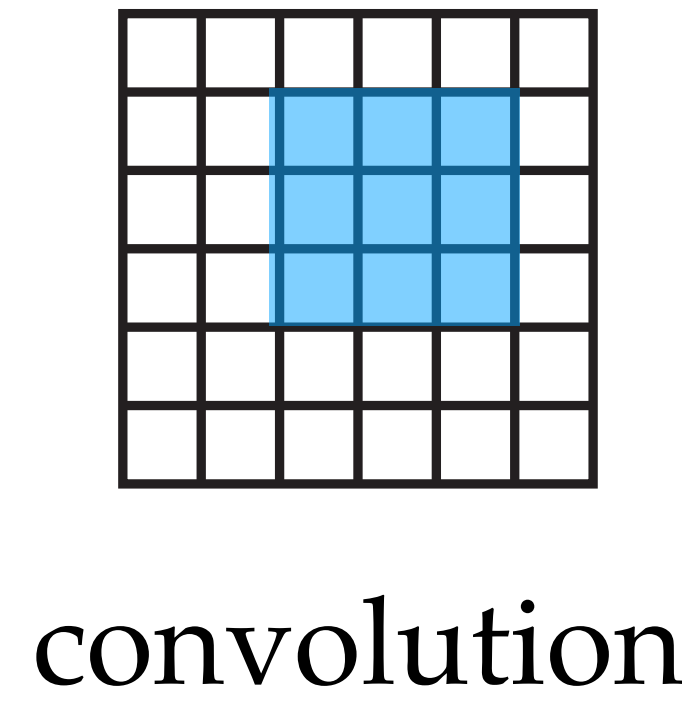




medical
imaging



scientific visualization

Neural networks: powerful visual processors




Beyond neural networks

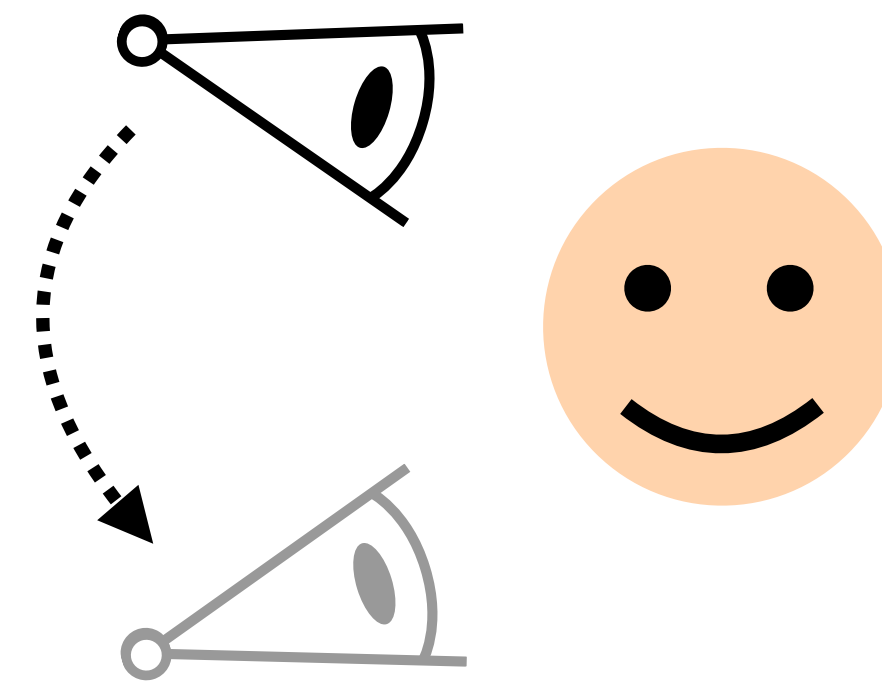
- embed prior knowledge e.g., 3D reasoning



3D scene




image



Beyond neural networks

- embed prior knowledge e.g., 3D reasoning
- enable speed

deeplab-res101-v2 on full HD: 2.6 TFLOPs, 40GB features



Chen et al.

<https://github.com/albanie/convnet-burden>


RTX 2080 Ti
theoretical peak perf:
6.7 TFLOP/s

OpenAI recently published GPT-3, the largest language model ever trained. GPT-3 has 175 billion parameters and would require **355 years** and \$4,600,000 to train even with the lowest priced GPU cloud on the market.
Jun 3, 2020



Beyond neural networks

- embed prior knowledge e.g., 3D reasoning
- enable speed
- debug & control




Classical methods: explicit modeling




3D scene




image




image



edges




sound



spectrogram




initial



simulated

Classical methods: explicit modeling

- ✓ embedded knowledge
- ✓ tailored to the application
- ✓ debug and control



Classical methods: explicit modeling

- ✓ embedded knowledge
- ✓ tailored to the application
- ✓ debug and control

- ✗ do not apply as broadly
- ✗ do not learn from data

sound

spectrogram

initial


simulated

Key to connect classical methods and deep learning: derivatives


derivatives enable intelligent decisions





3D scene




image





image



edges





sound



spectrogram

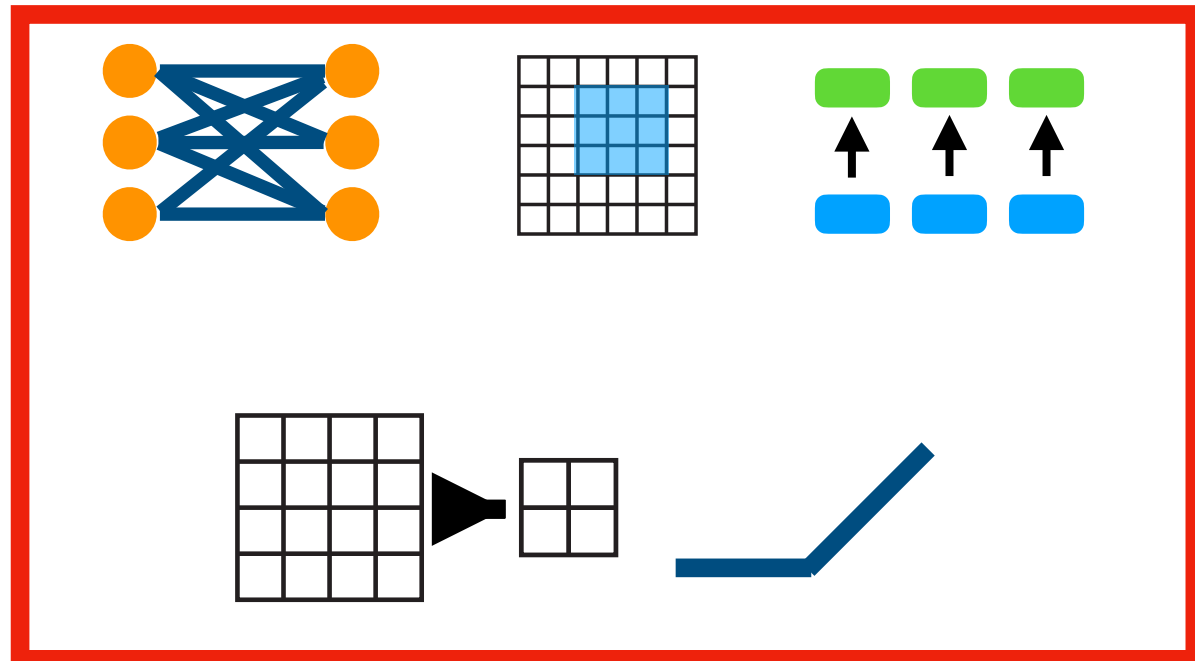


initial

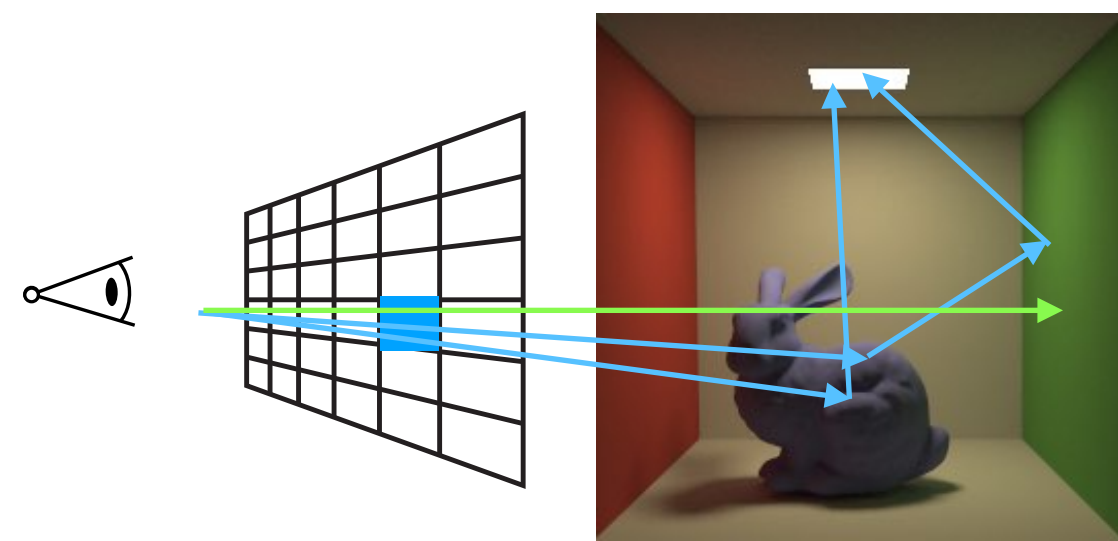


simulated


There is a huge world outside of the existing deep learning toolbox




deep learning
framework toolbox





ray tracing / rasterization





edge-aware filtering



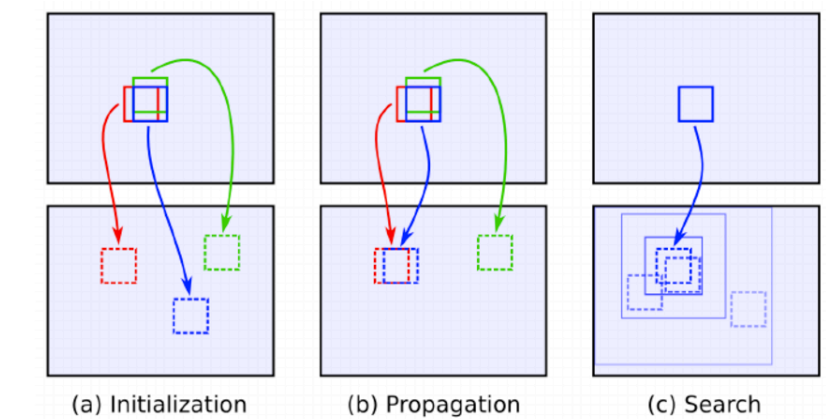
PDE solvers




meshing




mean-shift clustering



patchmatch



Markov random field




graph cut

Differentiable graphics


connects classical graphics algorithms with
modern data-driven methods

Differentiable graphics




connects classical graphics algorithms with
modern data-driven methods



geometry modeling



animation/
physical simulation



rendering



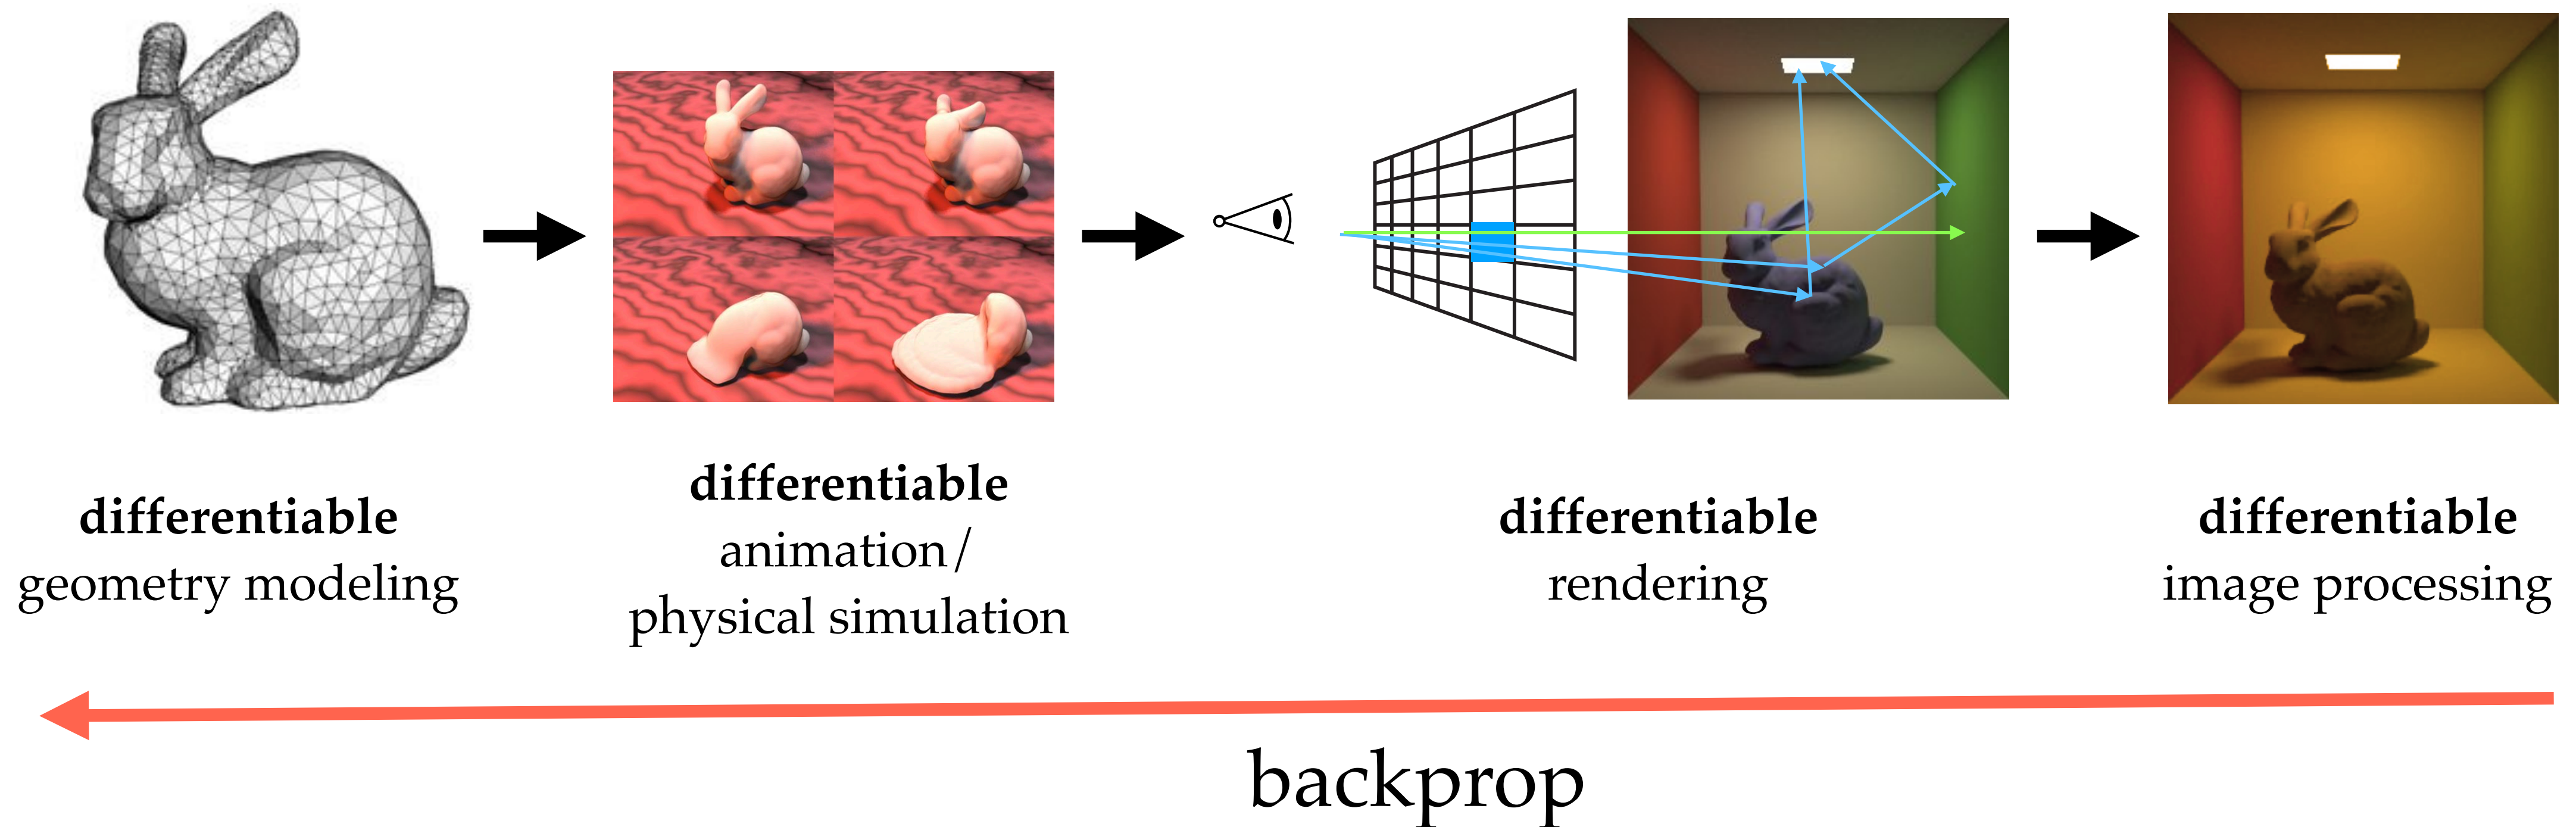


image processing

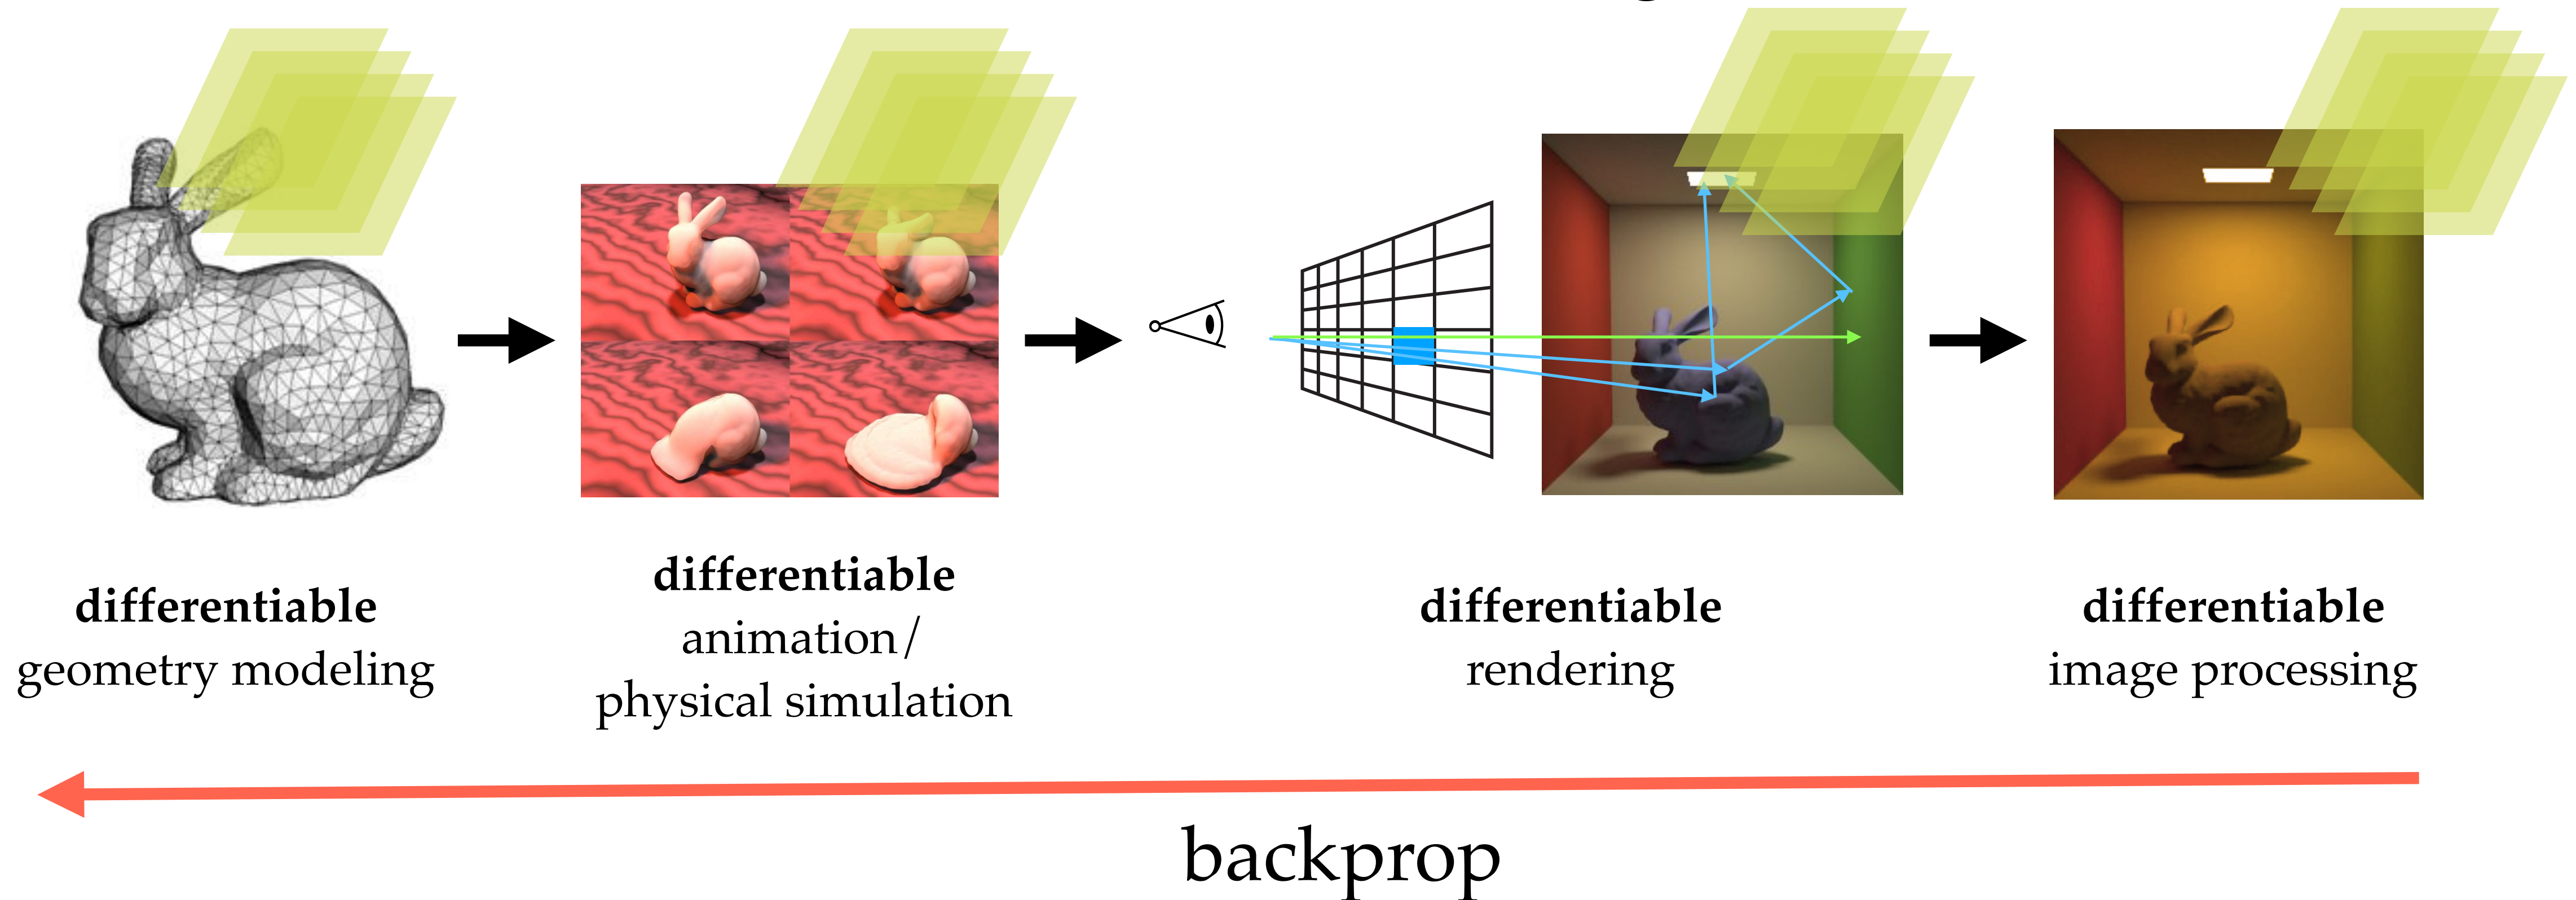
Differentiable graphics

connects classical graphics algorithms with modern data-driven methods **through derivatives**



Differentiable graphics

connects classical graphics algorithms with modern data-driven methods **through derivatives**




Differentiating graphics algorithms is hard

- beyond convolution
- discontinuities
- tedious derivation

$$\nabla_{\theta} \int_{A(\theta)} g(x) dx = \int_{A(\theta)} \nabla_{\theta} g(x) dx + \int_{\partial A(\theta)} (\nabla_{\theta} p(\theta) \cdot n) g(p) dp$$
$$\frac{\partial \mathbf{x}}{\partial t} = \frac{\partial E}{\partial \mathbf{p}} \quad \frac{\partial \mathbf{t}'}{\partial \mathbf{c}_i} = \frac{-1}{\frac{\partial f}{\partial \mathbf{c}_i}} \frac{\partial f}{\partial \mathbf{x}}$$

```
auto scatter_contrib = Vector3(0, 0, 0);
auto scatter_bsdf = Vector3(0, 0, 0);
if (bsdf_isect.valid()) {
    const auto &bsdf_shape = scene.shapes[bsdf_isect.shape_id];
    auto dir = bsdf_point.position - p;
    auto dist_sq = length_squared(dir);
    auto wo = dir / sqrt(dist_sq);
    auto pdf_bsdf = bsdf_pdf(material, shading_point, wi, wo, min_rough);
    if (dist_sq > 1e-20f && pdf_bsdf > 1e-20f) {
        auto bsdf_val = bsdf(material, shading_point, wi, wo, min_rough);
        if (bsdf_shape.light_id >= 0) {
            const auto &light = scene.area_lights[bsdf_shape.light_id];
            if (light.two_sided || dot(-wo, bsdf_point.shading_frame.n) > 0)
                light_contrib = light.intensity;
            auto light_pmf = scene.light_pmf[bsdf_shape.light_id];
            auto light_area = scene.light_areas[bsdf_shape.light_id];
            auto inv_area = 1 / light_area;
            auto geometry_term = fabs(dot(wo, bsdf_point.geom_normal));
            auto pdf_nee = (light_pmf * inv_area) / geometry_term;
            auto mis_weight = Real(1 / (1 + square((double)pdf_nee / pdf_bsdf)));
            scatter_contrib = (mis_weight / pdf_bsdf) * bsdf_val * light_contrib;
        }
    }
}
scatter_bsdf = bsdf_val / pdf_bsdf;
next_throughput = throughput * scatter_bsdf;
```


tf.conv2d v.s.



Challenges of both algorithms and systems

beyond traditional automatic differentiation

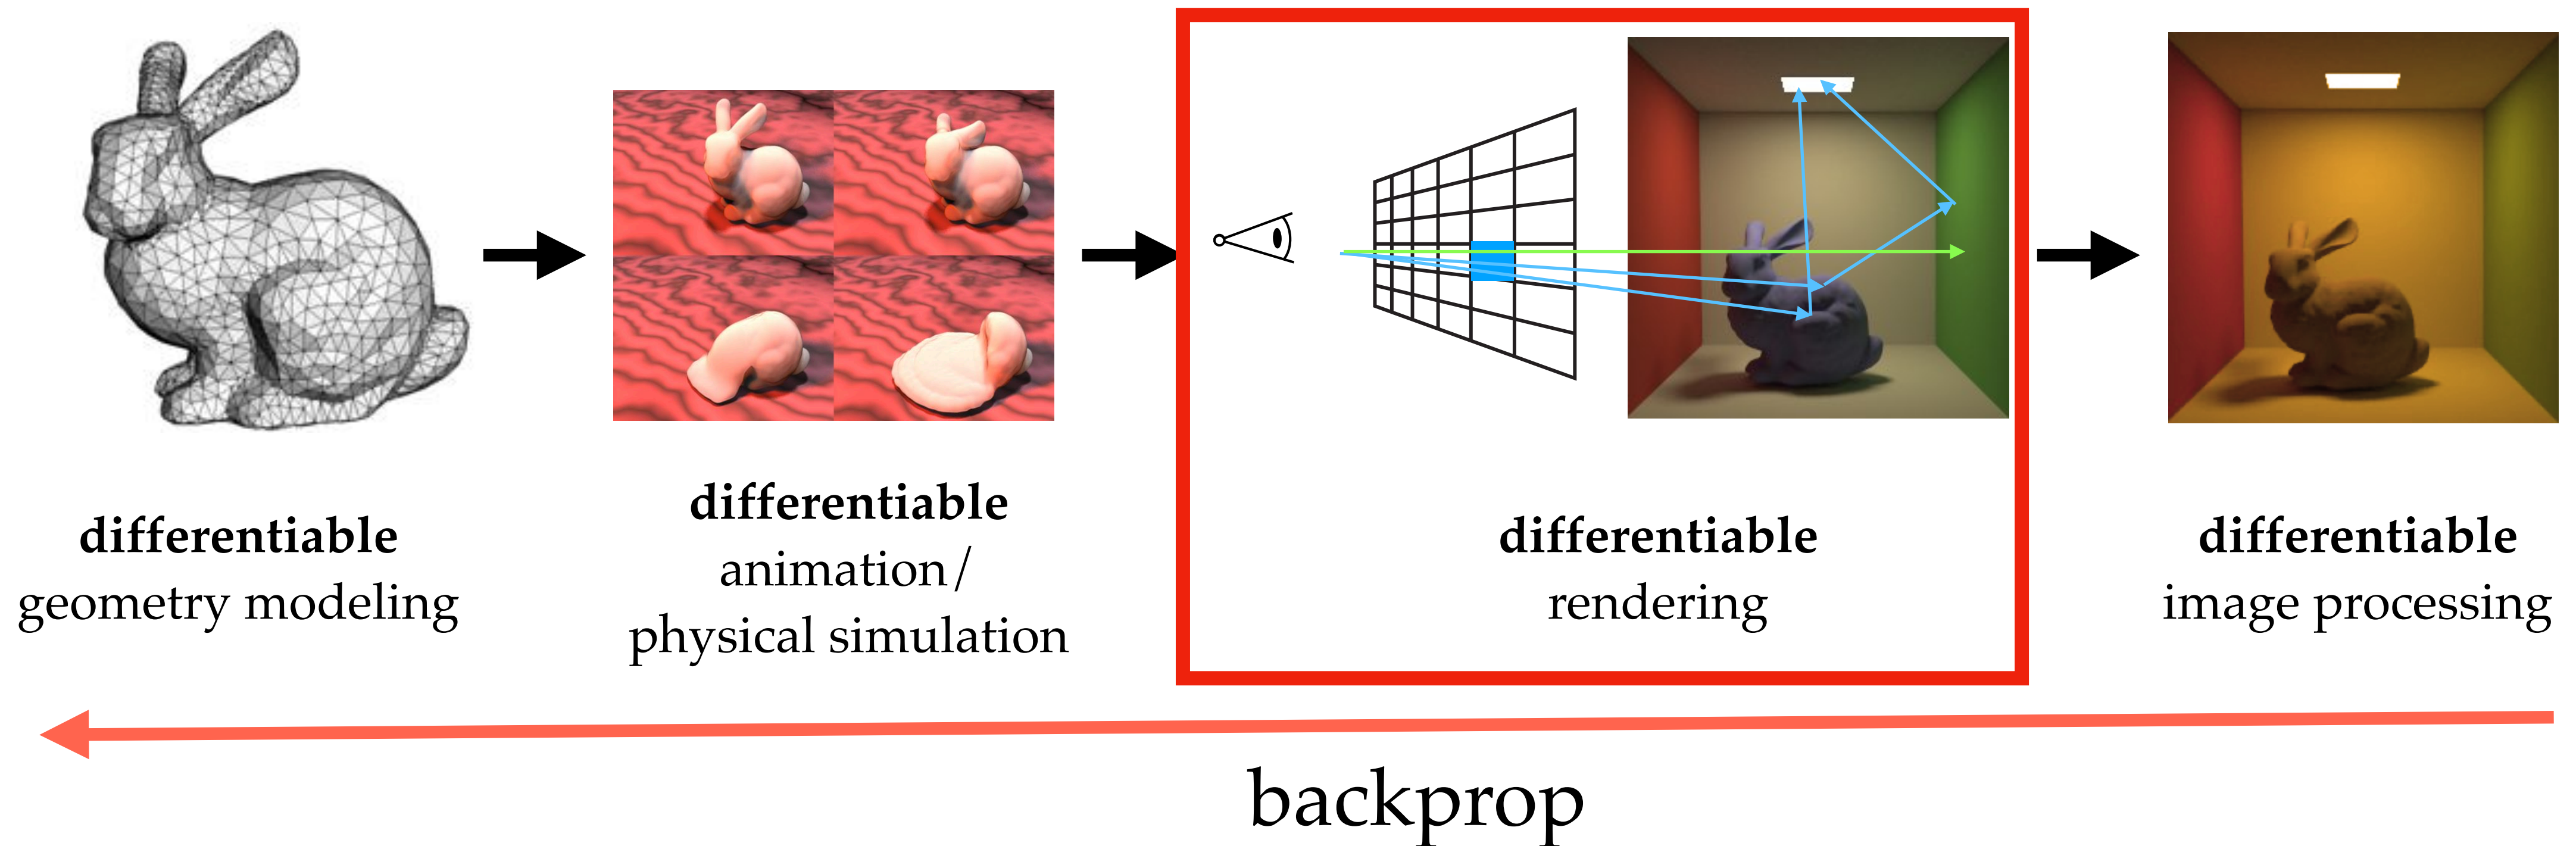
- need differentiable algorithms
- need differentiating compilers



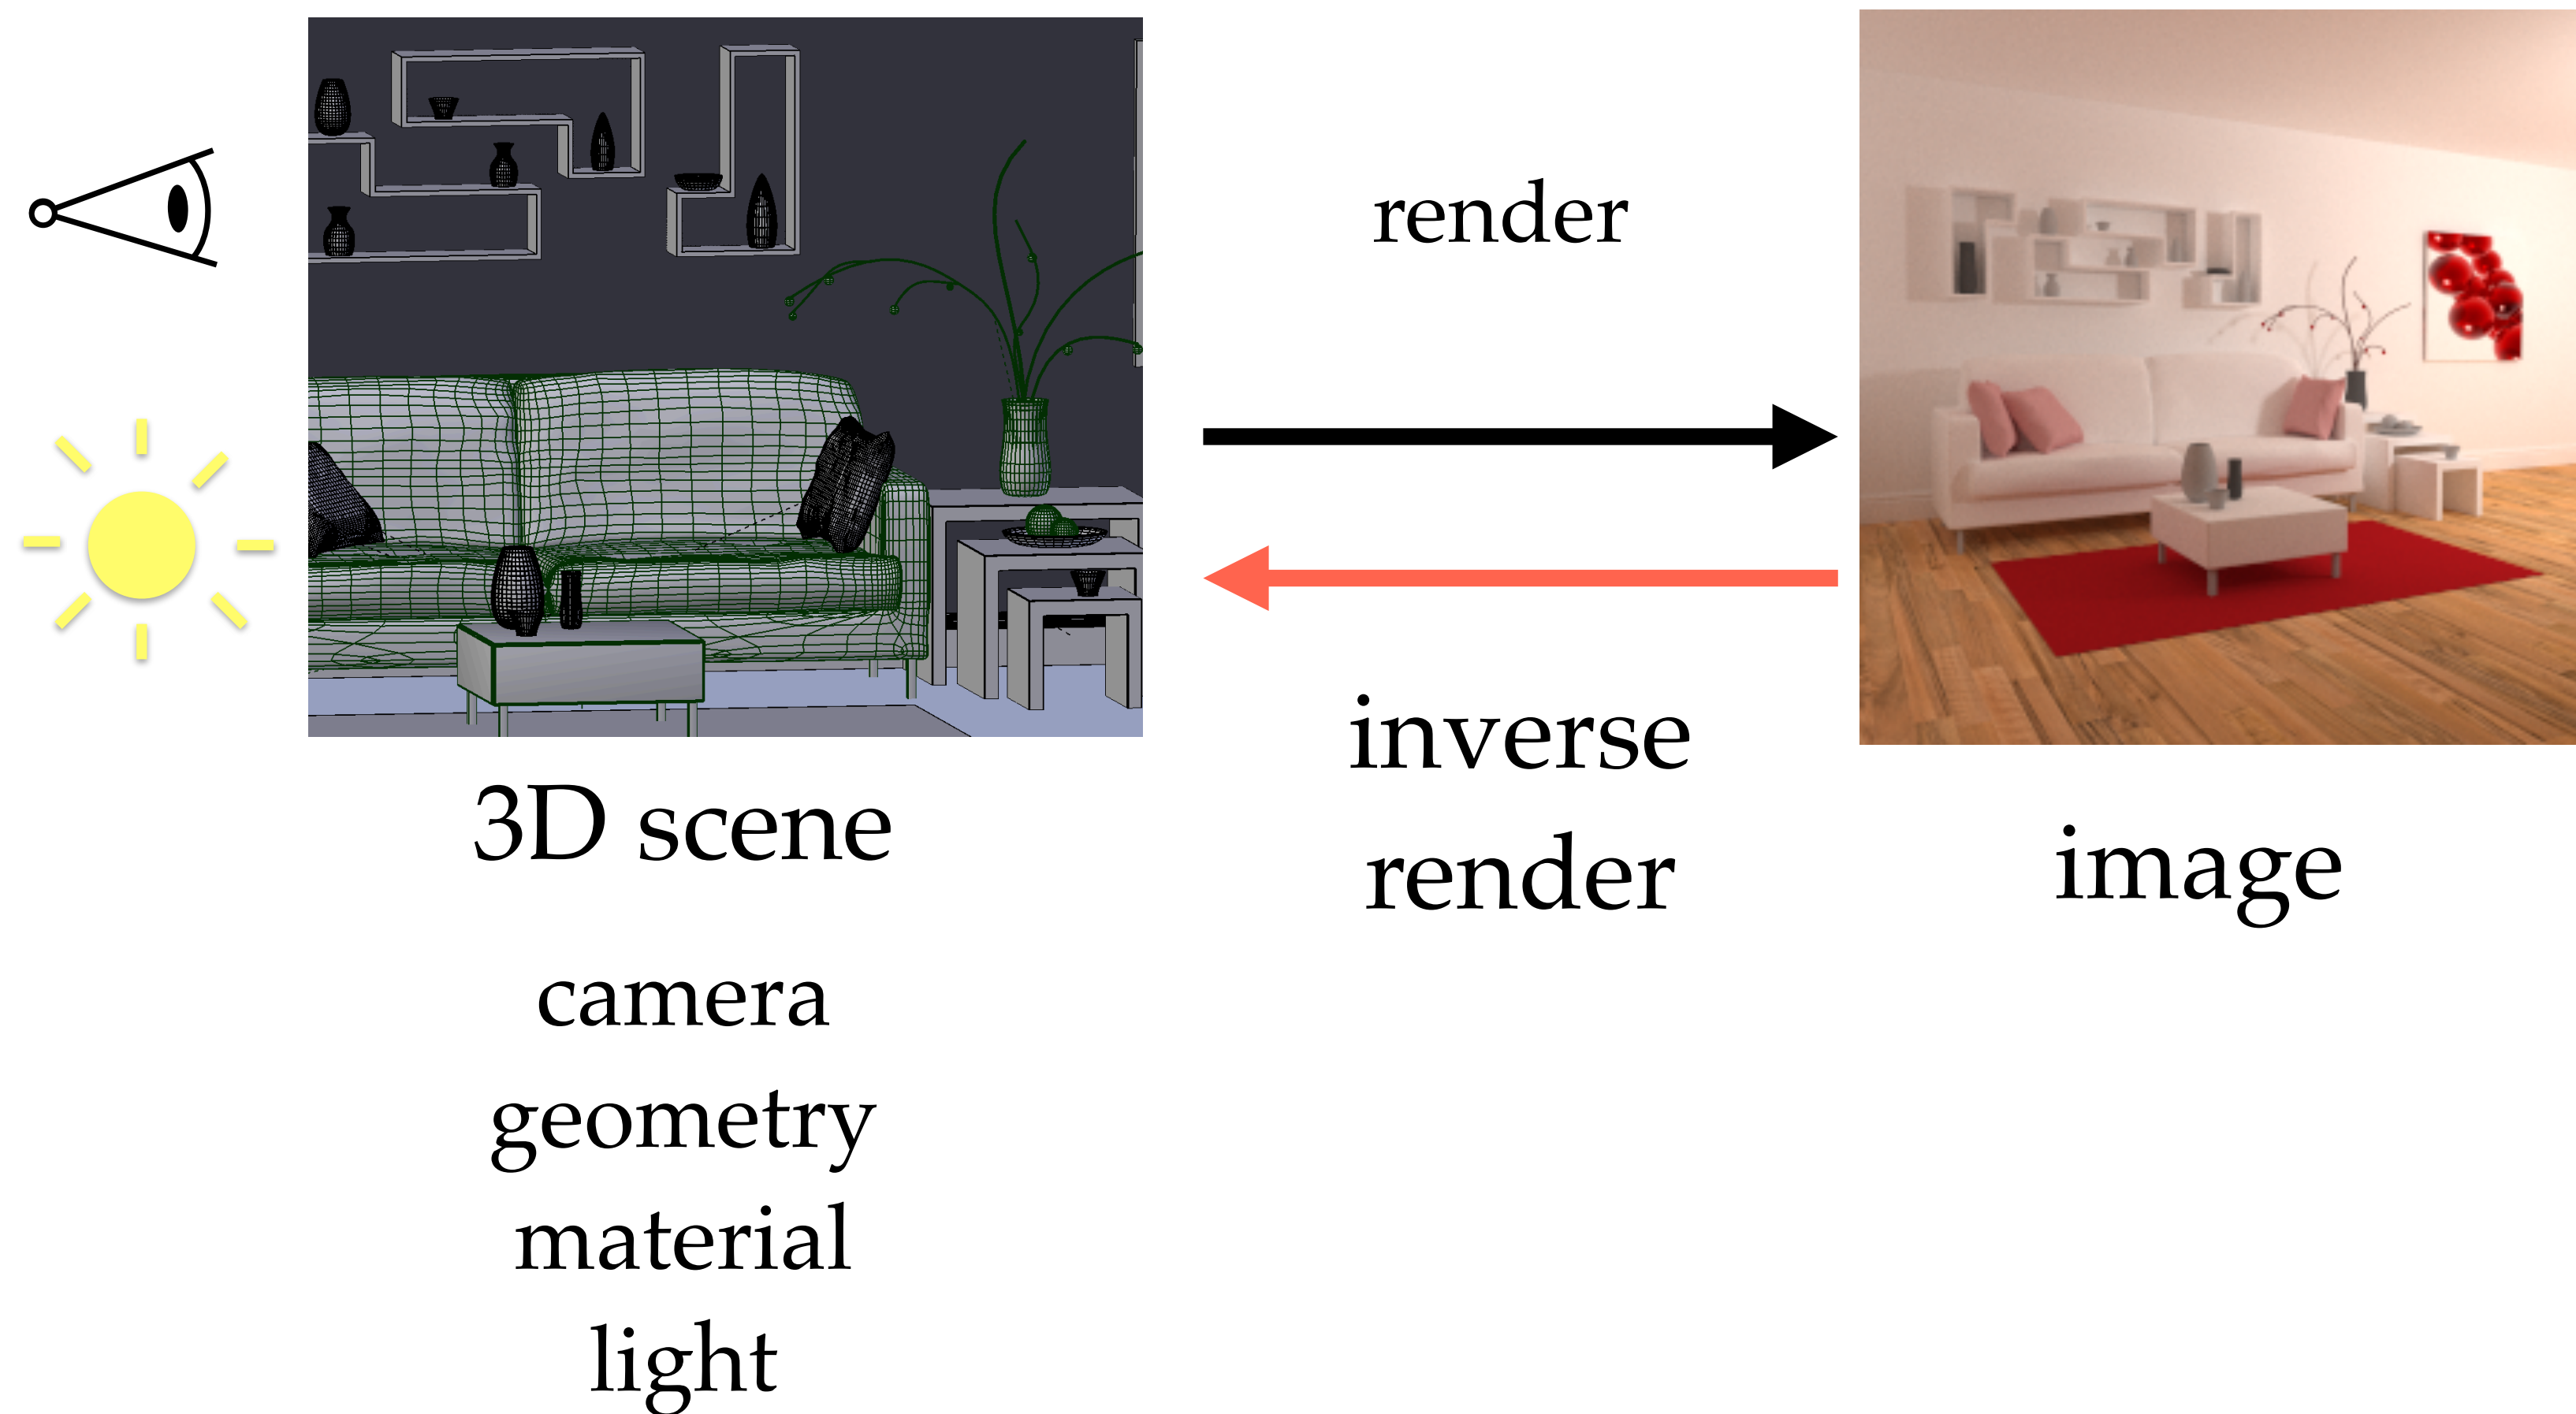
```
auto scatter_contrib = Vector3{0, 0, 0};
auto scatter_bsdf = Vector3{0, 0, 0};
if (bsdf_isect.valid()) {
  const auto &bsdf_shape = scene.shapes[bsdf_isect.shape_id];
  auto dir = bsdf_point.position - p;
  auto dist_sq = length_squared(dir);
  auto wo = dir / sqrt(dist_sq);
  auto pdf_bsdf = bsdf_pdf(material, shading_point, wi, wo, min_rough);
  if (dist_sq > 1e-20f && pdf_bsdf > 1e-20f) {
    auto bsdf_val = bsdf(material, shading_point, wi, wo, min_rough);
    if (bsdf_shape.light_id >= 0) {
      const auto &light = scene.area_lights[bsdf_shape.light_id];
      if (light.two_sided || dot(-wo, bsdf_point.shading_frame.n) > 0)
        auto light_contrib = light.intensity;
      auto light_pmf = scene.light_pmf[bsdf_shape.light_id];
      auto light_area = scene.light_areas[bsdf_shape.light_id];
      auto inv_area = 1 / light_area;
      auto geometry_term = fabs(dot(wo, bsdf_point.geom_normal));
      auto pdf_nee = (light_pmf * inv_area) / geometry_term;
      auto mis_weight = Real(1 / (1 + square((double)pdf_nee) / pdf_bsdf));
      scatter_contrib = (mis_weight / pdf_bsdf) * bsdf_val * light_contrib;
    }
  }
}
scatter_bsdf = bsdf_val / pdf_bsdf;
next_throughput = throughput * scatter_bsdf;
```

Differentiable graphics

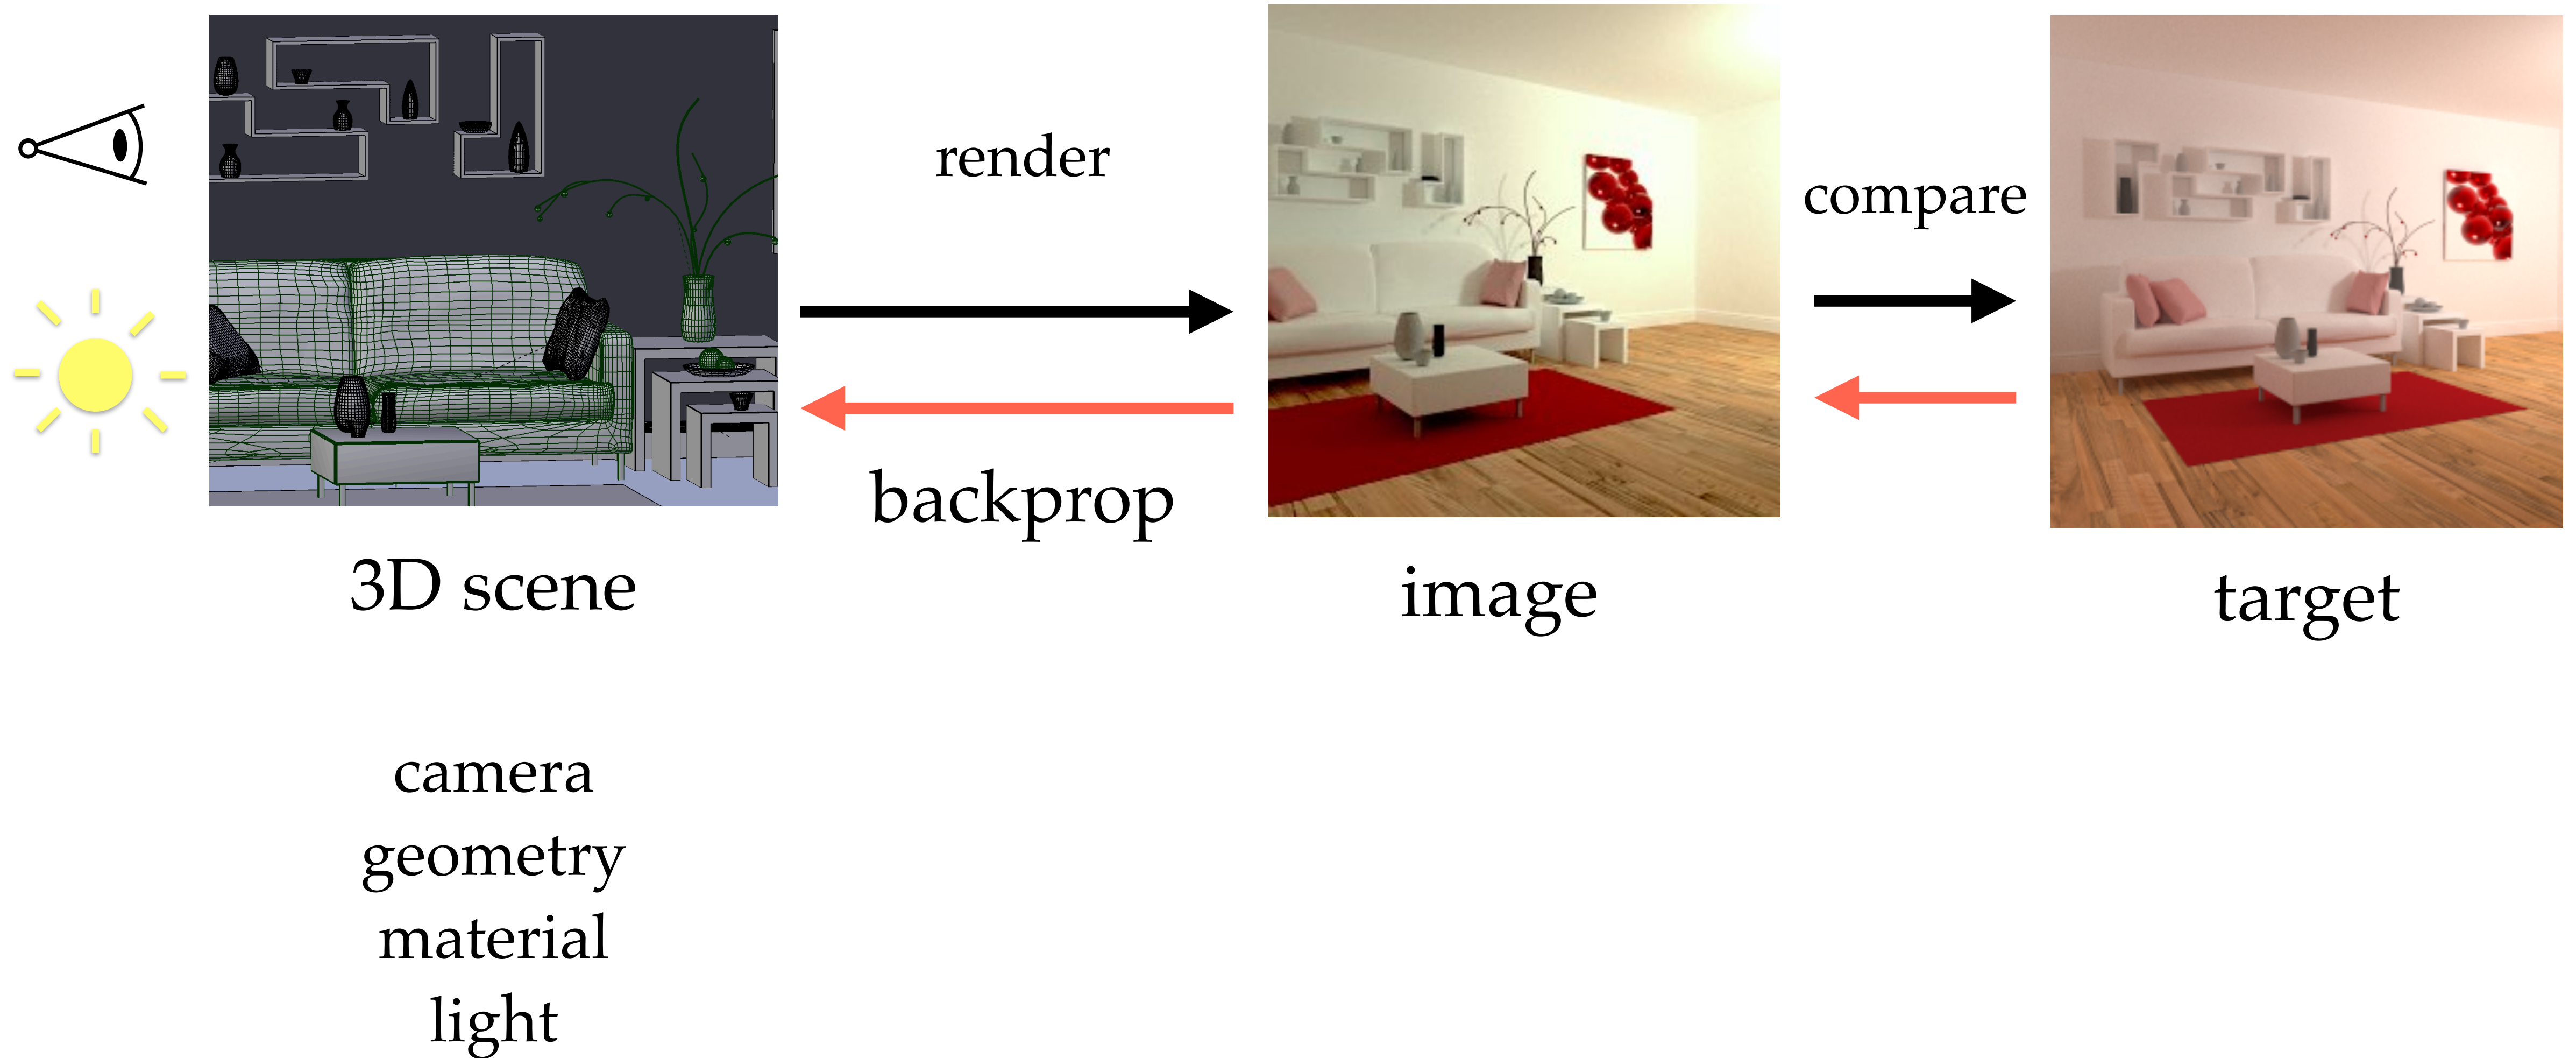
connects classical graphics algorithms with modern data-driven methods **through derivatives**



Motivation 1: inverse rendering

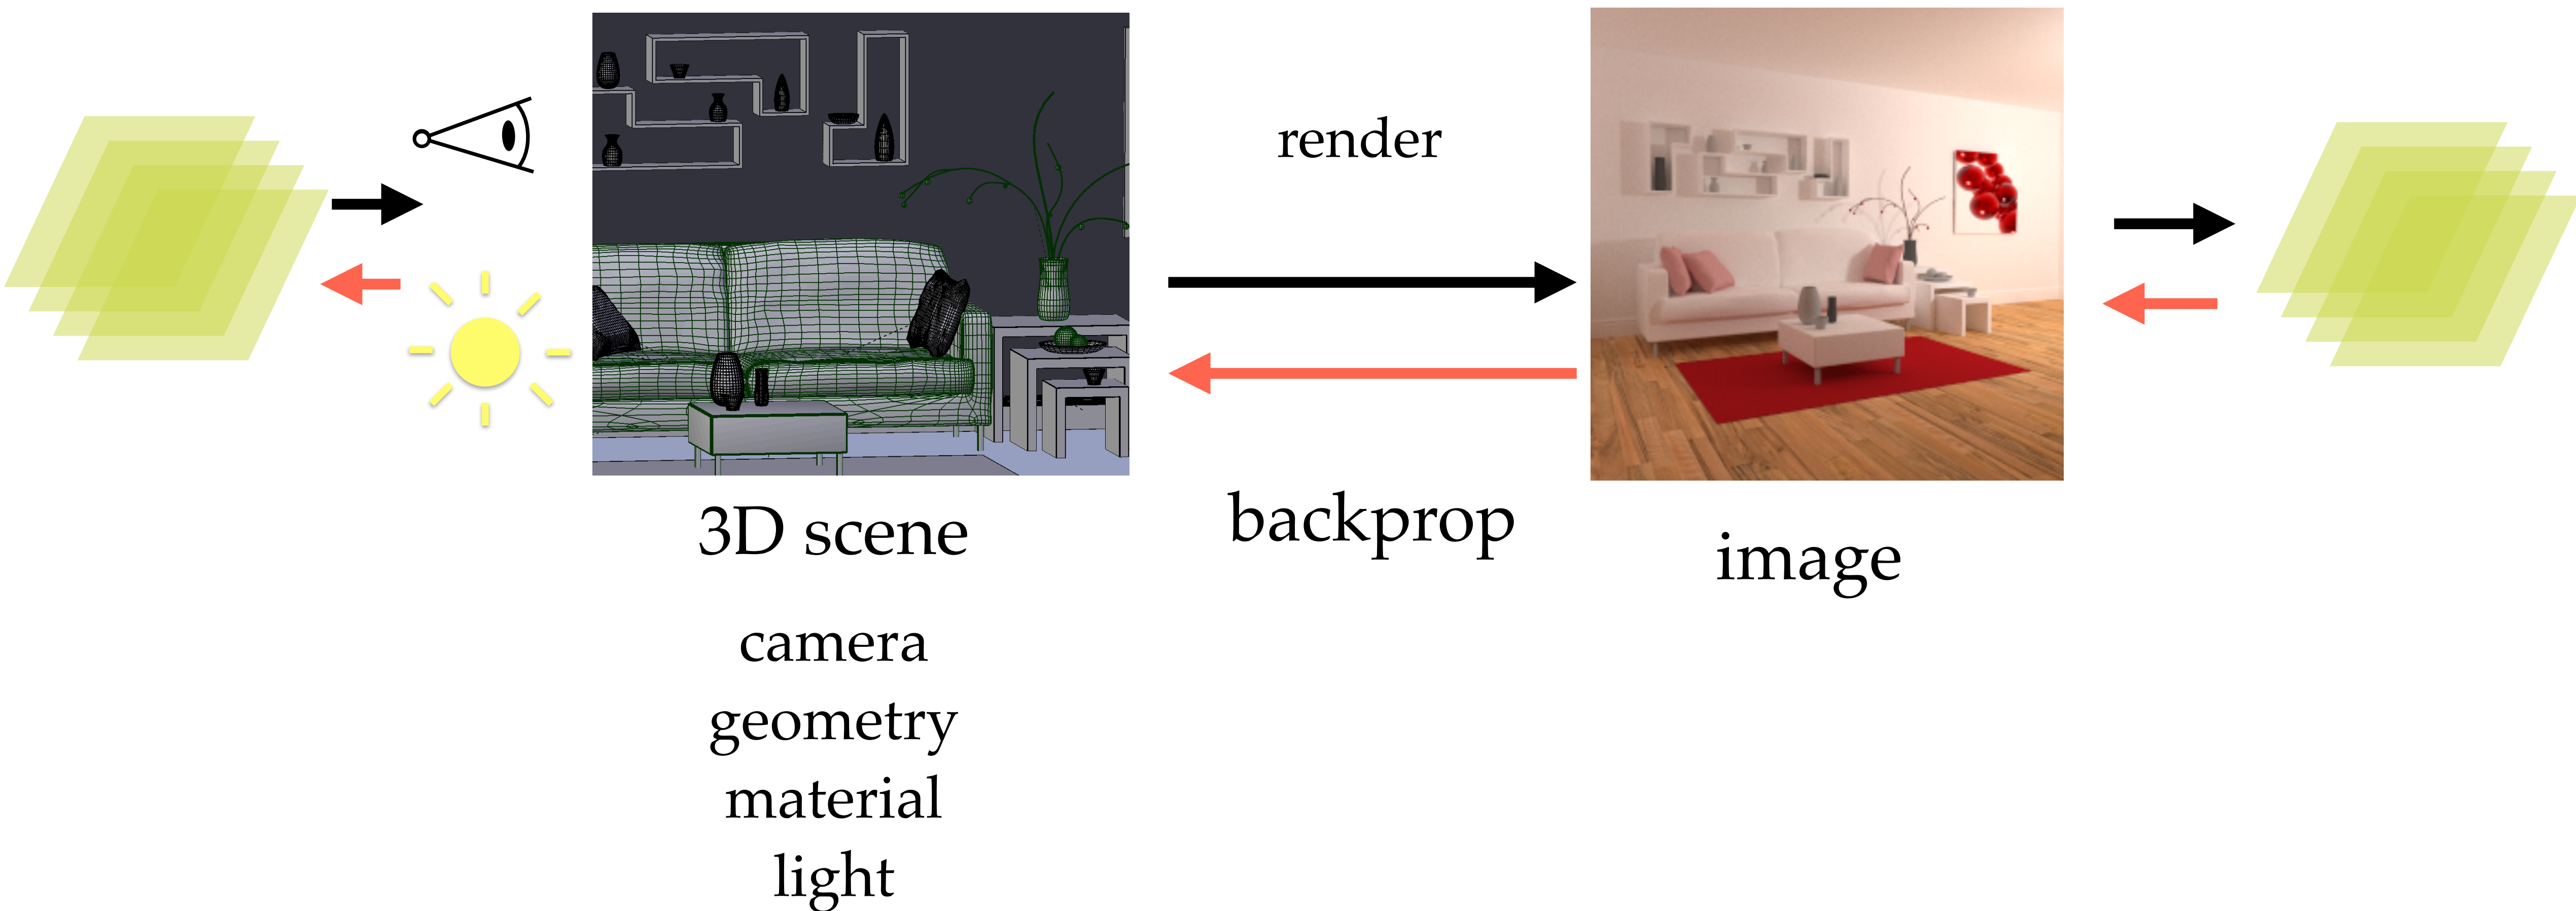


Use gradients to update 3D scene



Motivation 2: deep learning

adversarial robustness, self-supervised learning, etc



Many applications




virtual conferencing




robotics




fabrication



3D content creation




visual effects



biomedical optics

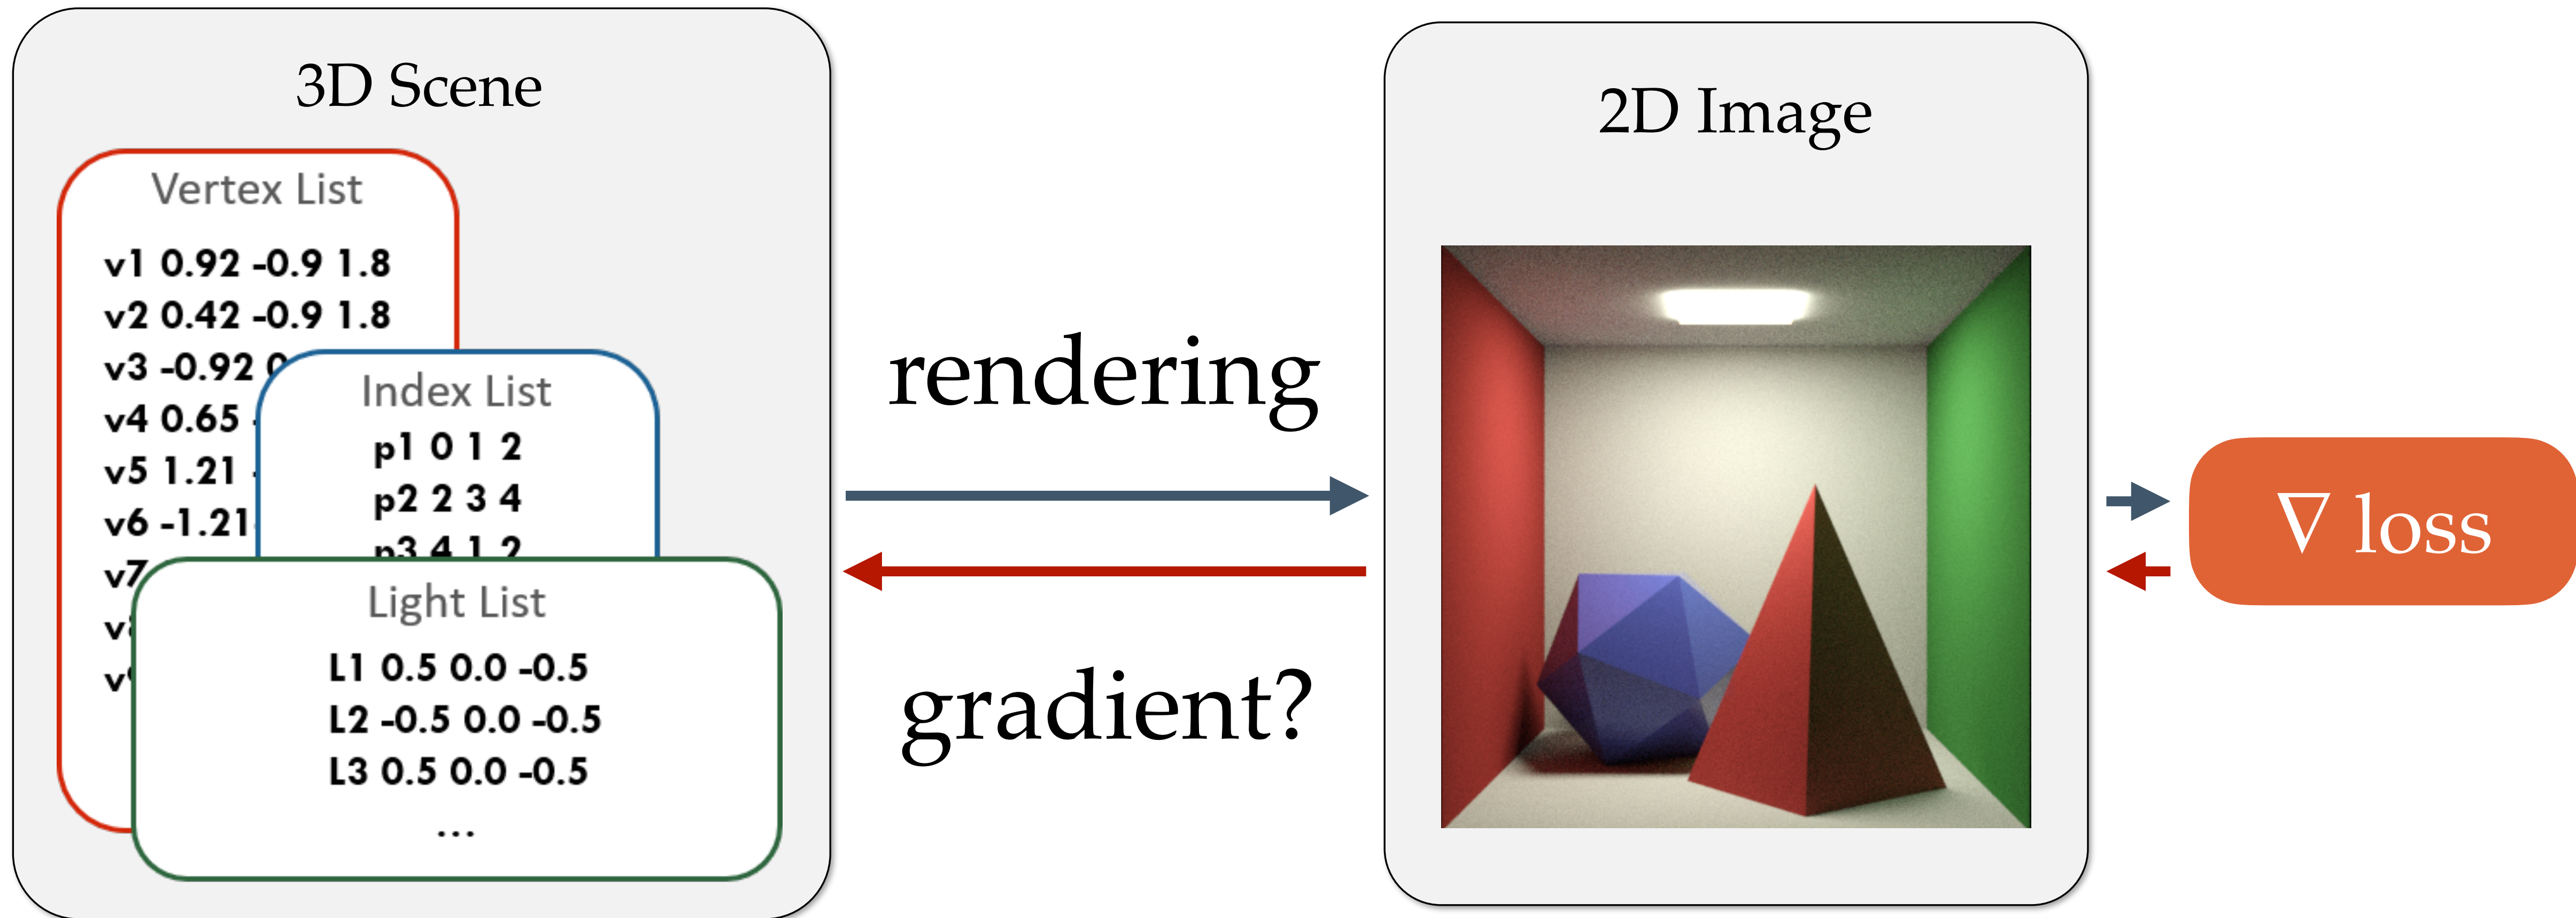


architectural design






computational imaging

Task: compute the rendering gradient



Task: compute the rendering gradient

image I






θ : object translation

θ : vertex position


θ : camera rotation

gradient $\frac{\partial I}{\partial \theta}$




Challenge: visibility



both camera visibility and shadow matter




Goal: correct, general, differentiable, and physically-based rendering



camera visibility



glossy
reflection



shadow




mirror
reflection

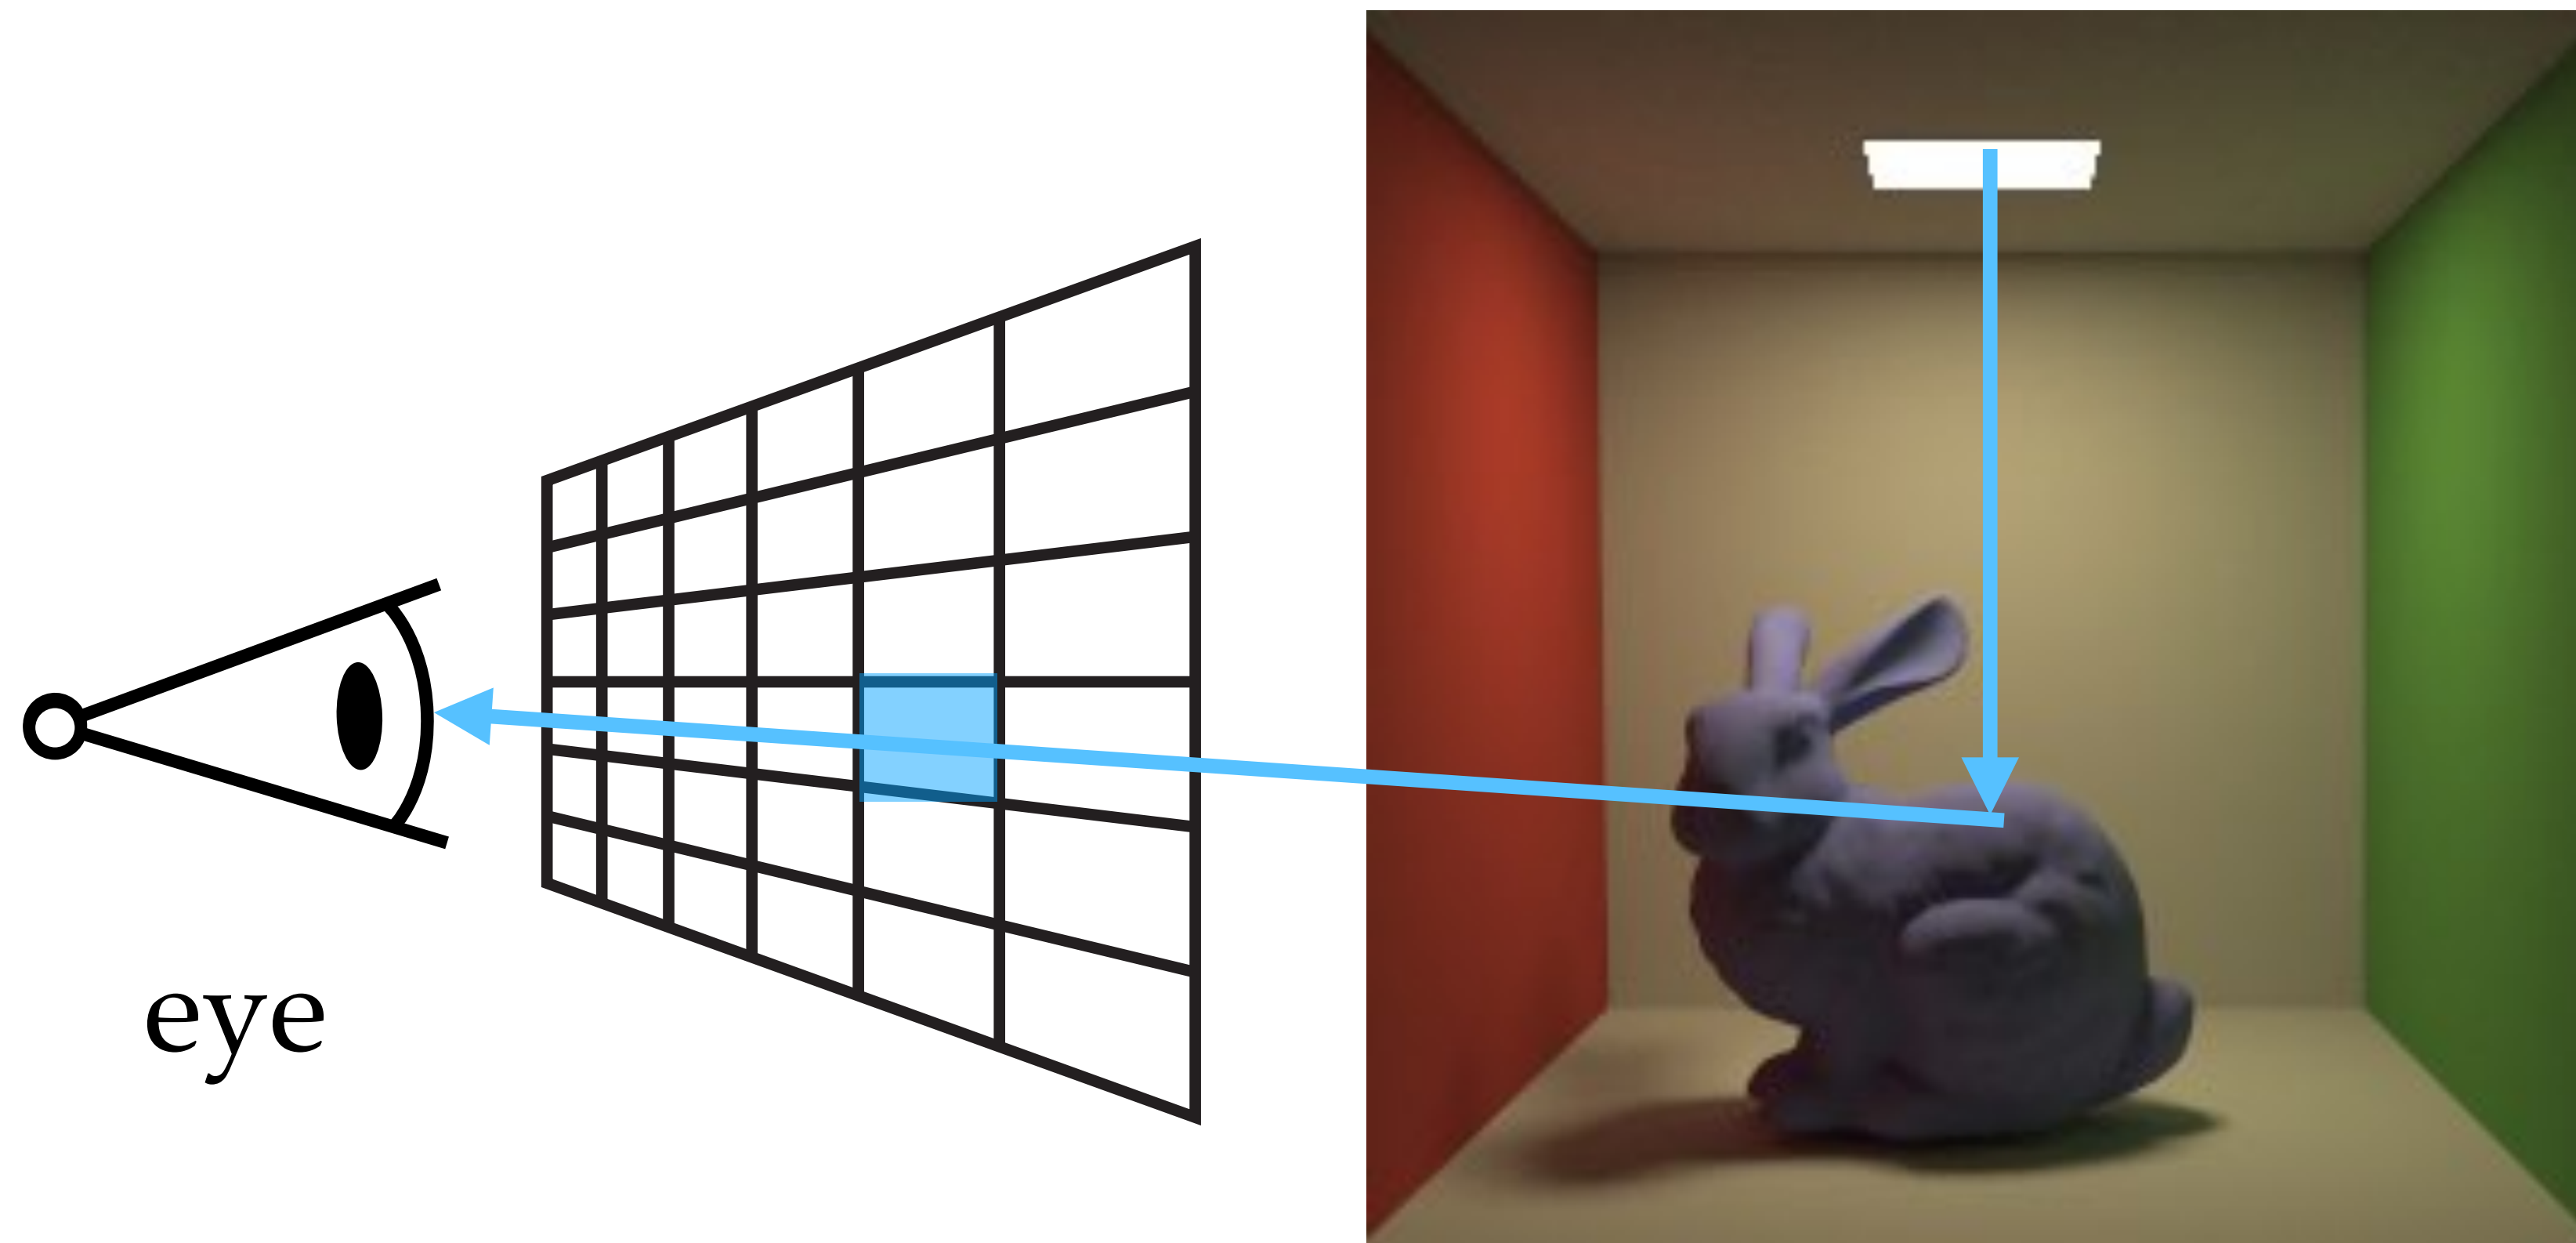


global
illumination

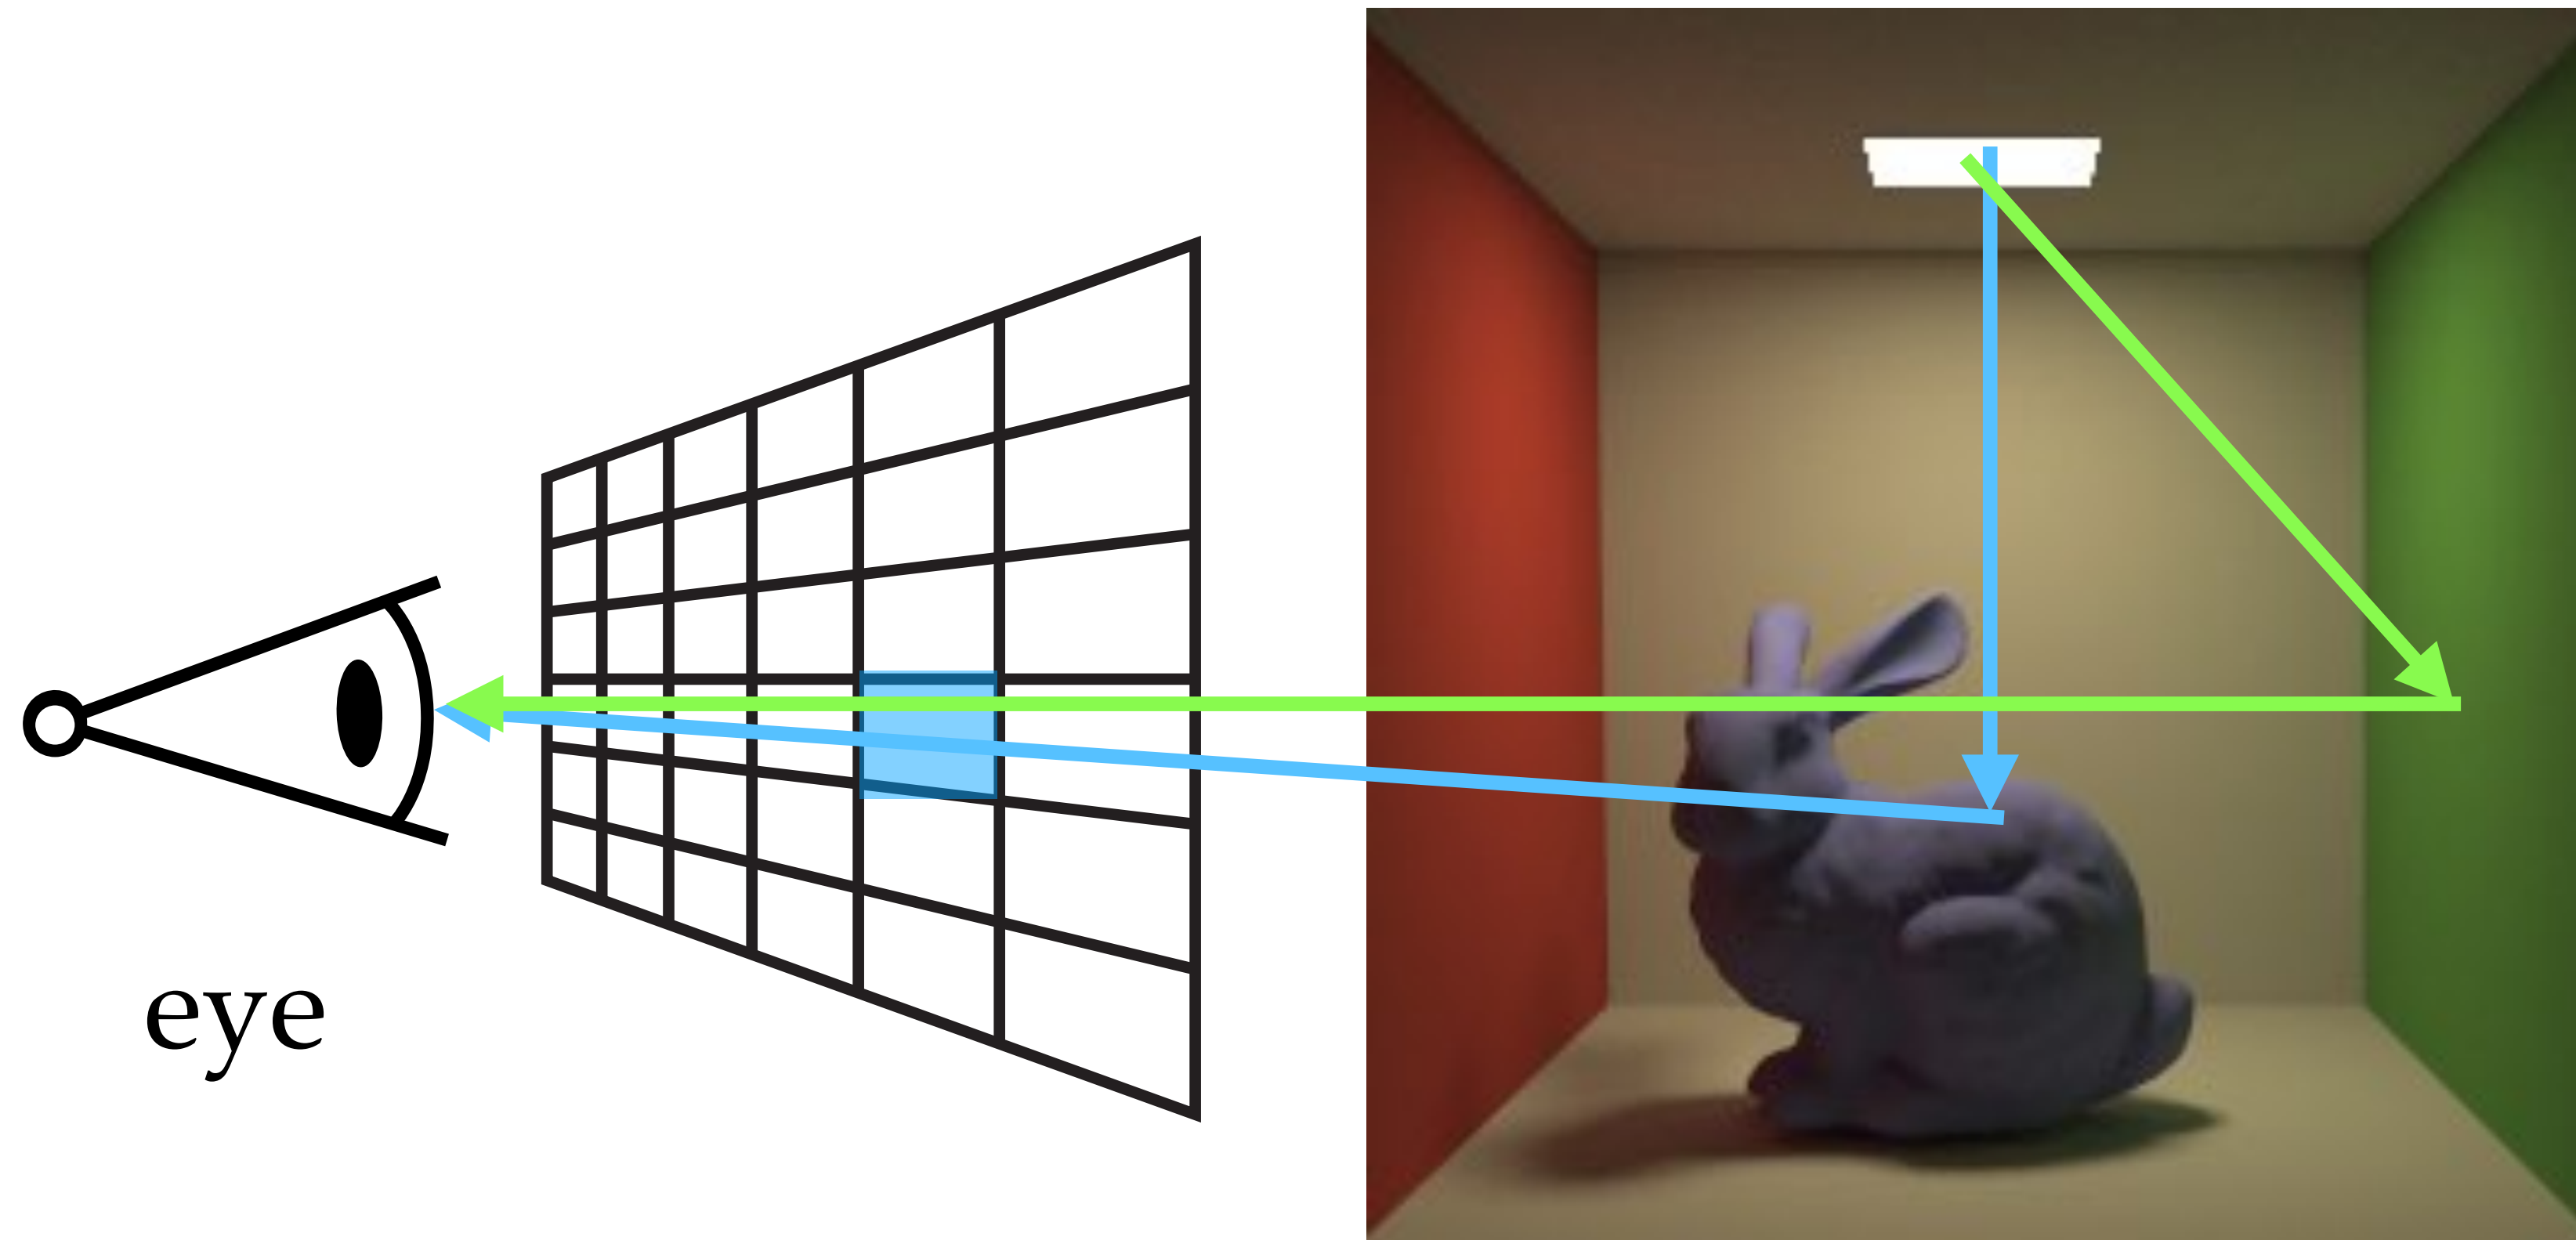
Renderers sample light paths




Renderers sample light paths




Renderers sample light paths




Renderers sample light paths




Renderers sample light paths



Renderers sample light paths

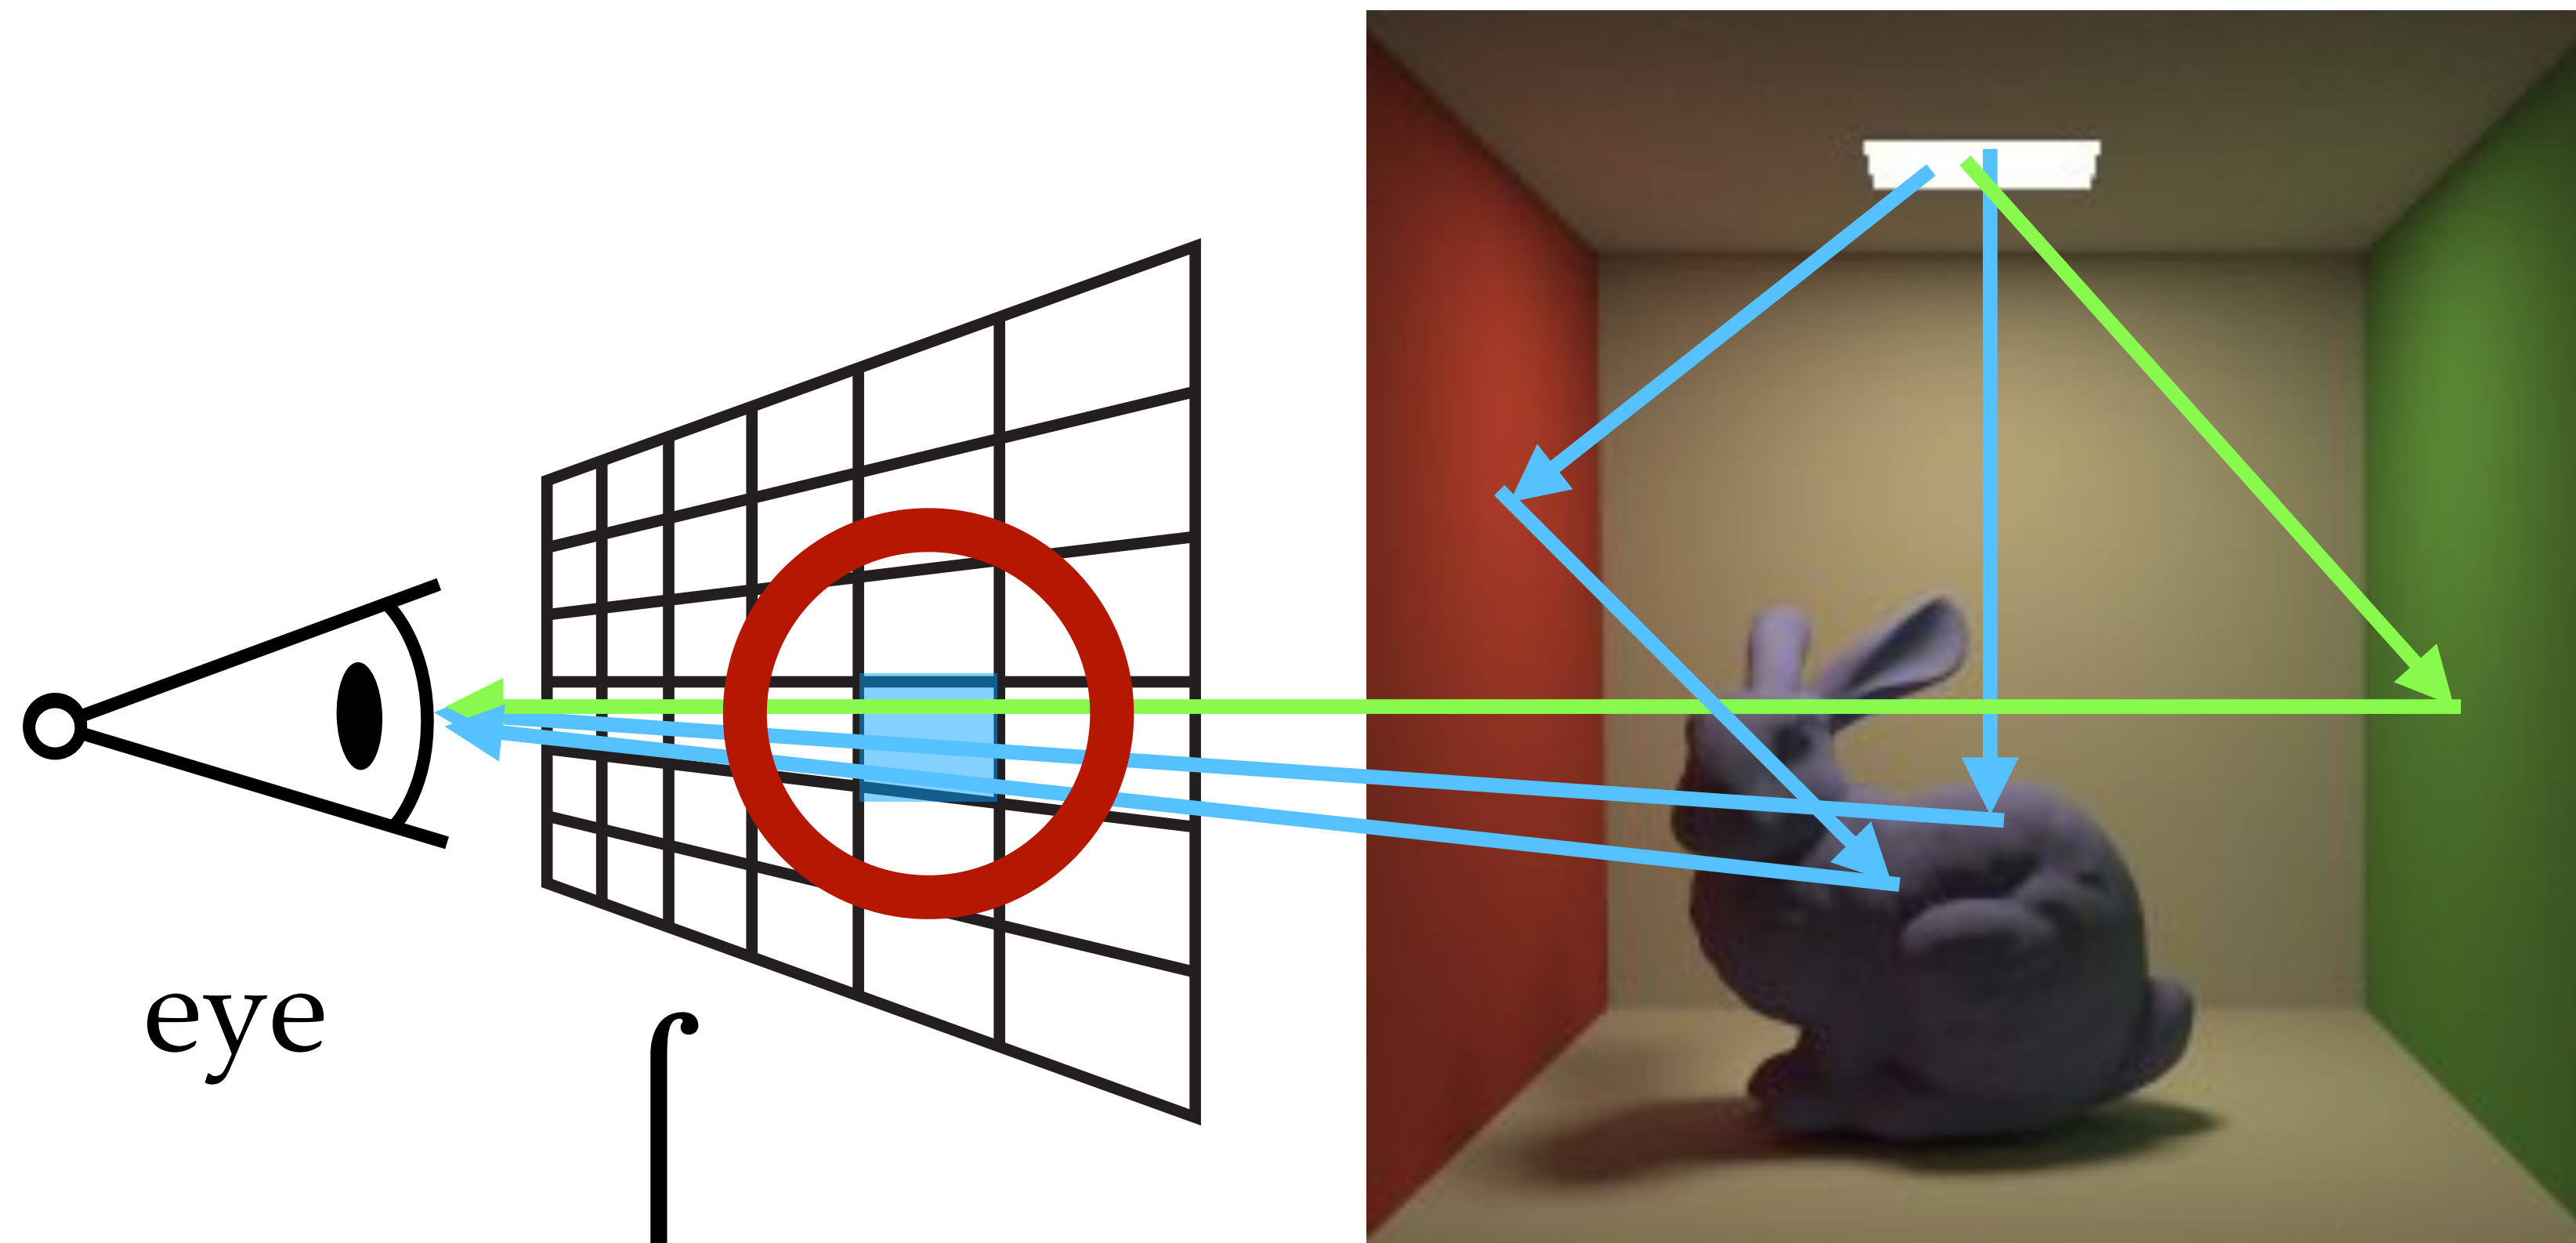


eye




key observation: the **integrand** is discontinuous, but the **integral** is differentiable

Renderers sample light paths





eye



key observation: the **integrand** is discontinuous, but the **integral** is differentiable



Toy example

single pixel






rendering = sample integral

single pixel



Differential of shape parameters =
boundary changes

single pixel





=the change of pixel color after we scale the shapes

Area sampling misses the boundaries


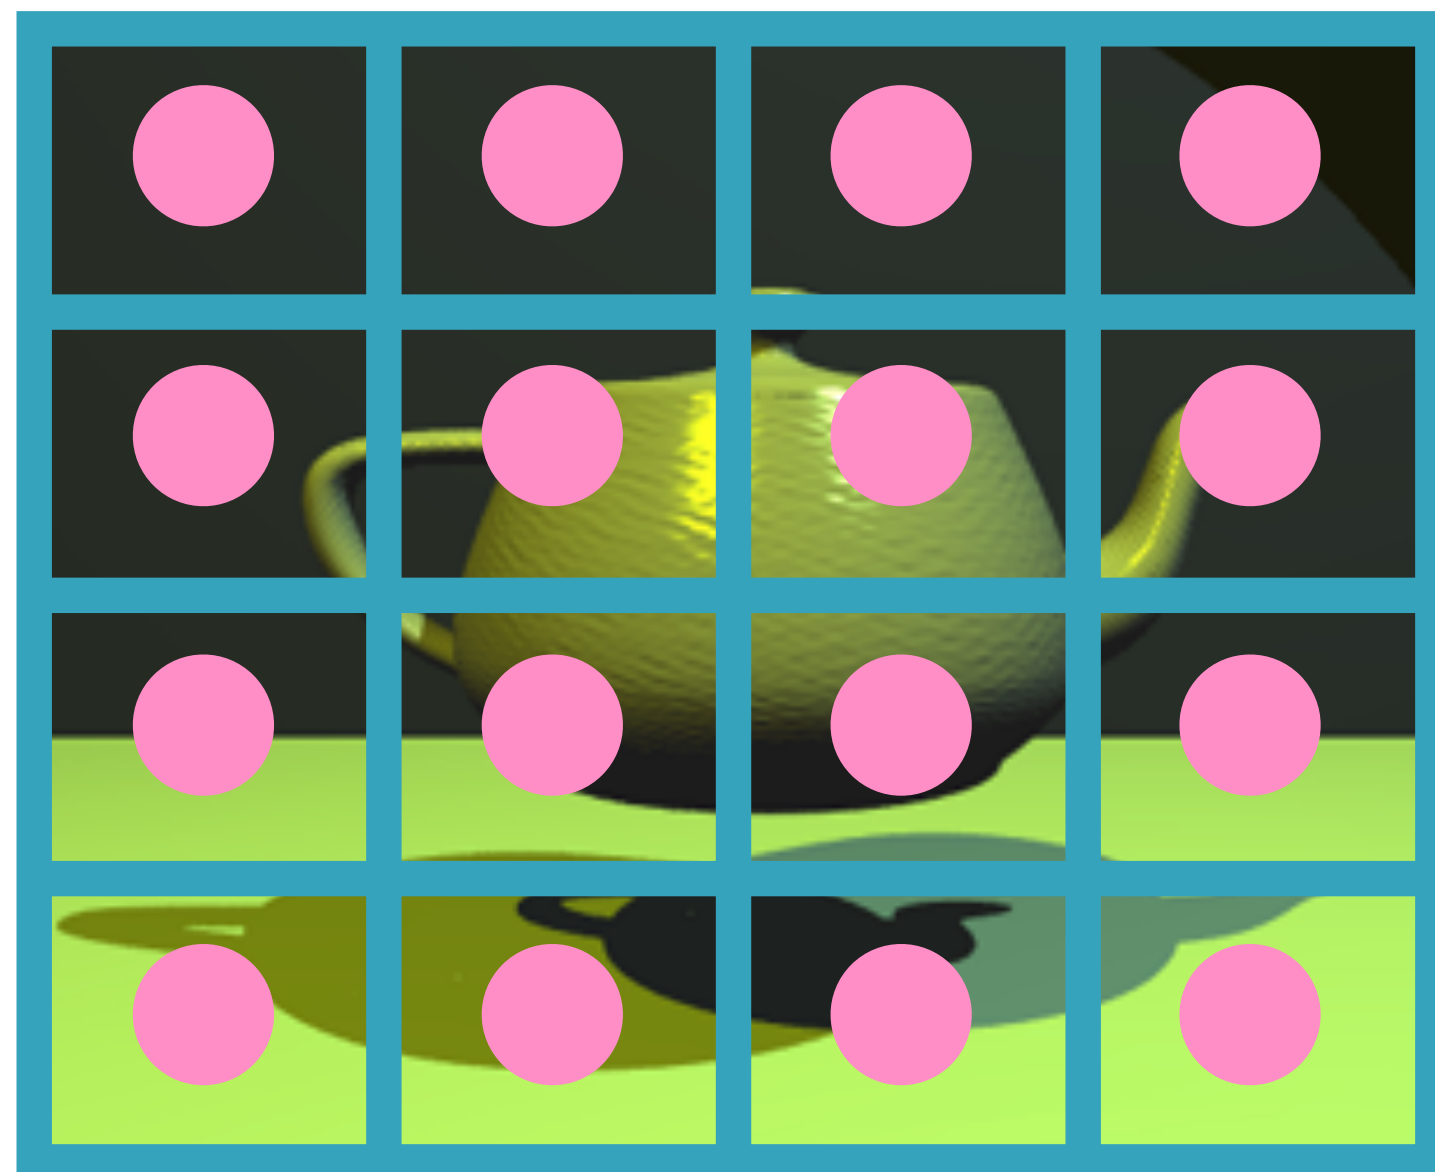
automatic differentiation can't compute correct gradients!

single pixel





Existing techniques are all based on area sampling

need a fundamentally different method



Our key idea: integrate over the boundaries

single pixel



Mathematically

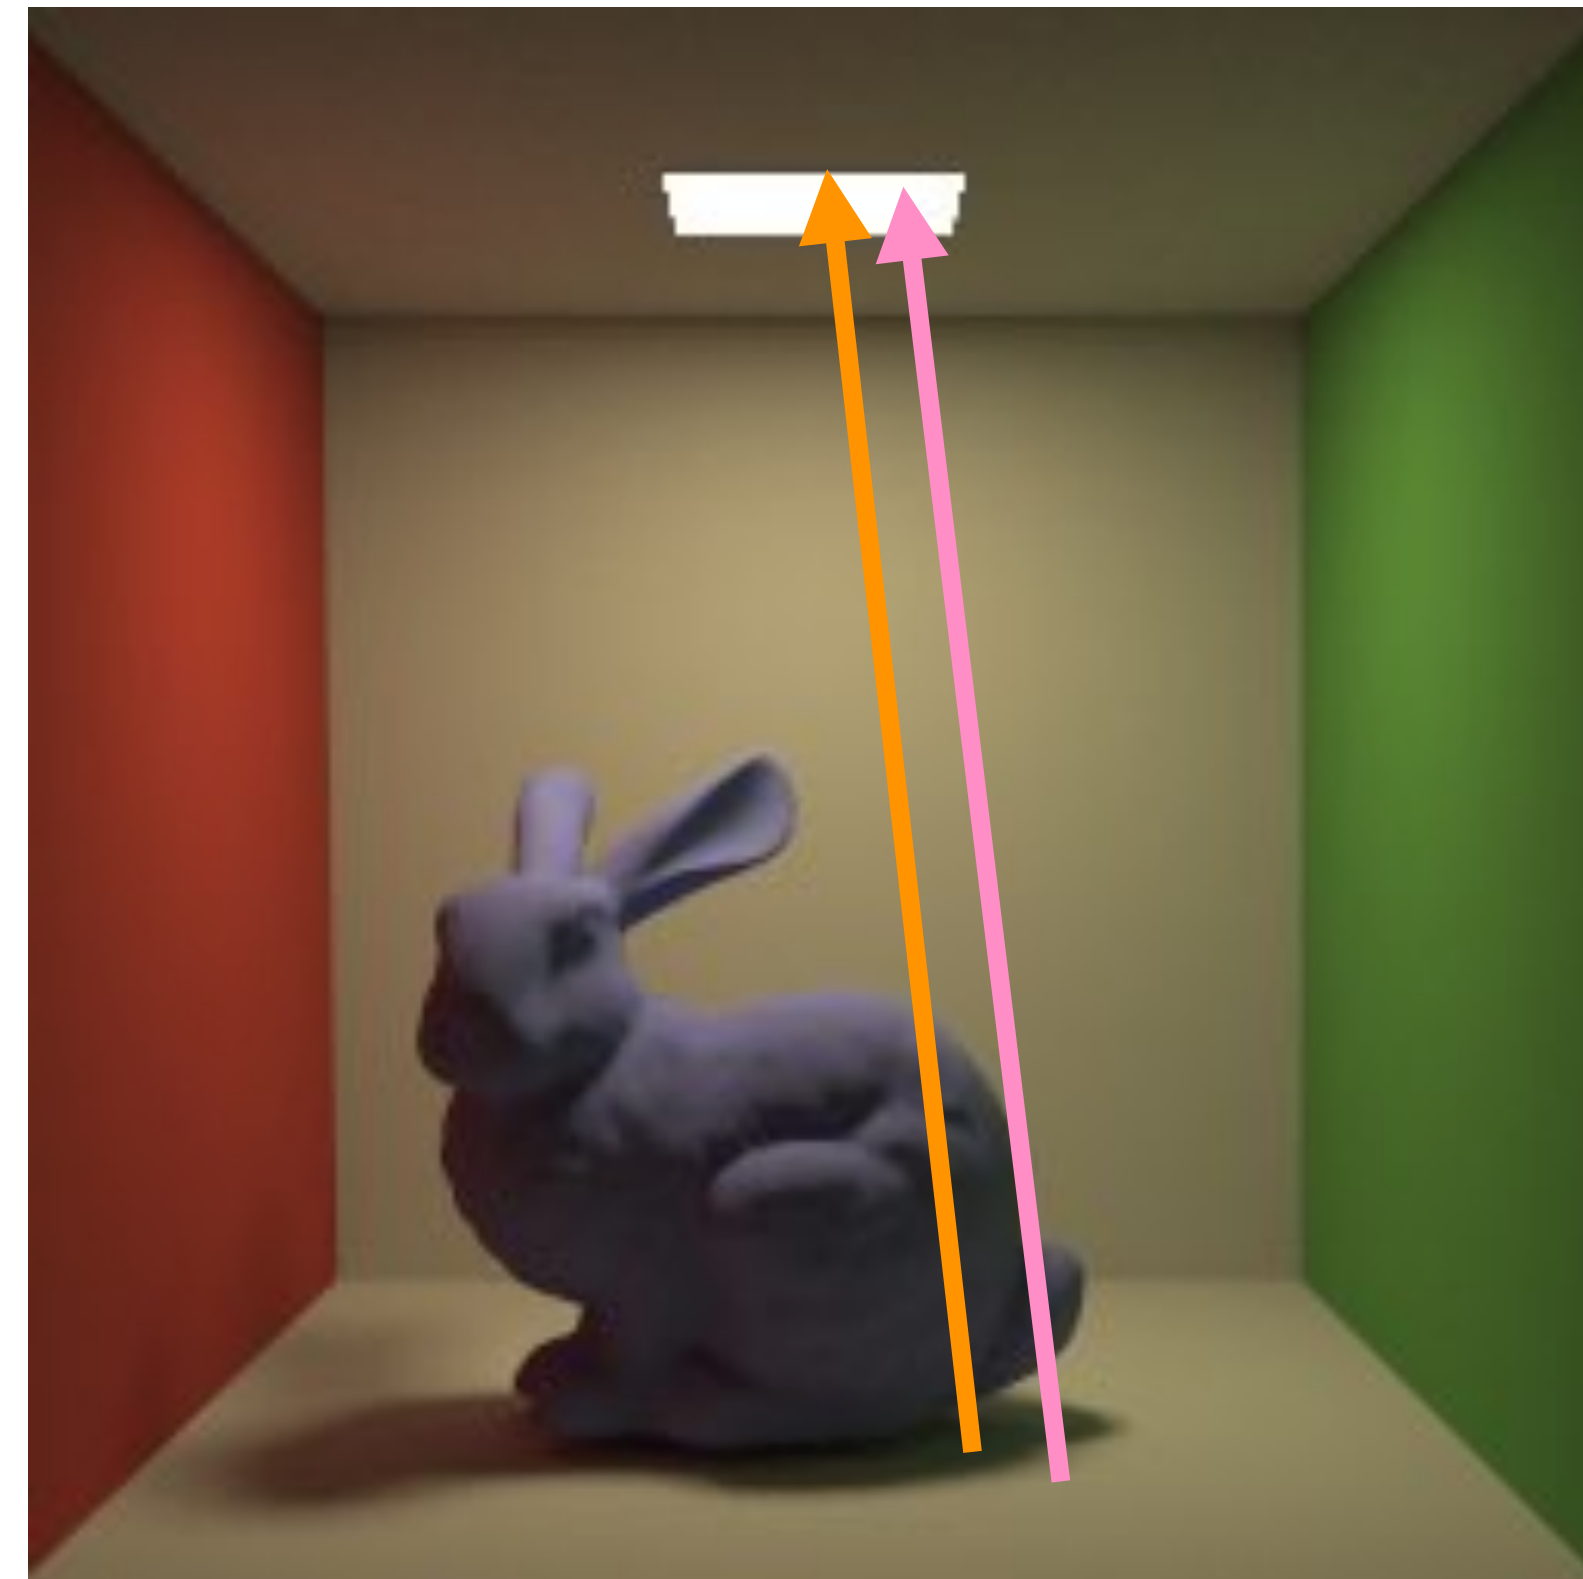

$$\nabla \iint \boxed{\text{blue triangle and red rectangle}} = \iint_{\text{area}} \nabla \boxed{\text{blue triangle and red rectangle with yellow dots}} + \int_{\text{boundary}} \boxed{\text{blue triangle and red rectangle with pink dots}}$$

The diagram illustrates the divergence theorem. On the left, a downward-pointing triangle symbol ∇ is followed by a double integral \iint over a square box containing a blue triangle and a red rectangle. This is set equal to a double integral \iint over the area of the box, with a downward-pointing triangle symbol ∇ inside, plus a single integral \int over the boundary of the box. The first box on the right contains yellow dots inside the shapes, and the second box on the right contains pink dots on the boundaries of the shapes.

derived through Dirac delta or Reynolds transport theorem

Generalize to shadow & interreflection

integrating a different domain



Challenge: scalability

sampling triangle edges requires different acceleration data structures



0.15M triangles



0.8M triangles



4.8M triangles

We can convert boundary integrals back to area integrals (divergence theorem)

we can then reuse data structures from traditional rendering


$$\int_{\text{boundary}} = \iint_{\text{area}} \nabla \cdot$$

boundary

area

$$\nabla \iint \boxed{\text{blue triangle red square}} = \iint \nabla \boxed{\text{blue triangle red square with yellow dots}} + \int \boxed{\text{blue triangle red square with pink dots}}$$


Applications



inverse rendering






analyzing vision systems



"A drawing of a cat".


bridging vector and raster graphics



inverse shader design

Inverse rendering

differentiable rendering unifies 3D reconstruction



- geometry
- material
- light
- camera

Inverse rendering with real photos

optimize camera pose, light emission & materials



initial guess



target



reconstructed

Inverse rendering with real photos

optimize camera pose, light emission & materials



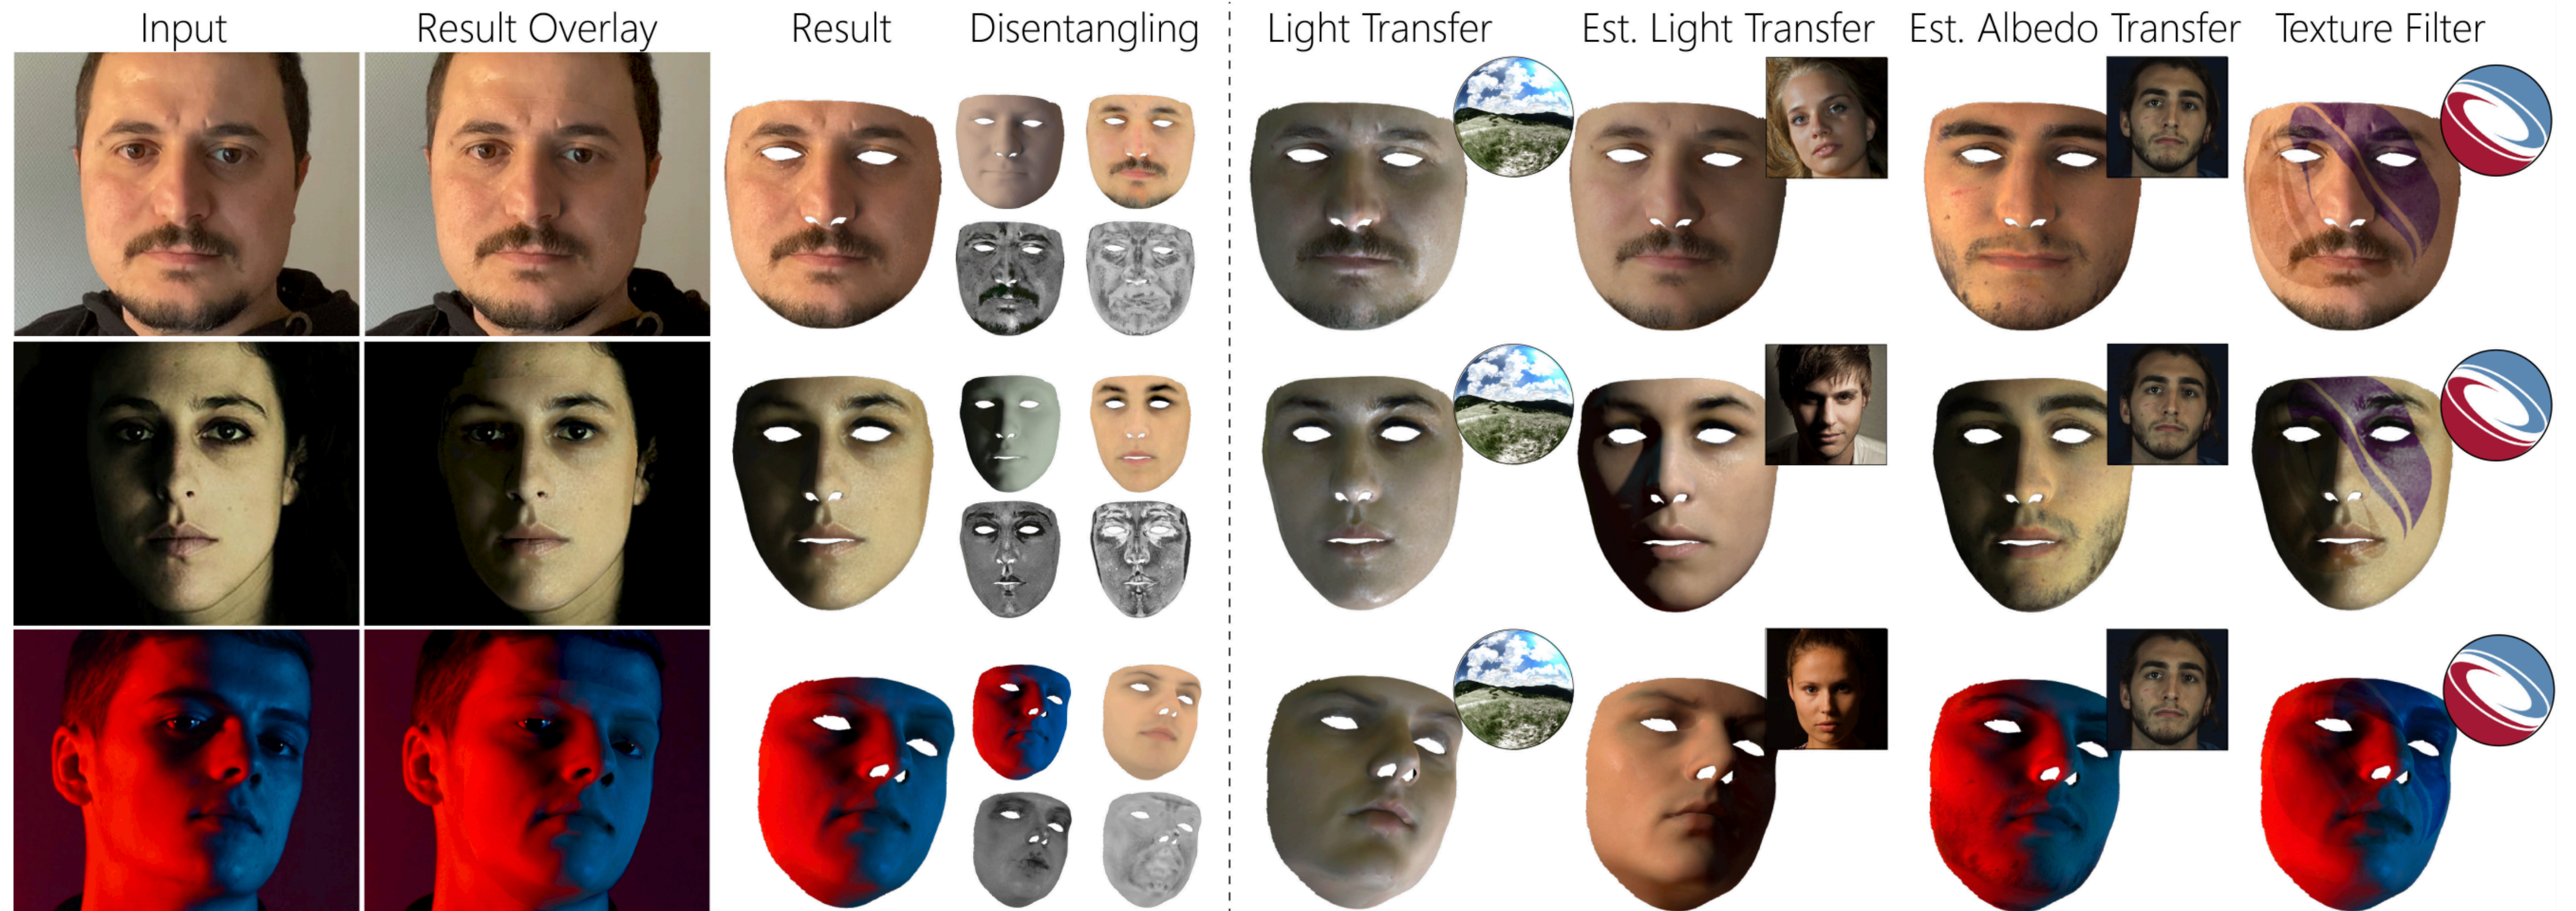
optimization video



target

Face reconstruction

from Dib et al., using our differentiable renderer (not my work!)




Multi-view 3D reconstruction

from Goel et al., using our differentiable renderer (not my work!)

Two Input Views



Our Mesh + SVBRDF

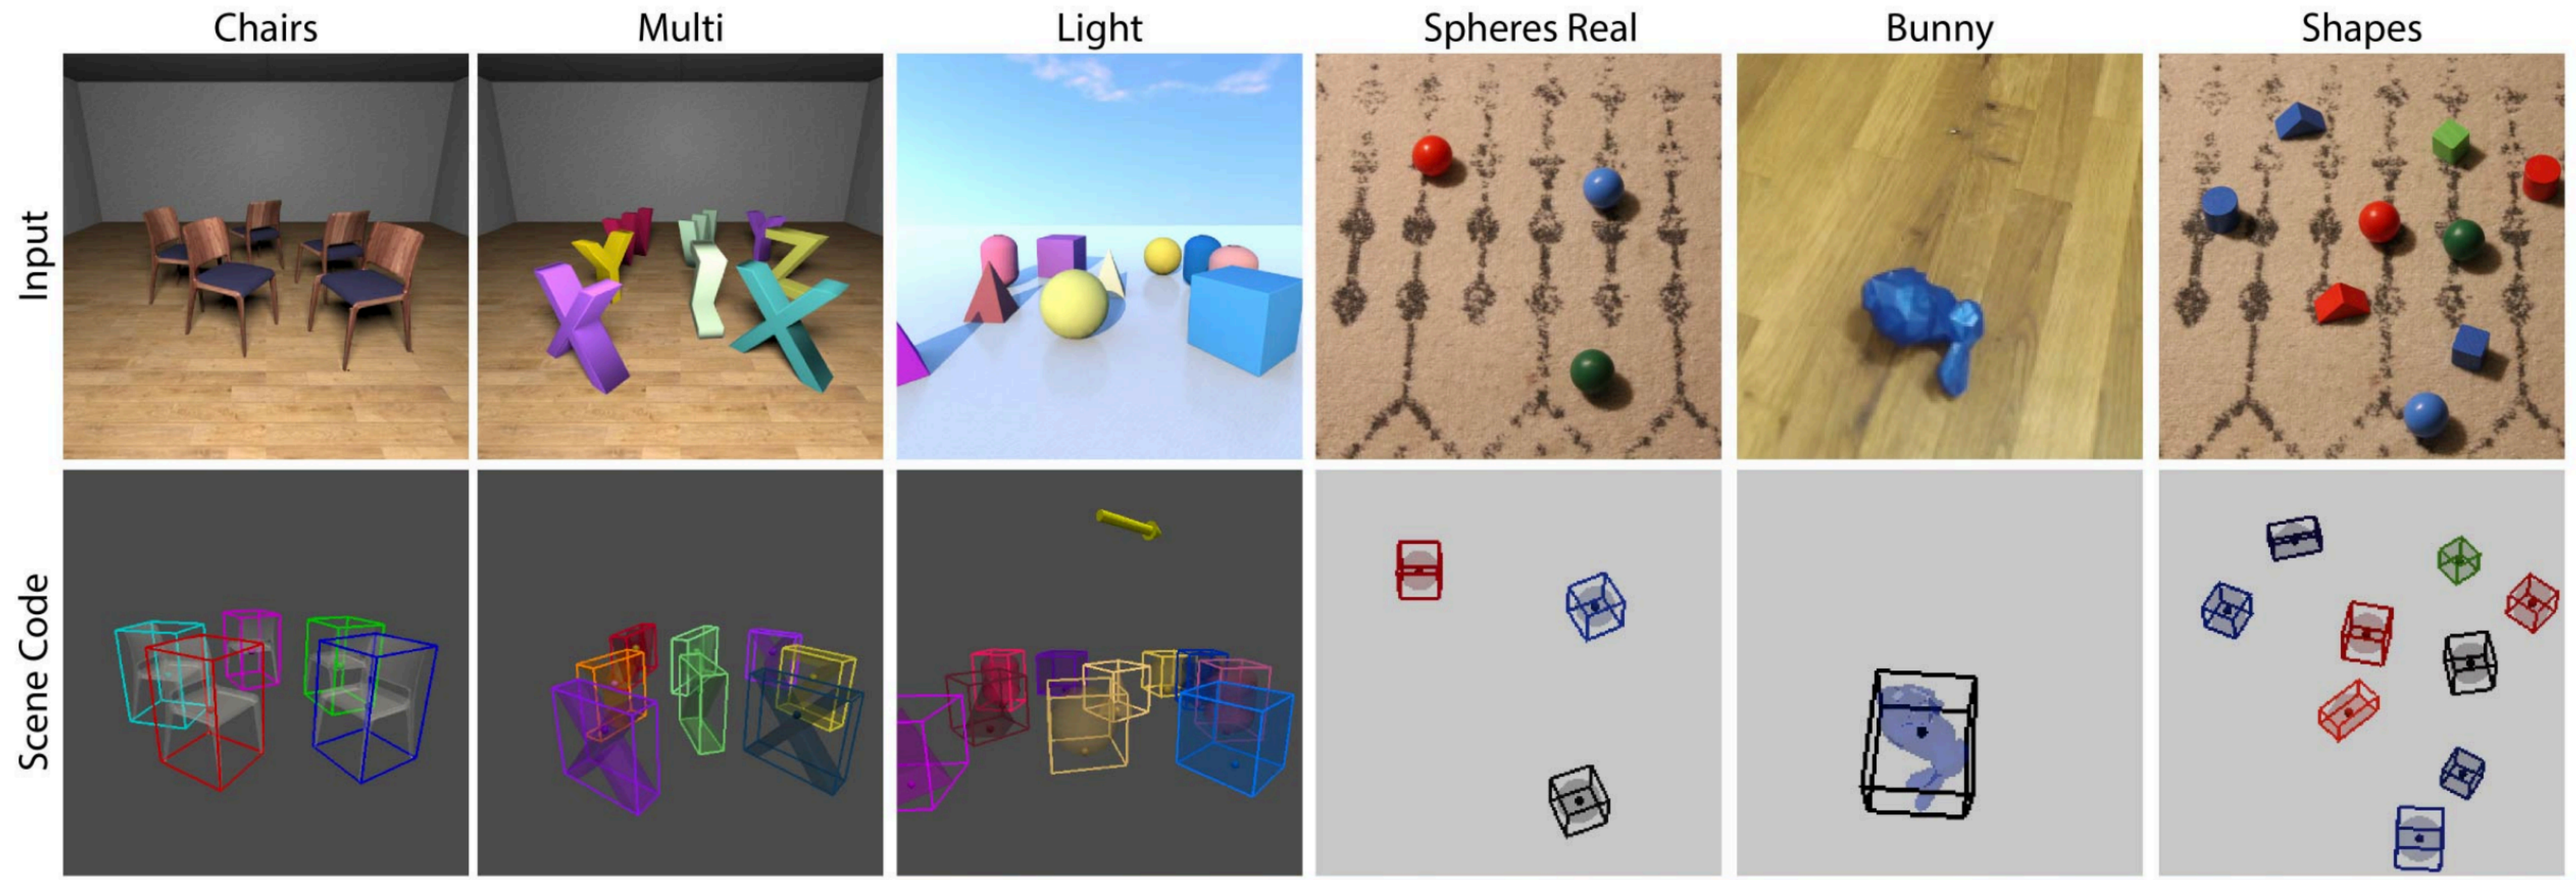


Ours (re-lit)

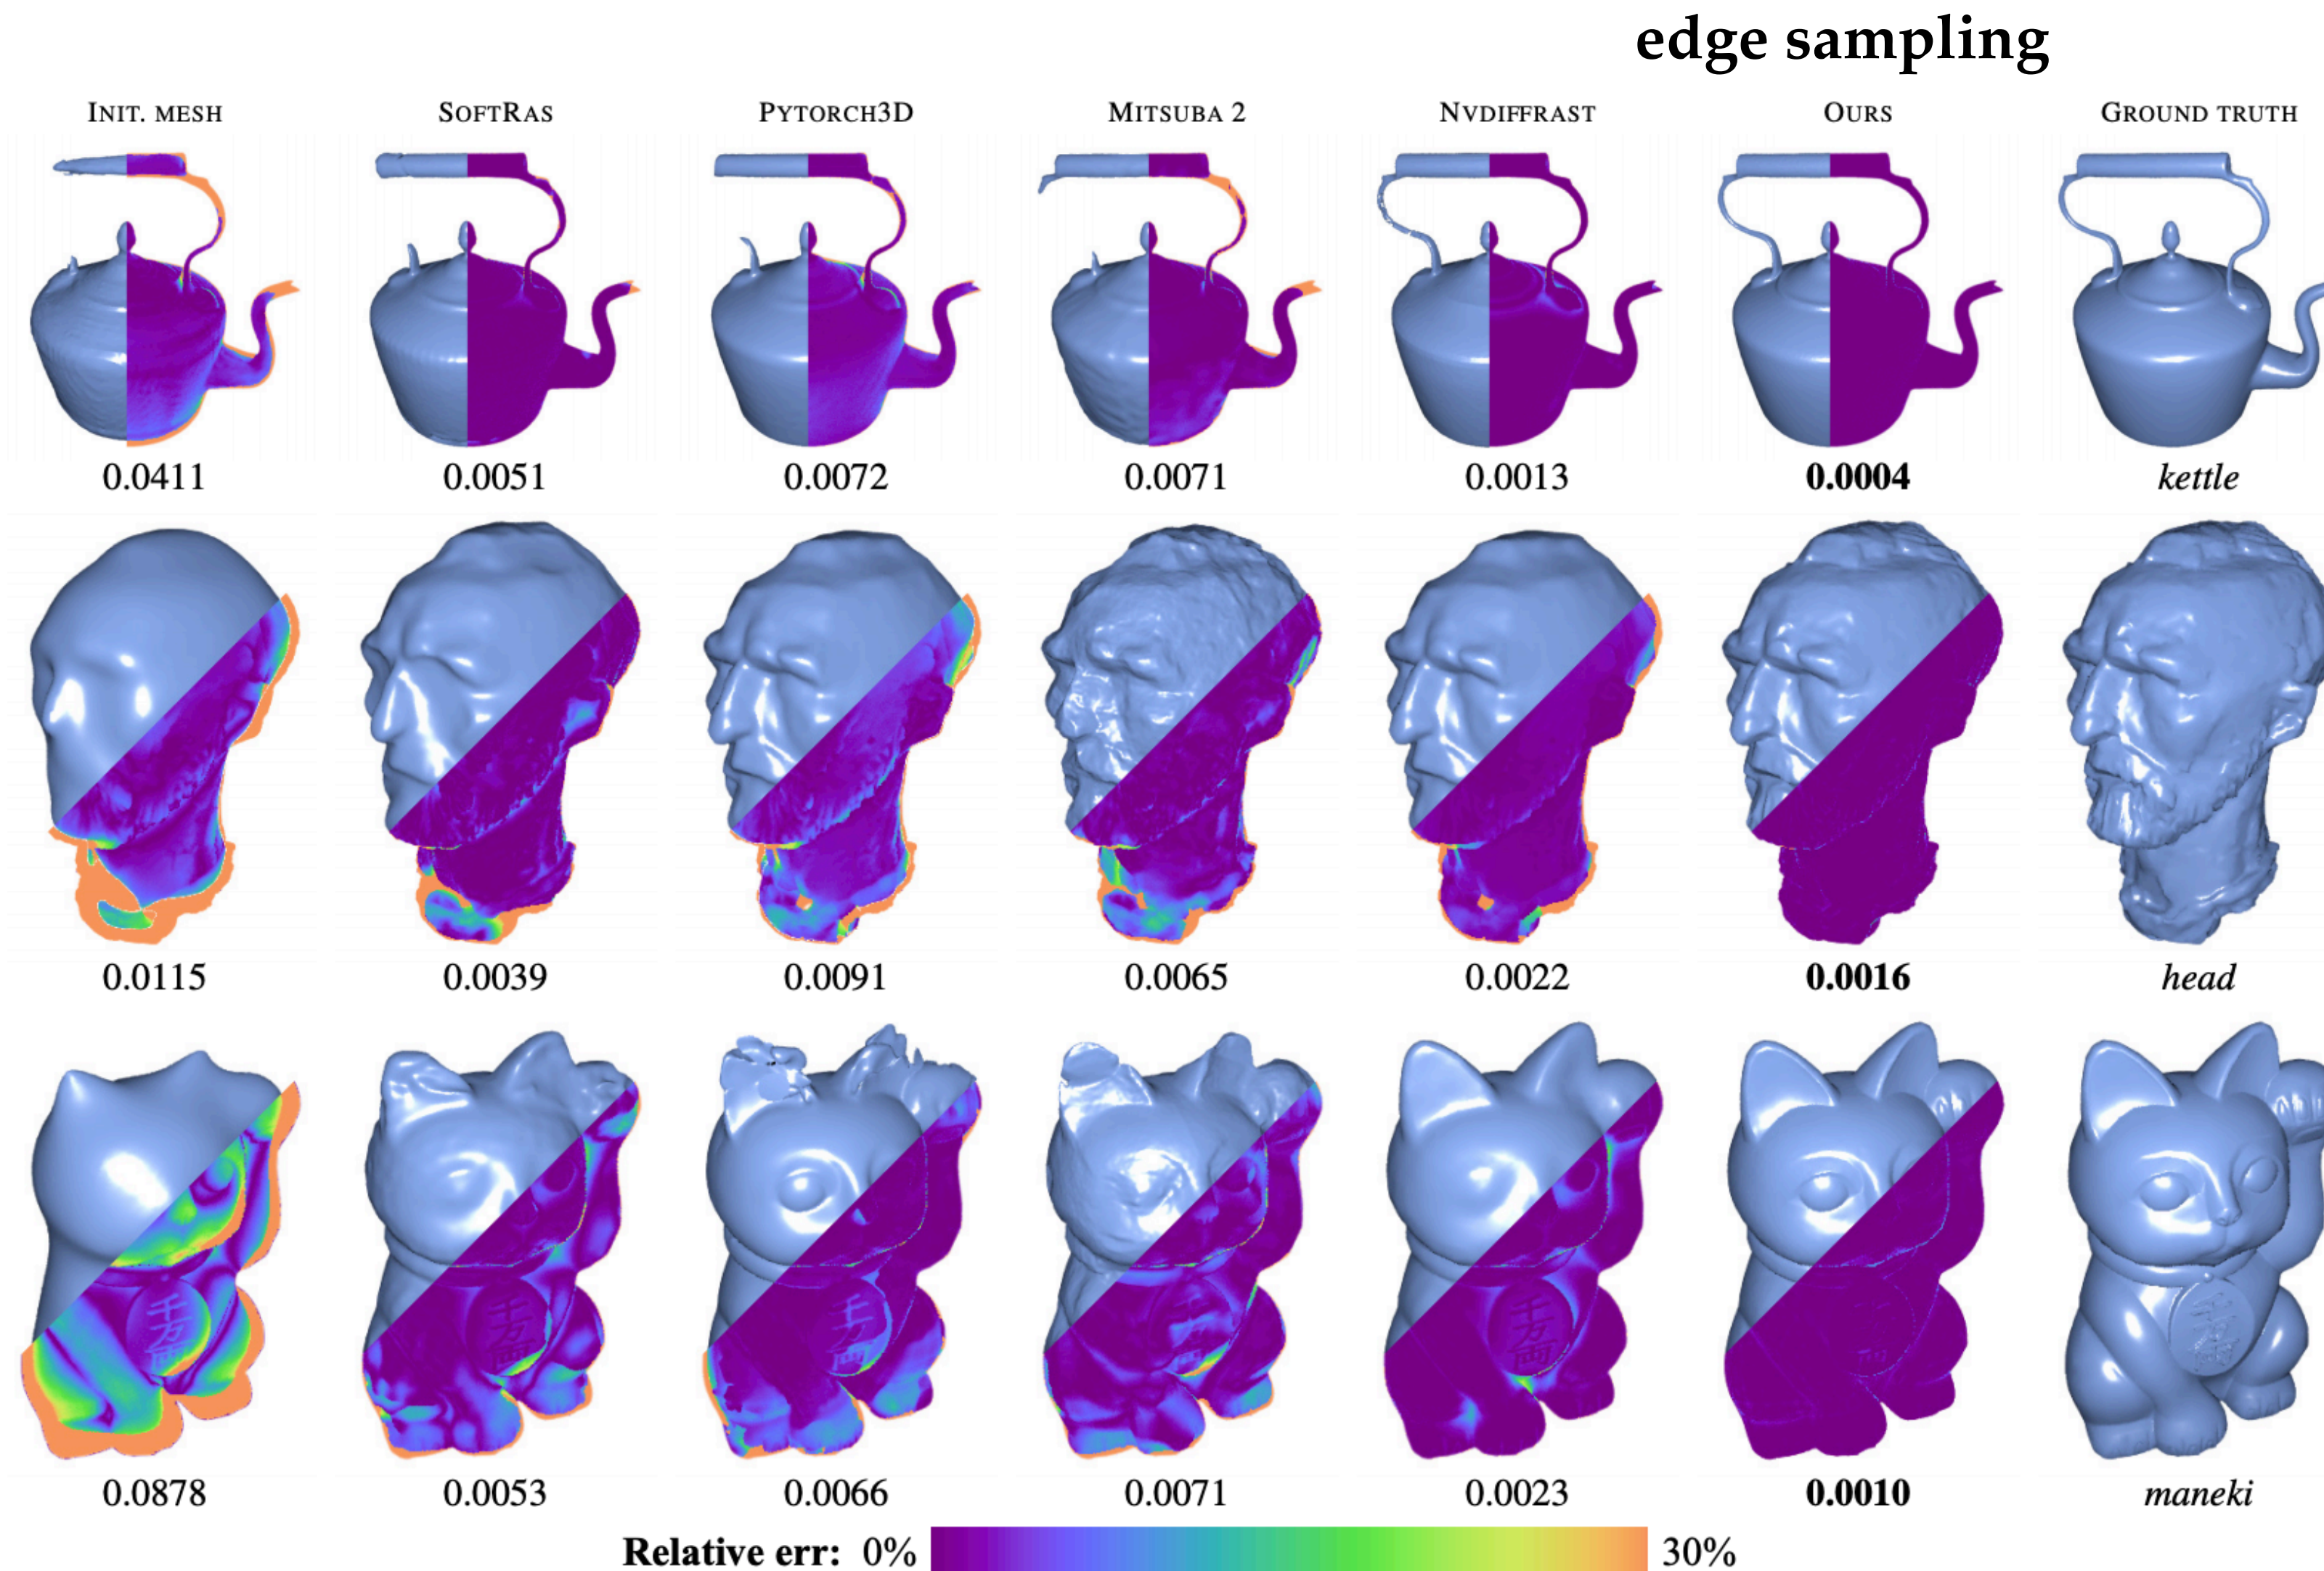


Learning to generate 3D scenes

from Griffiths et al., using our differentiable renderer (not my work!)




Having correct gradients matters



from Luan et al. 2021
(not my work!)

Better lighting model = better reconstruction

rendering reflectance rendering reflectance



Maier 2017

Ours

photo

Augmented reality application

3D reconstruction is useful for editing images



real photo




edited

Analyzing the behavior of vision systems

differentiable rendering allows us to ask 3D questions

Analyzing the behavior of vision systems


differentiable rendering allows us to ask 3D questions



Analyzing the behavior of vision systems

differentiable rendering allows us to ask 3D questions


- lighting



Analyzing the behavior of vision systems

differentiable rendering allows us to ask 3D questions

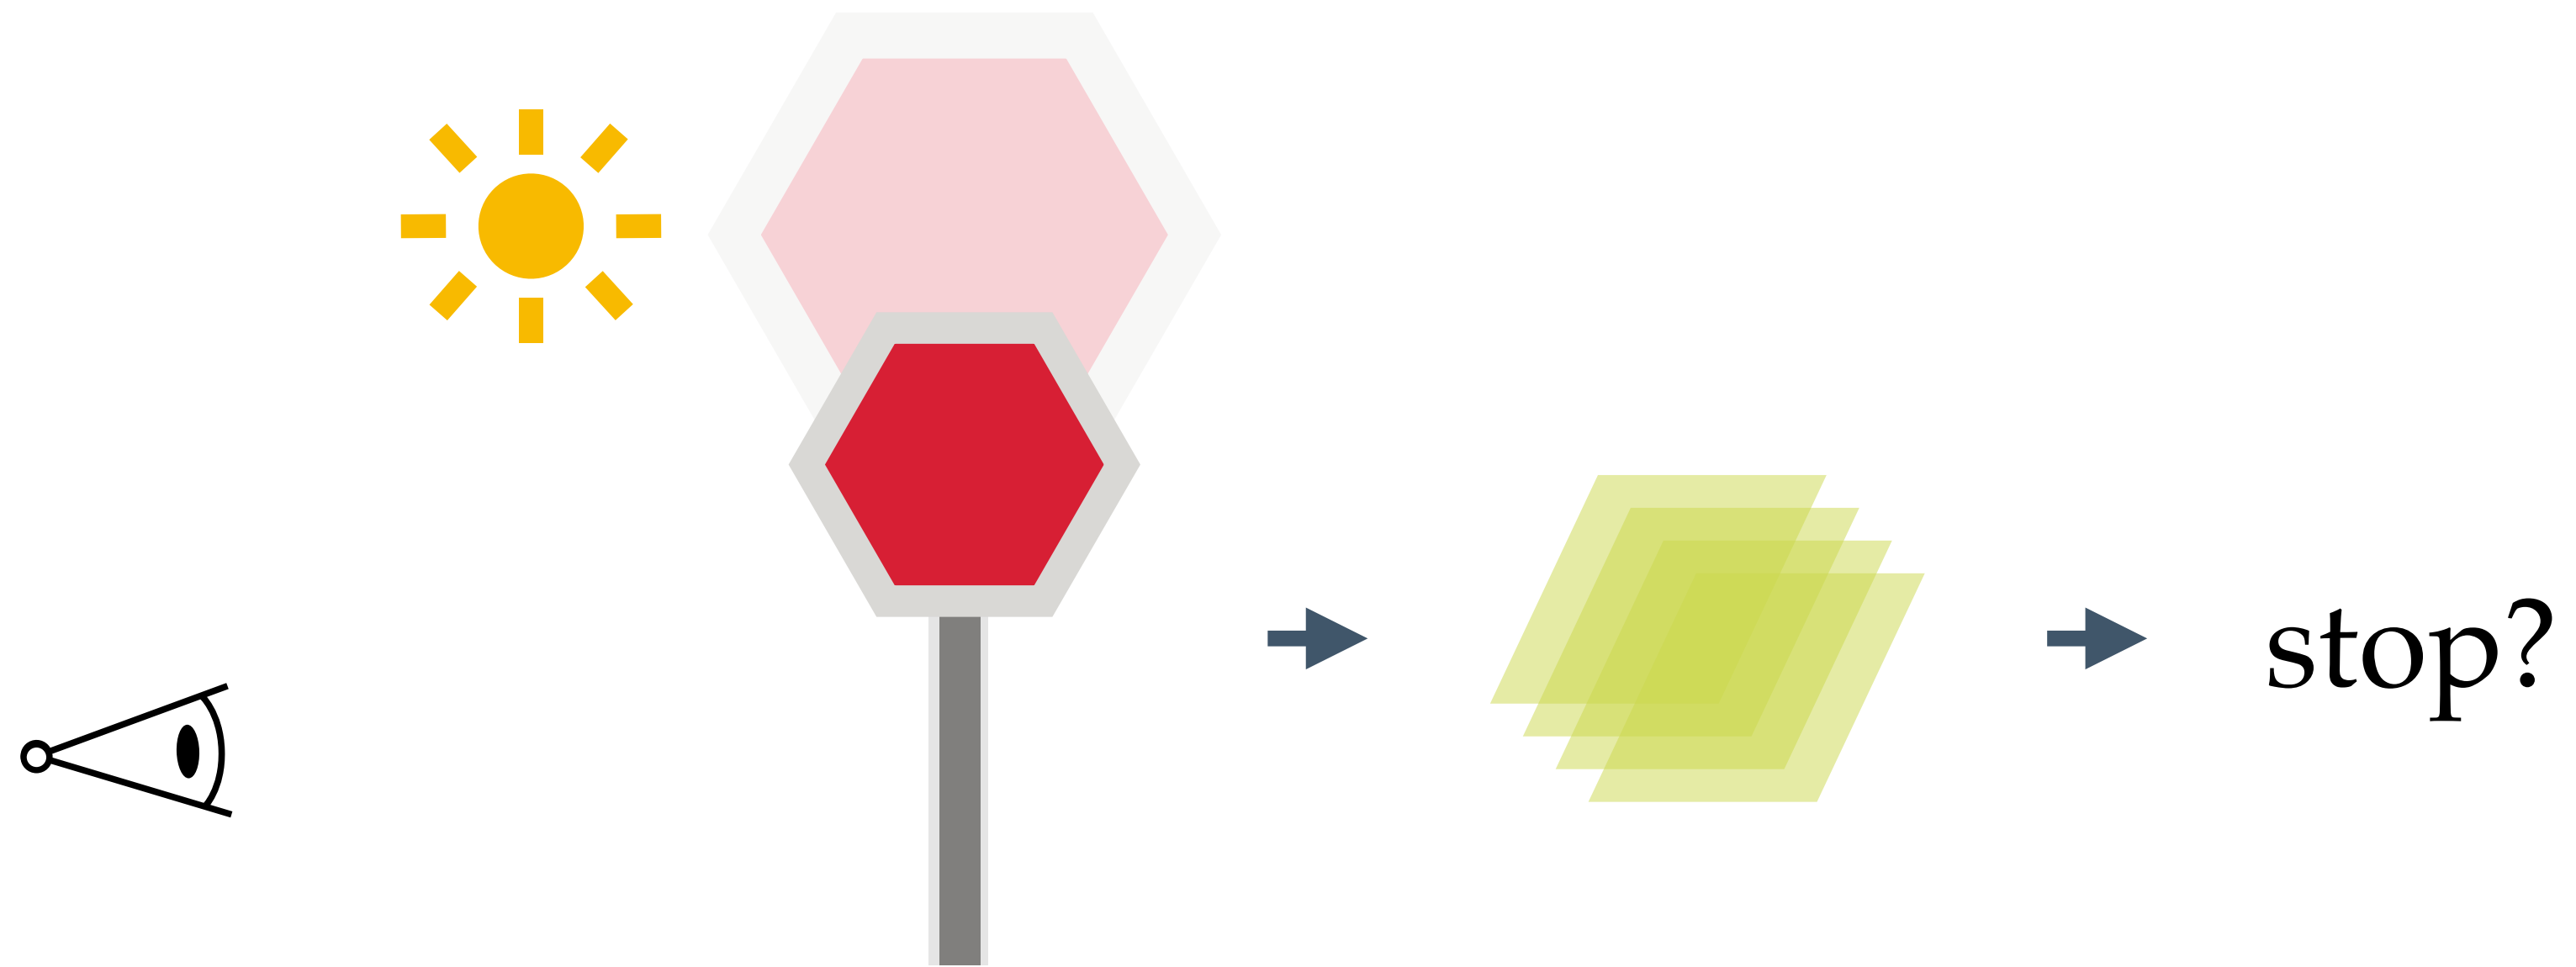
- lighting
- view



Analyzing the behavior of vision systems

differentiable rendering allows us to ask 3D questions


- lighting
- view
- shape



Analyzing the behavior of vision systems

differentiable rendering allows us to ask 3D questions

- lighting
- view
- shape
- material



3D adversarial examples

optimize for vertex position, camera pose,
light intensity, position



VGG 16:
53% street sign
6.7% handrail

3D adversarial examples

optimize for vertex position, camera pose,
light intensity, position



VGG 16:
53% street sign
6.7% handrail



5 iterations:
26.8% handrail
20.2% street sign

3D adversarial examples

optimize for vertex position, camera pose,
light intensity, position



VGG 16:
53% street sign
6.7% handrail




5 iterations:
26.8% handrail
20.2% street sign



25 iterations:
23.3% handrail
3.4% street sign

Bridging vector and raster graphics

differentiable rendering enables
neural network operations on vector graphics



Bridging vector and raster graphics

painterly rendering


ours



target

Bridging vector and raster graphics

image processing for vector graphics




Bridging vector and raster graphics

unsupervised learning for vector graphics




raster training data
(MNIST)



our vector output

Bridging vector and raster graphics

unsupervised learning for vector graphics



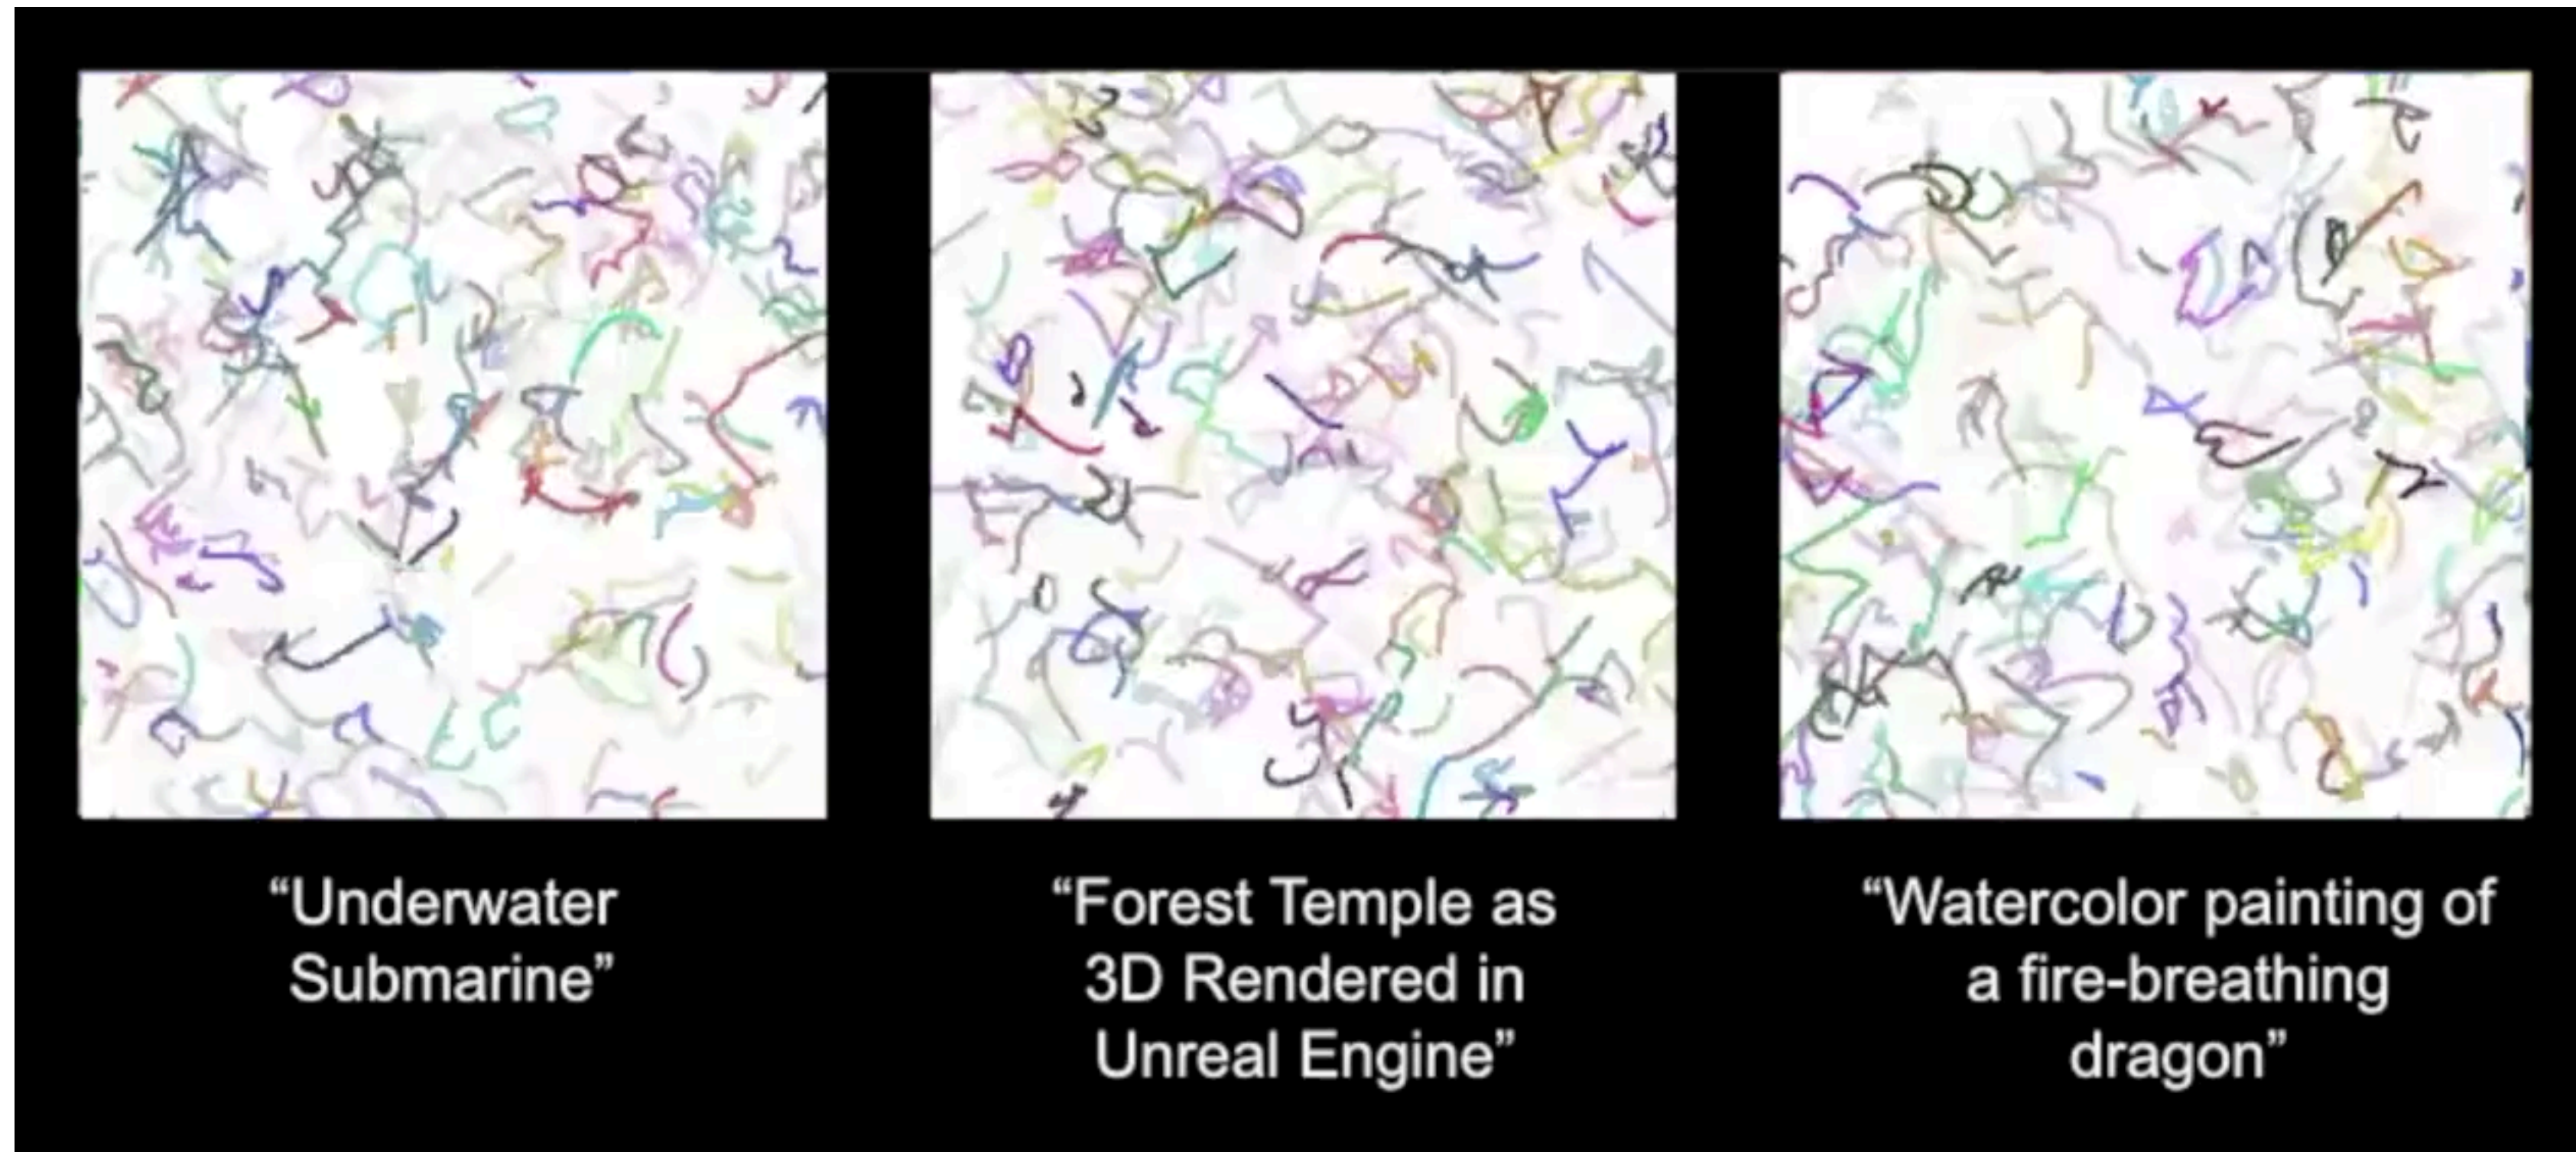
from Reddy et al.,
interpolating emojis using our renderer

(not my work!)

CLIPDraw [Frans et al.]

(not my work!)

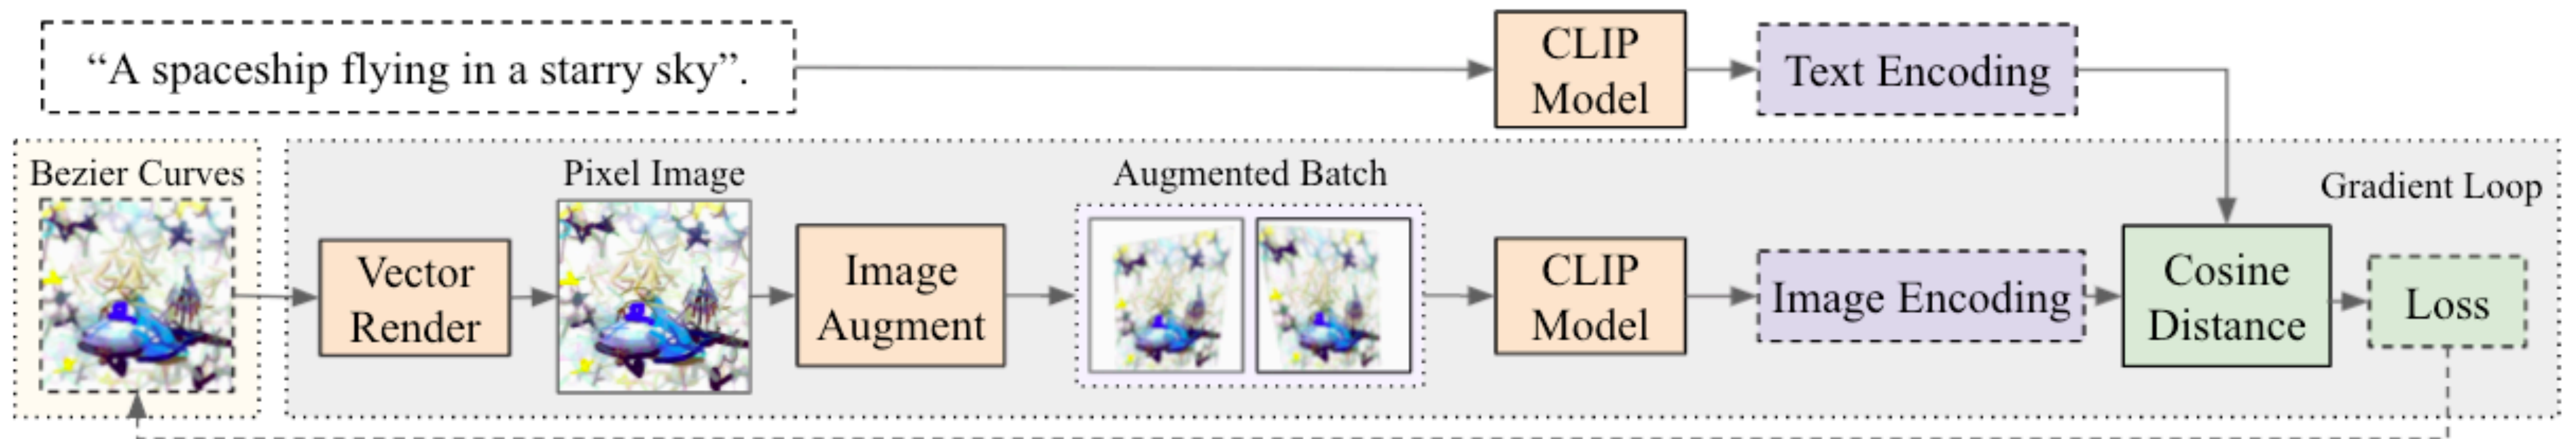
- text to drawing using CLIP & differentiable rasterization using our renderer



CLIPDraw [Frans et al.]

(not my work!)

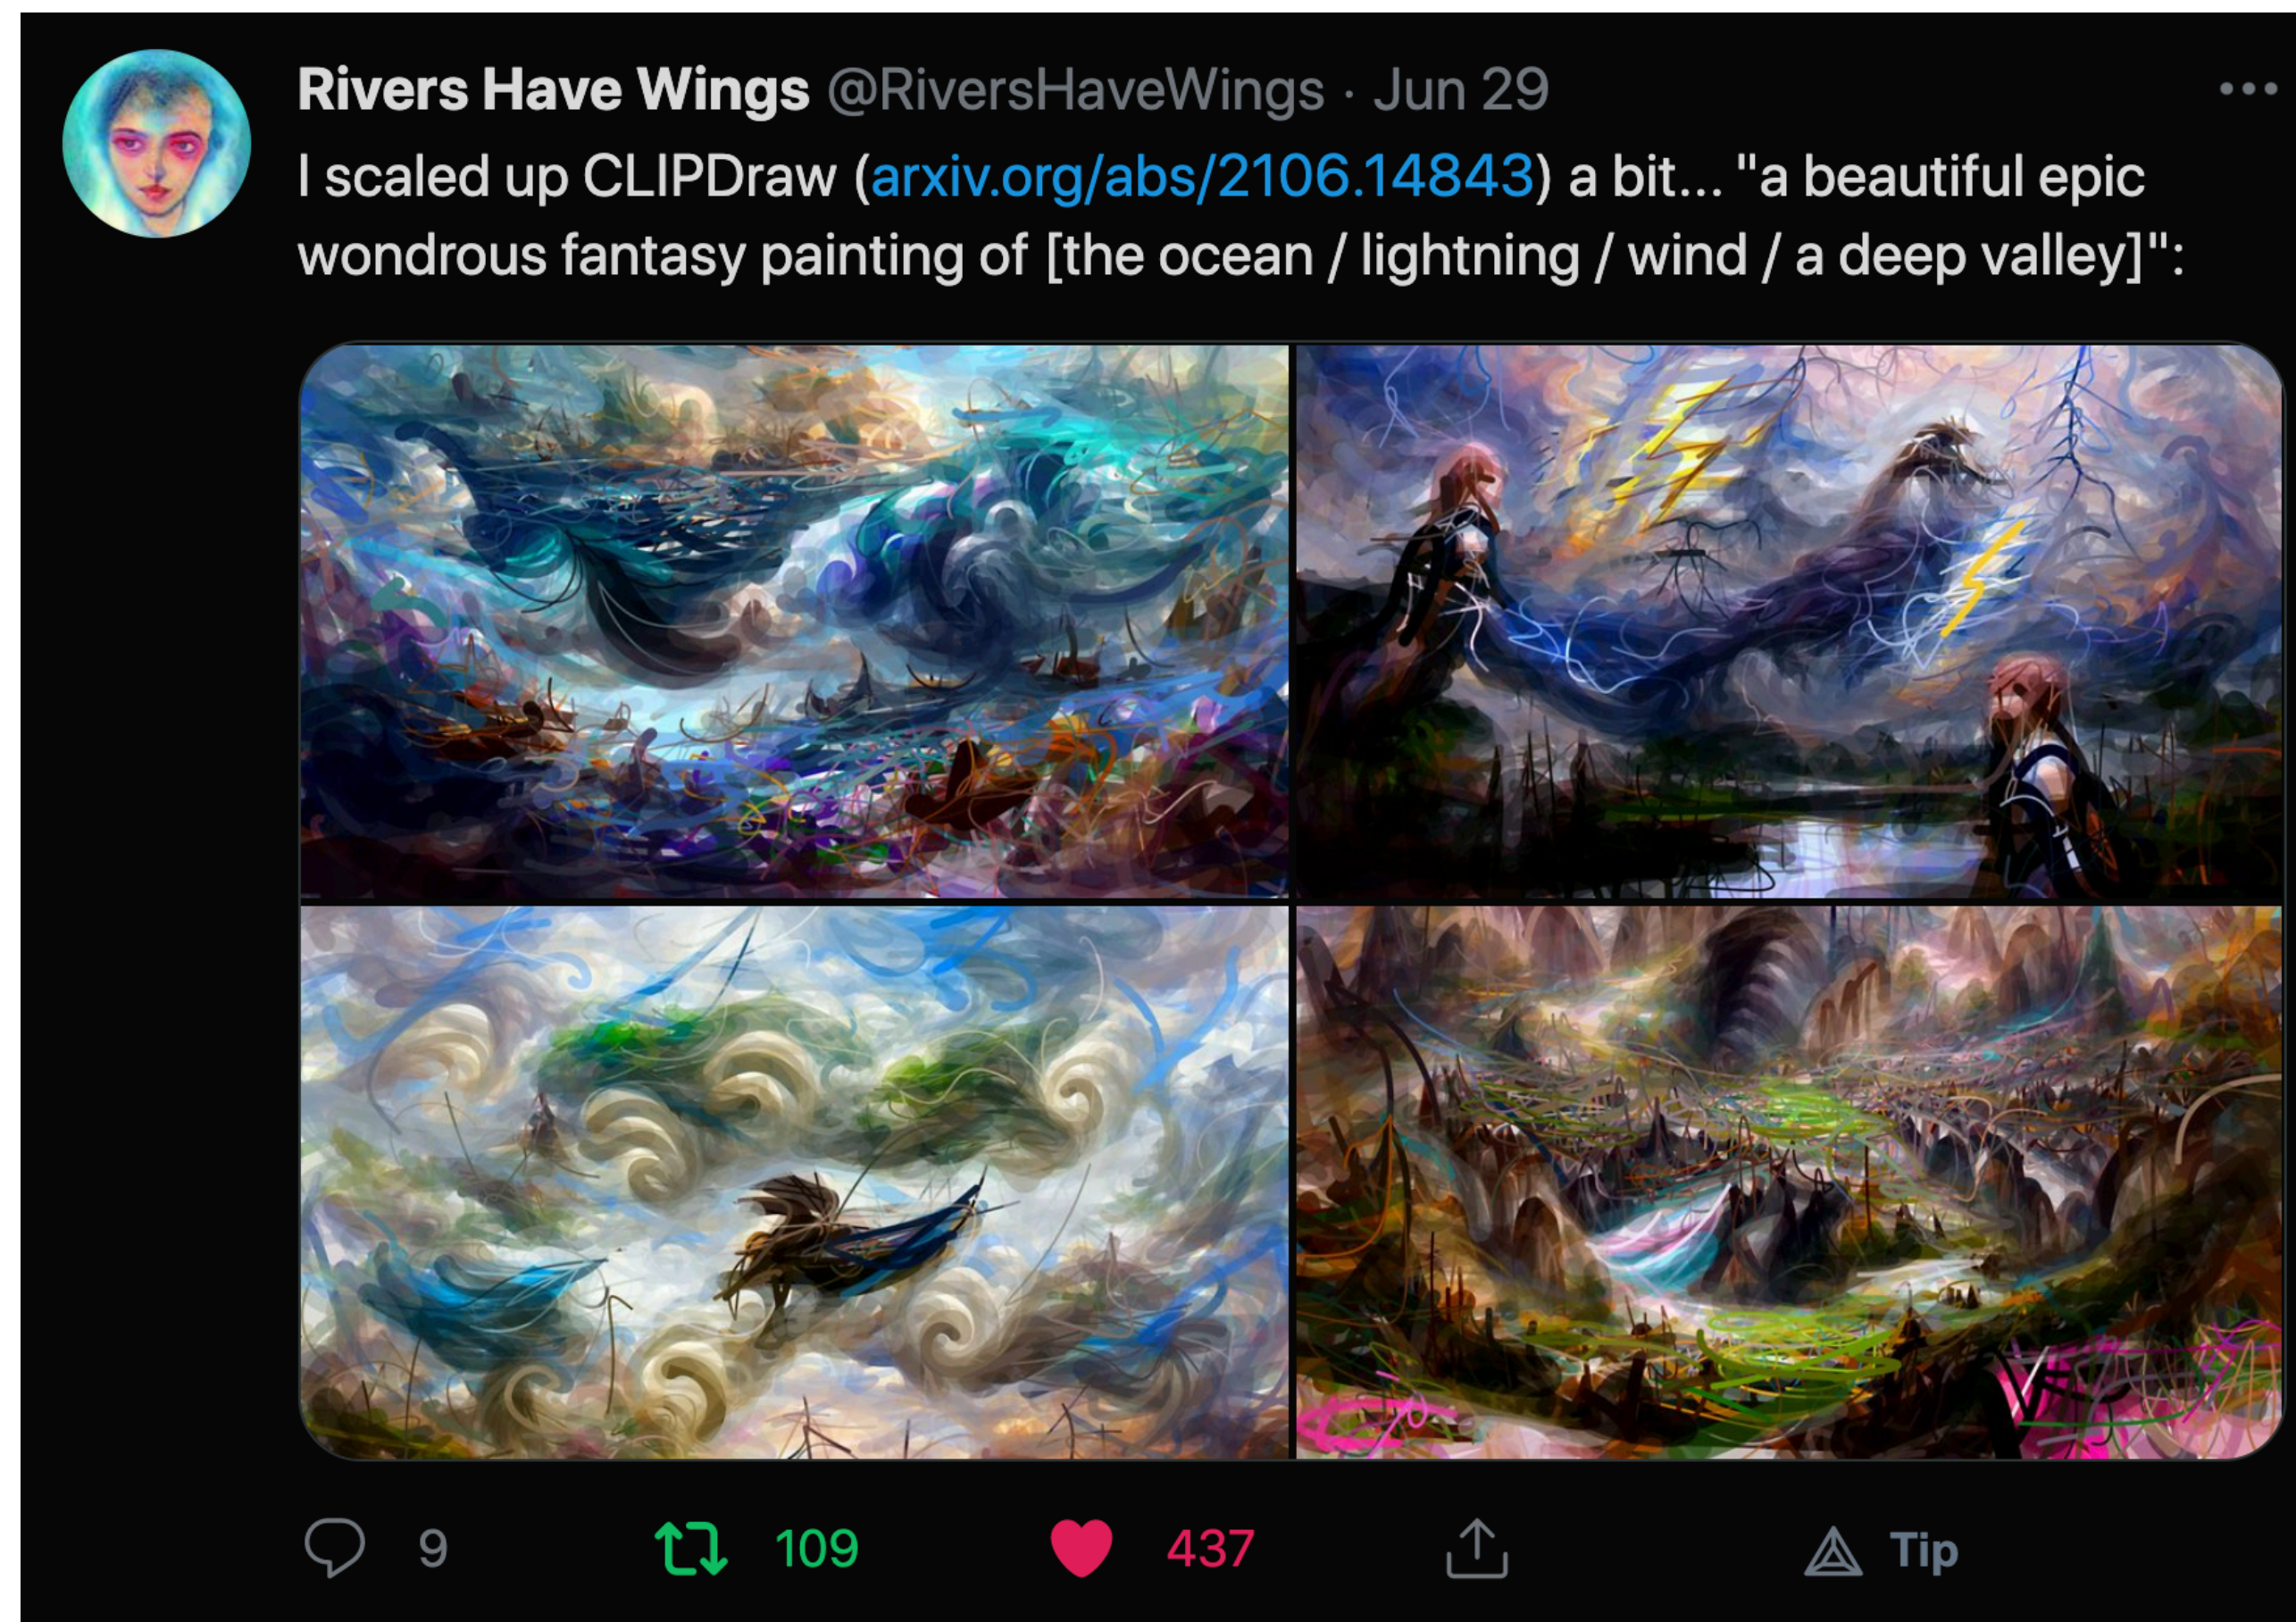
- text to drawing using CLIP & differentiable rasterization using our renderer



CLIPDraw [Frans et al.]

(not my work!)


- text to drawing using CLIP & differentiable rasterization using our renderer




Inverse shader design

- thresholding Perlin noise leads to discontinuities

Perlin noise shader



Thresholded Perlin noise




Threshold
... * [noise > t] * ...

Inverse shader design



target image




our shader optimization

Inverse shader design


- ignoring discontinuities lead to worse / incorrect results



target image




our shader optimization




naive autodiff

Inverse shader design


- ignoring discontinuities lead to worse / incorrect results



target image




our shader optimization



naive autodiff


Recap

differentiable rendering connects 2D and 3D (and vector/raster)



3D



**differentiable
rendering**



2D

Recap

automatically differentiating a renderer computes incorrect results!



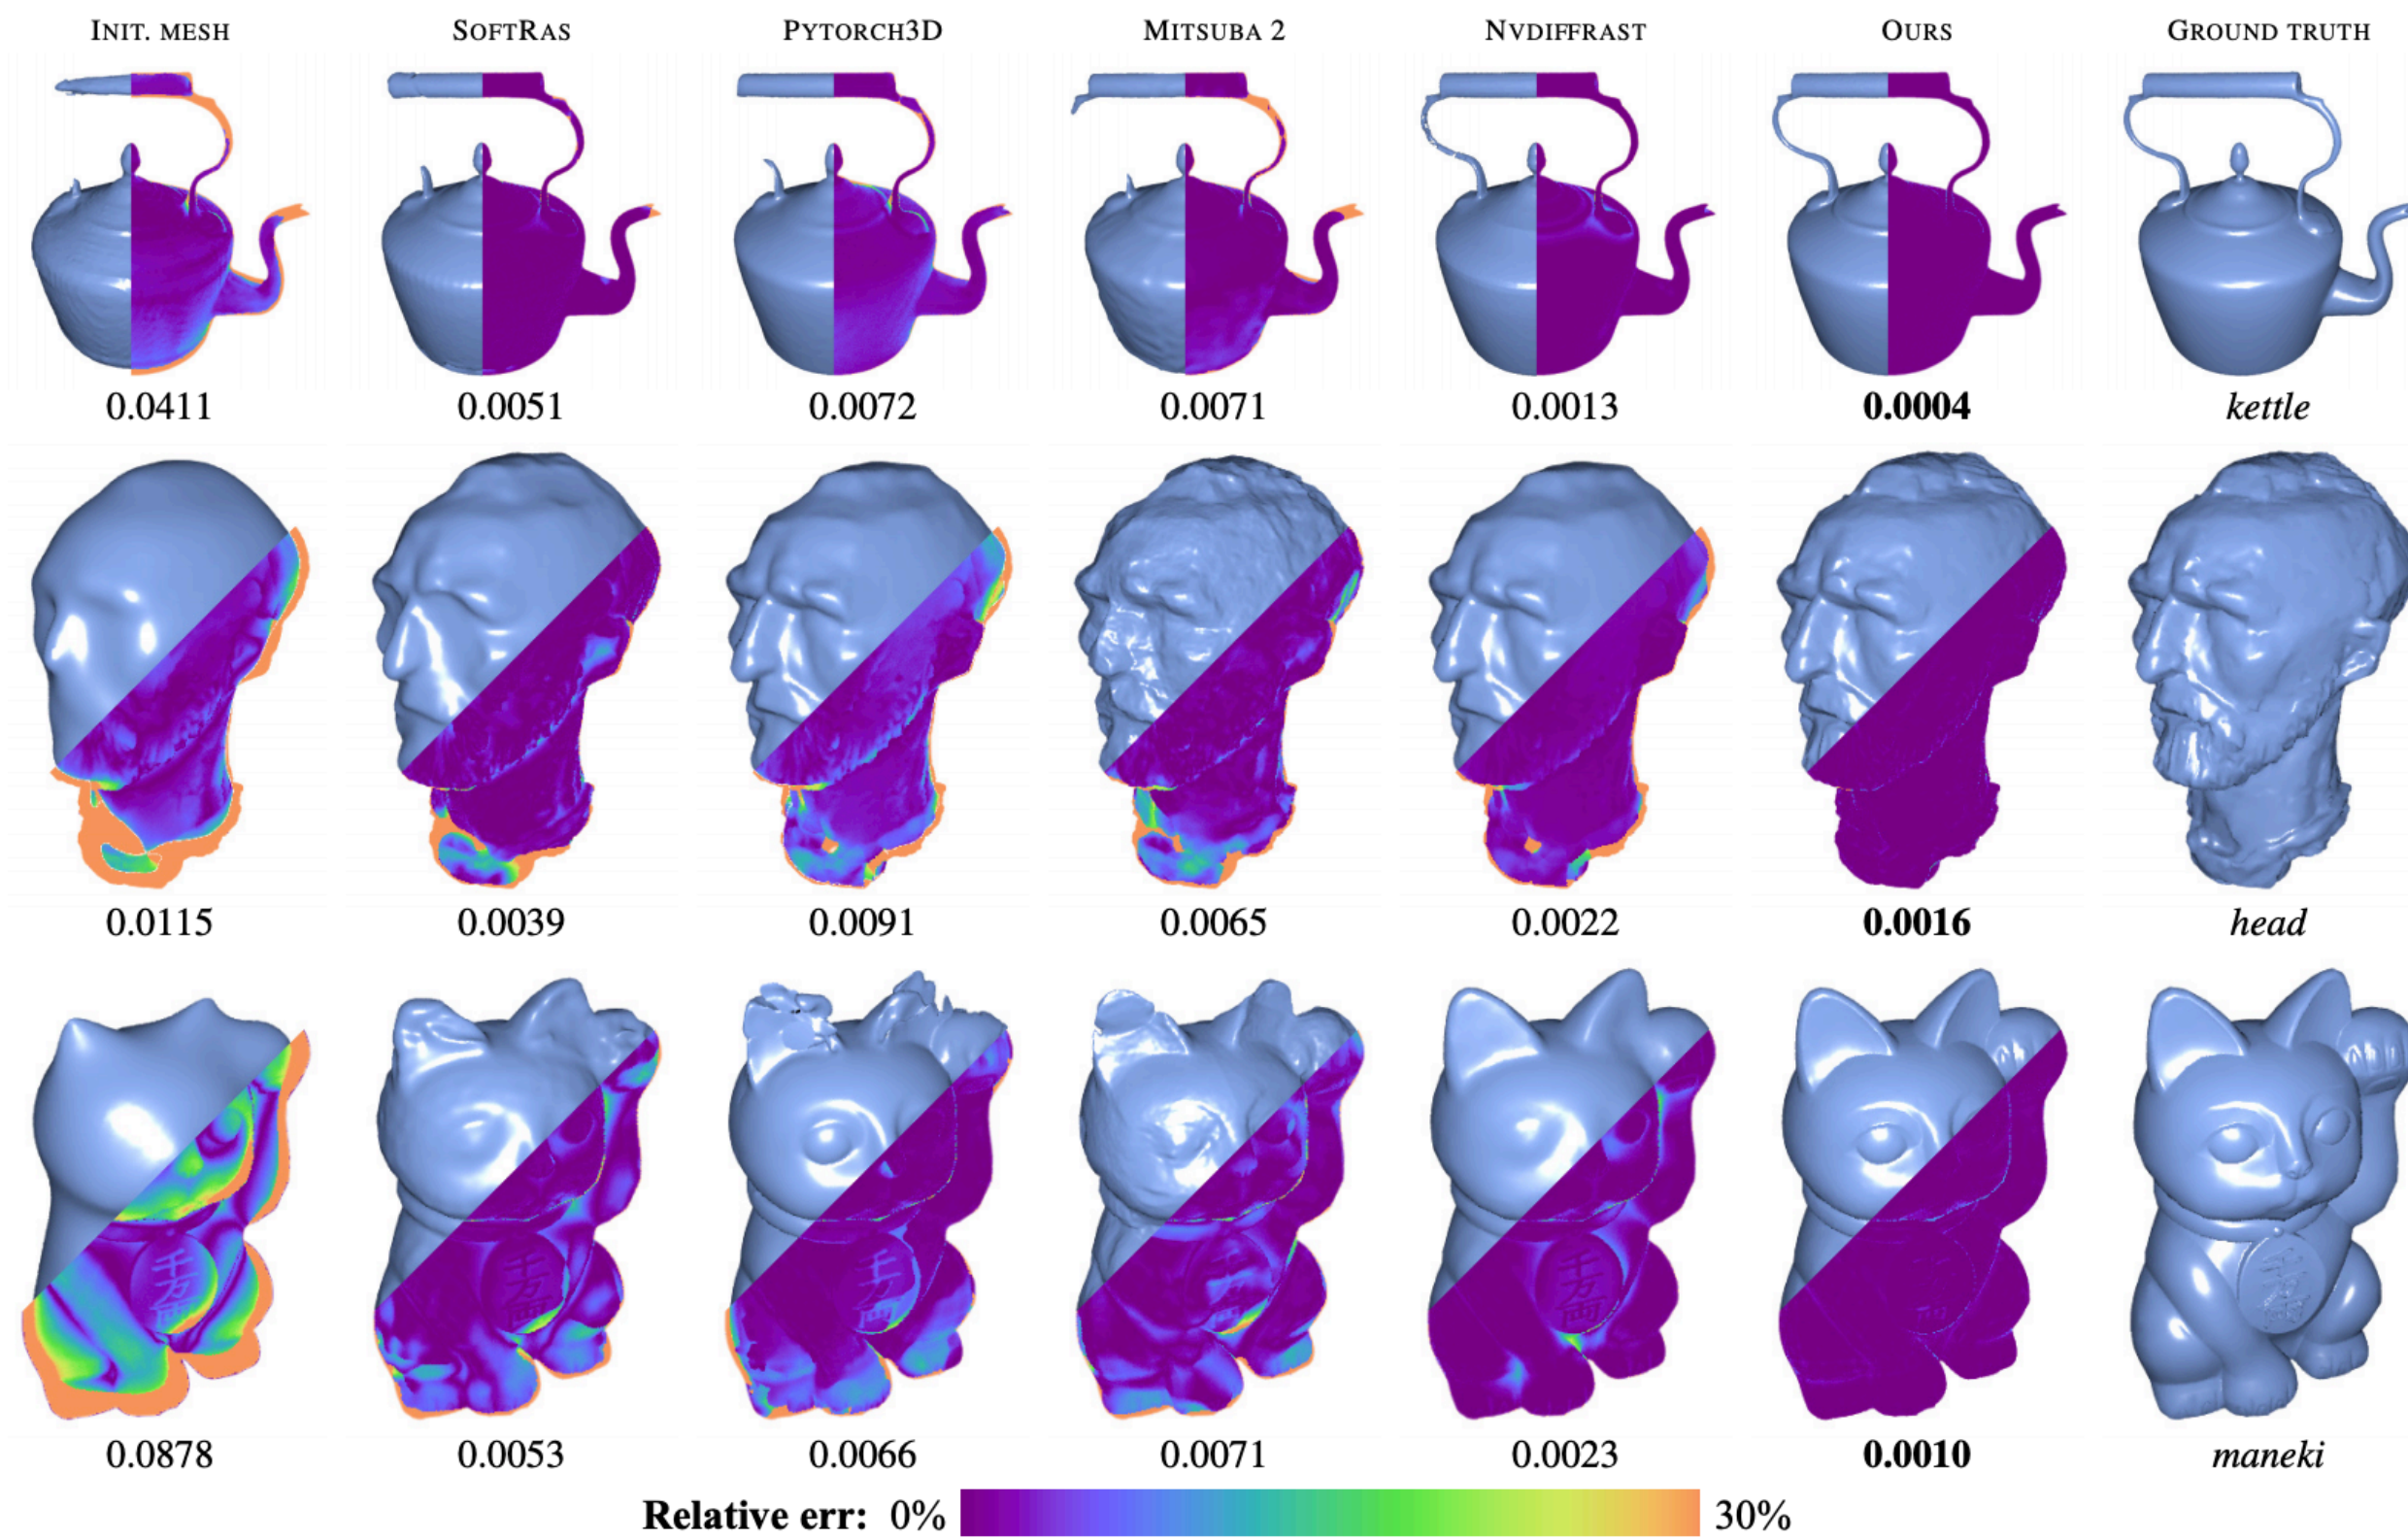
Recap

- what autodiff is missing:
the domain knowledge that rendering equation is an integral
- we derive the correct derivatives from first principles

$$\nabla \iint \text{[Image]} = \iint_{\text{area}} \nabla \text{[Image]} + \int \text{[Image]}$$

area boundary

Correct gradients & lighting matters




Maier 2017

Ours

photo

Beyond Neural Networks

differentiable rendering enables 3D & vector reasoning in optimization and training





inverse rendering



analyzing vision systems



"A drawing of a cat".

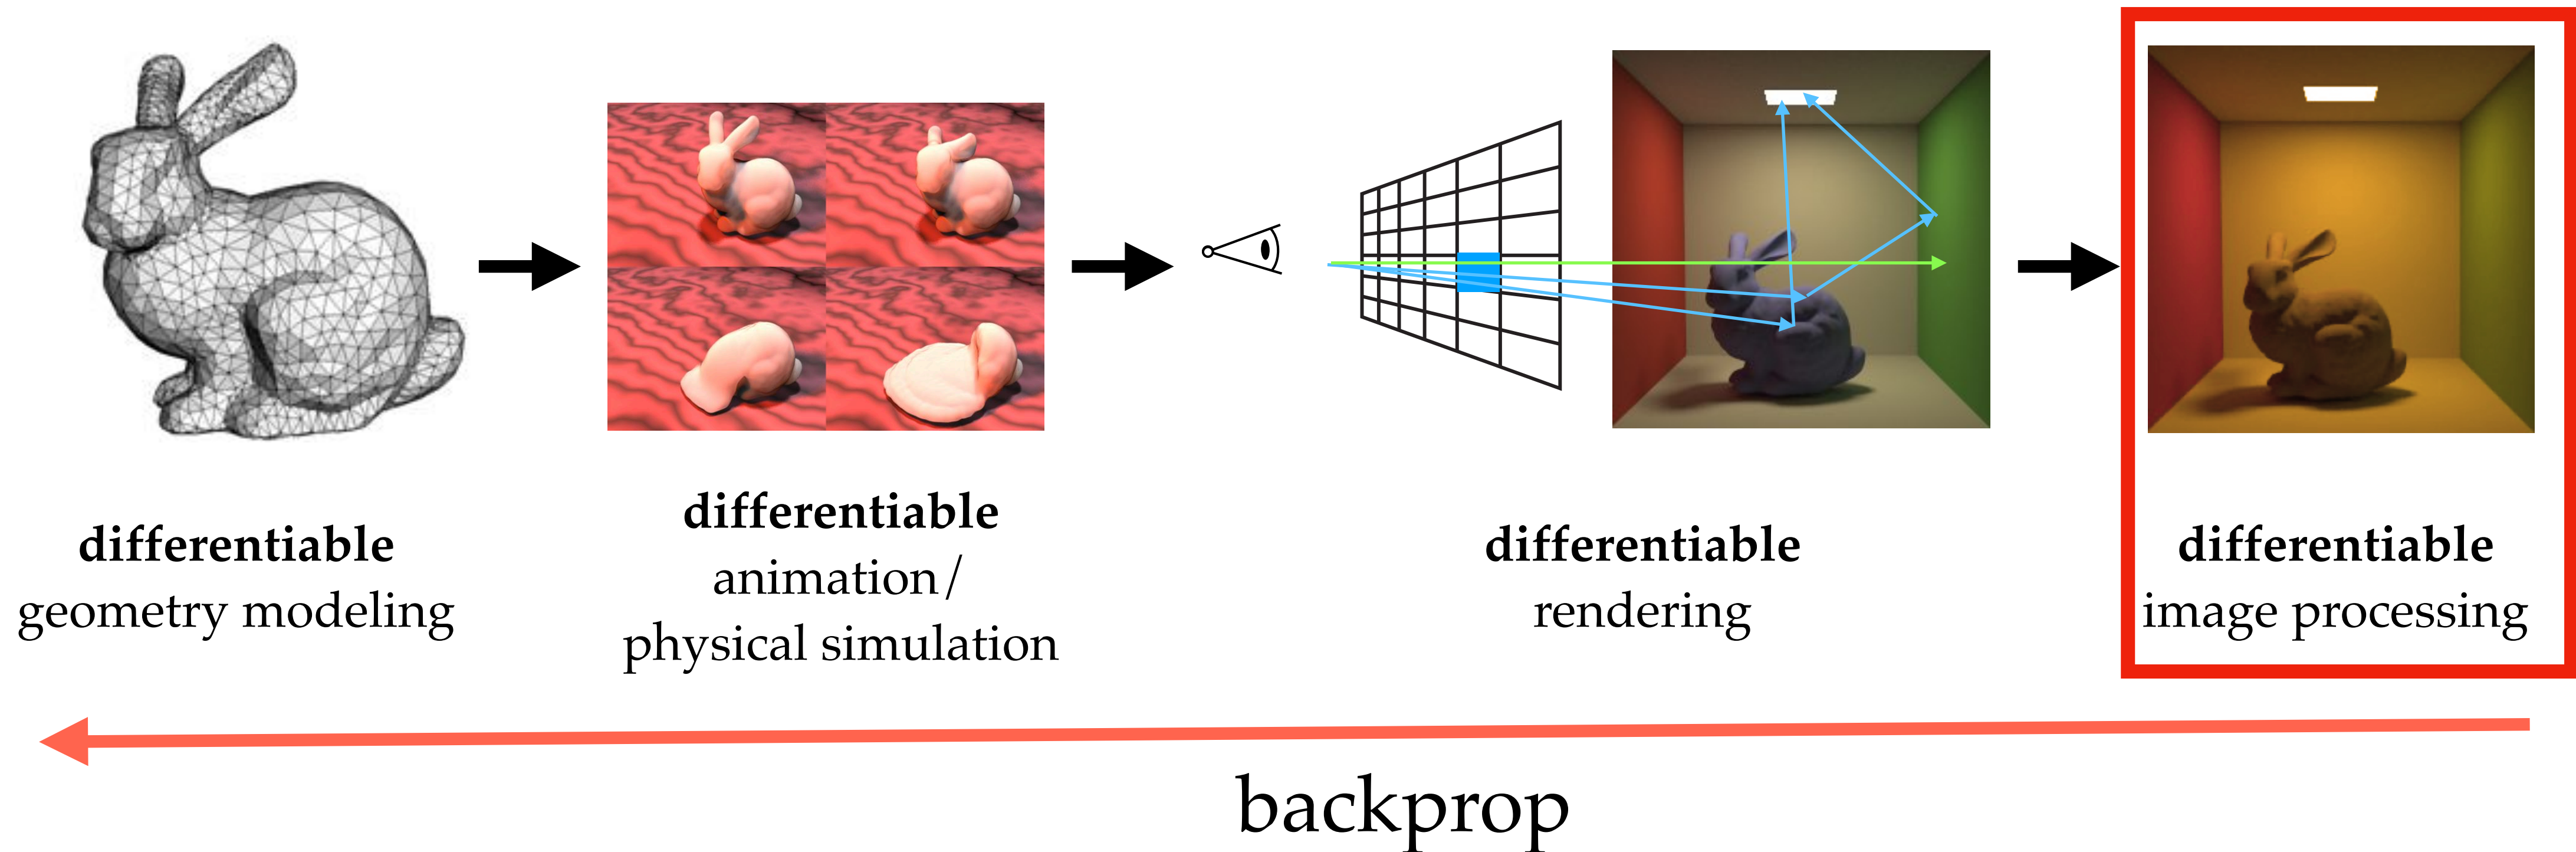


bridging vector and raster graphics

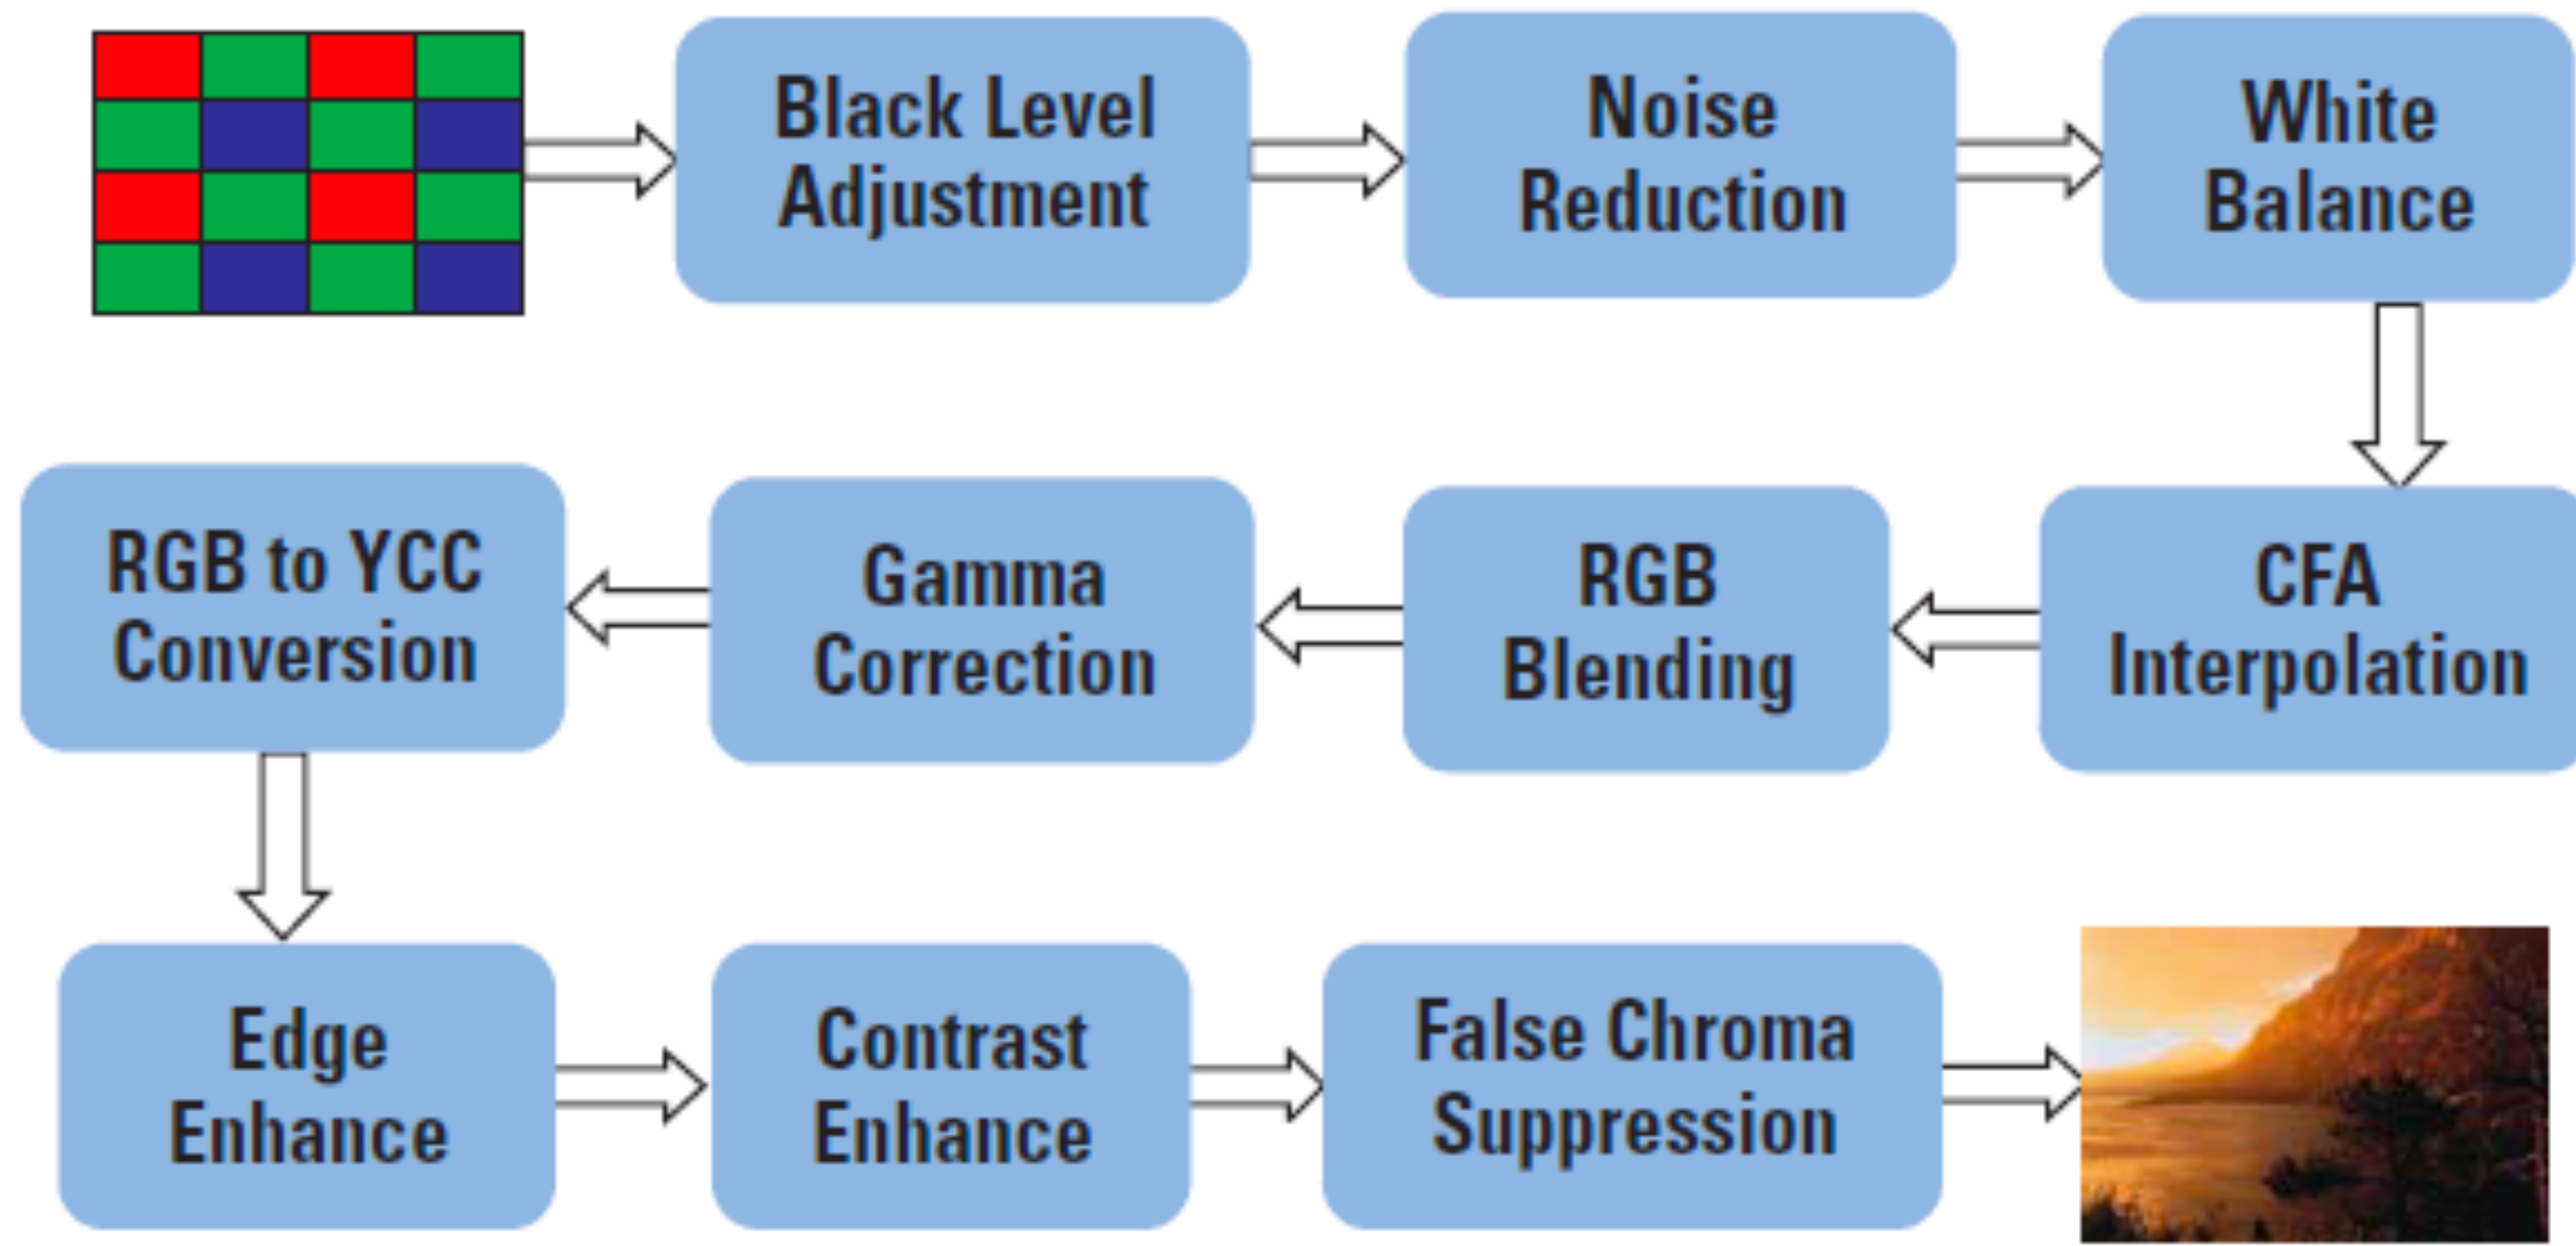
inverse shader design

Differentiable graphics

connects classical graphics algorithms with modern data-driven methods **through derivatives**



Modern cameras have a complex pipeline



Convolutional neural networks: powerful but expensive

designing our own algorithms enables speed & control


deeplab-res101-v2 on full HD: 2.6 TFLOPs, 40GB features



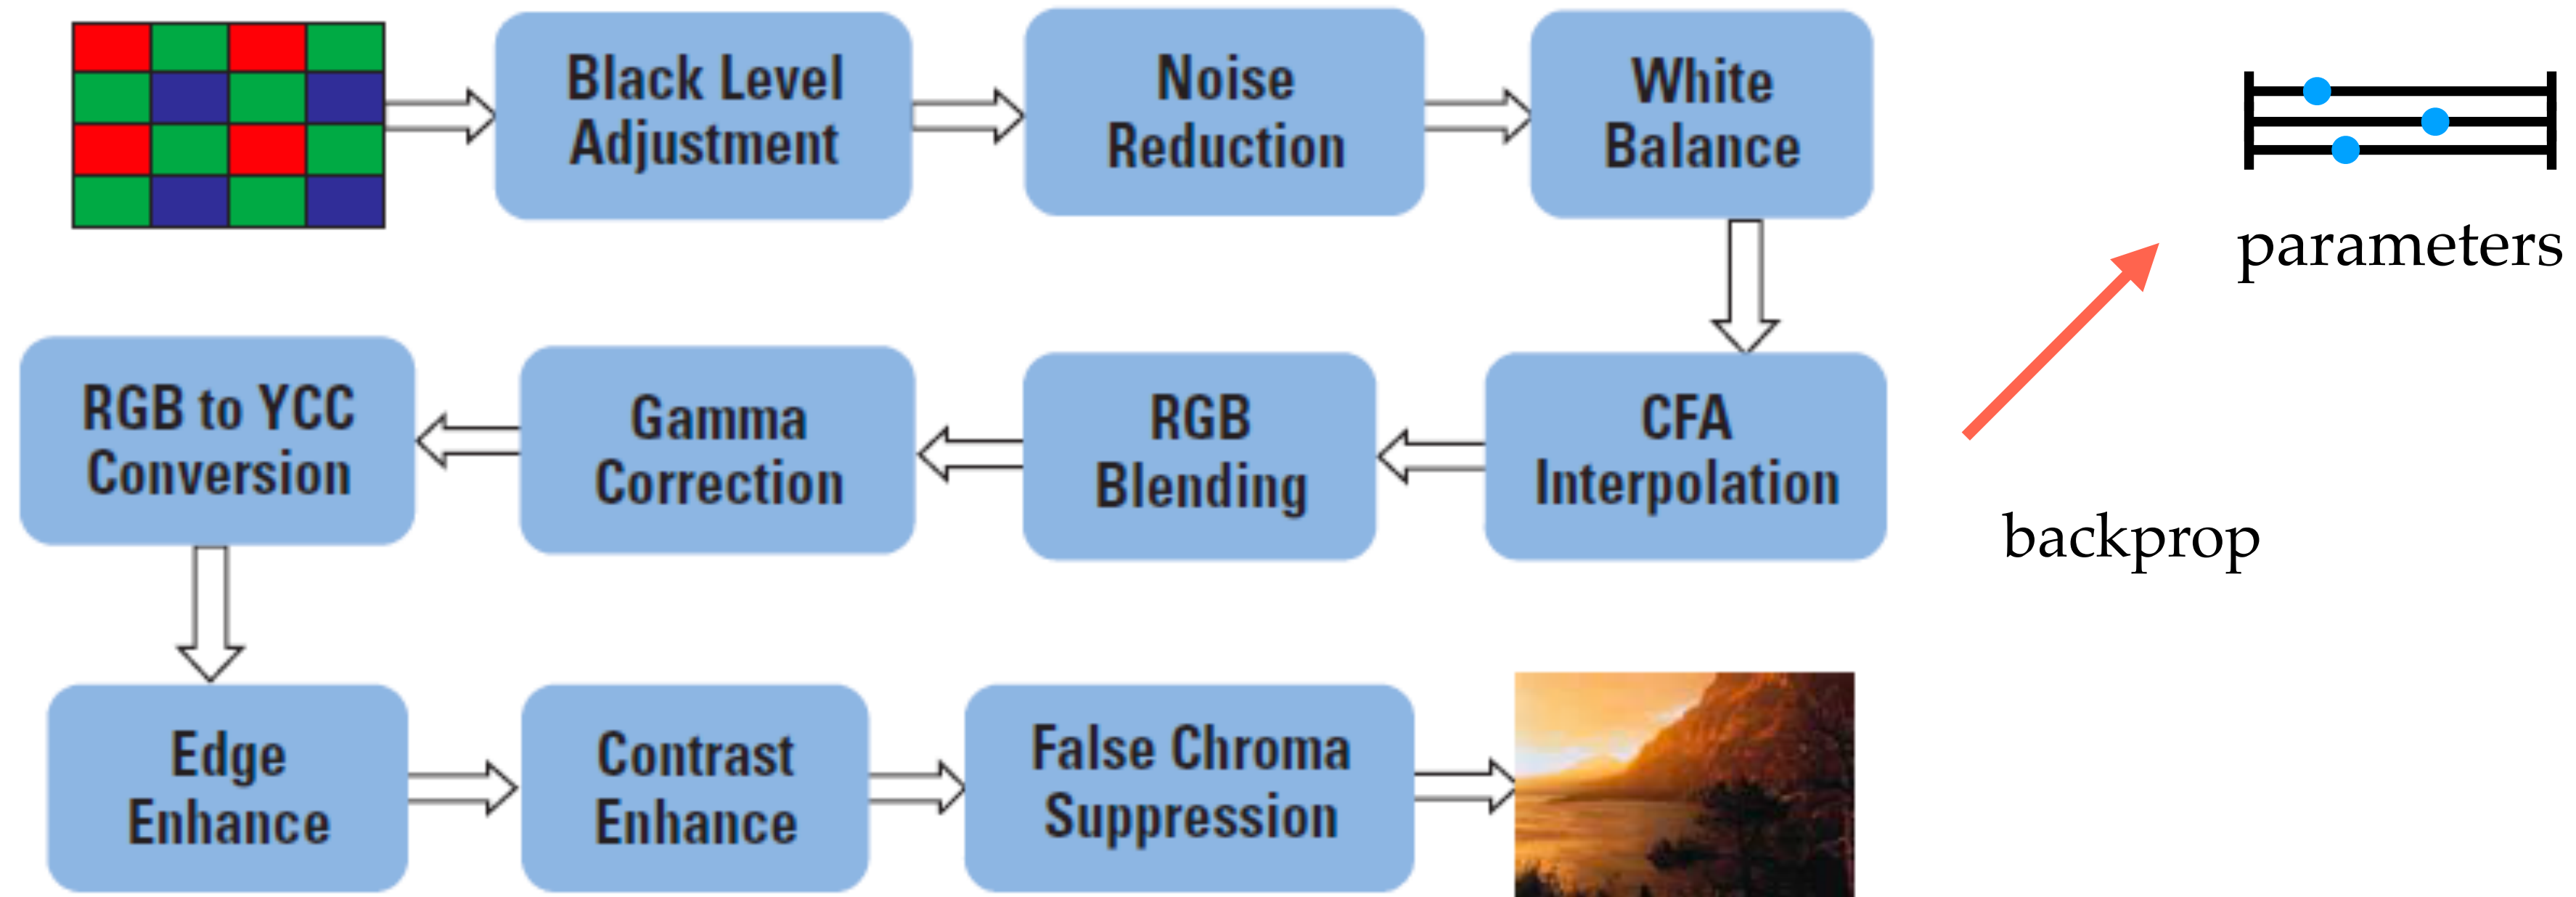
Chen et al.

Google Pixel 4 XL
theoretical peak perf:
954.7 GFLOP/s

The missing Pareto frontier of quality & efficiency trade-offs



Idea: treating traditional image processing building blocks as neural networks



Domain-specific building blocks are more efficient than generic convolutional layers


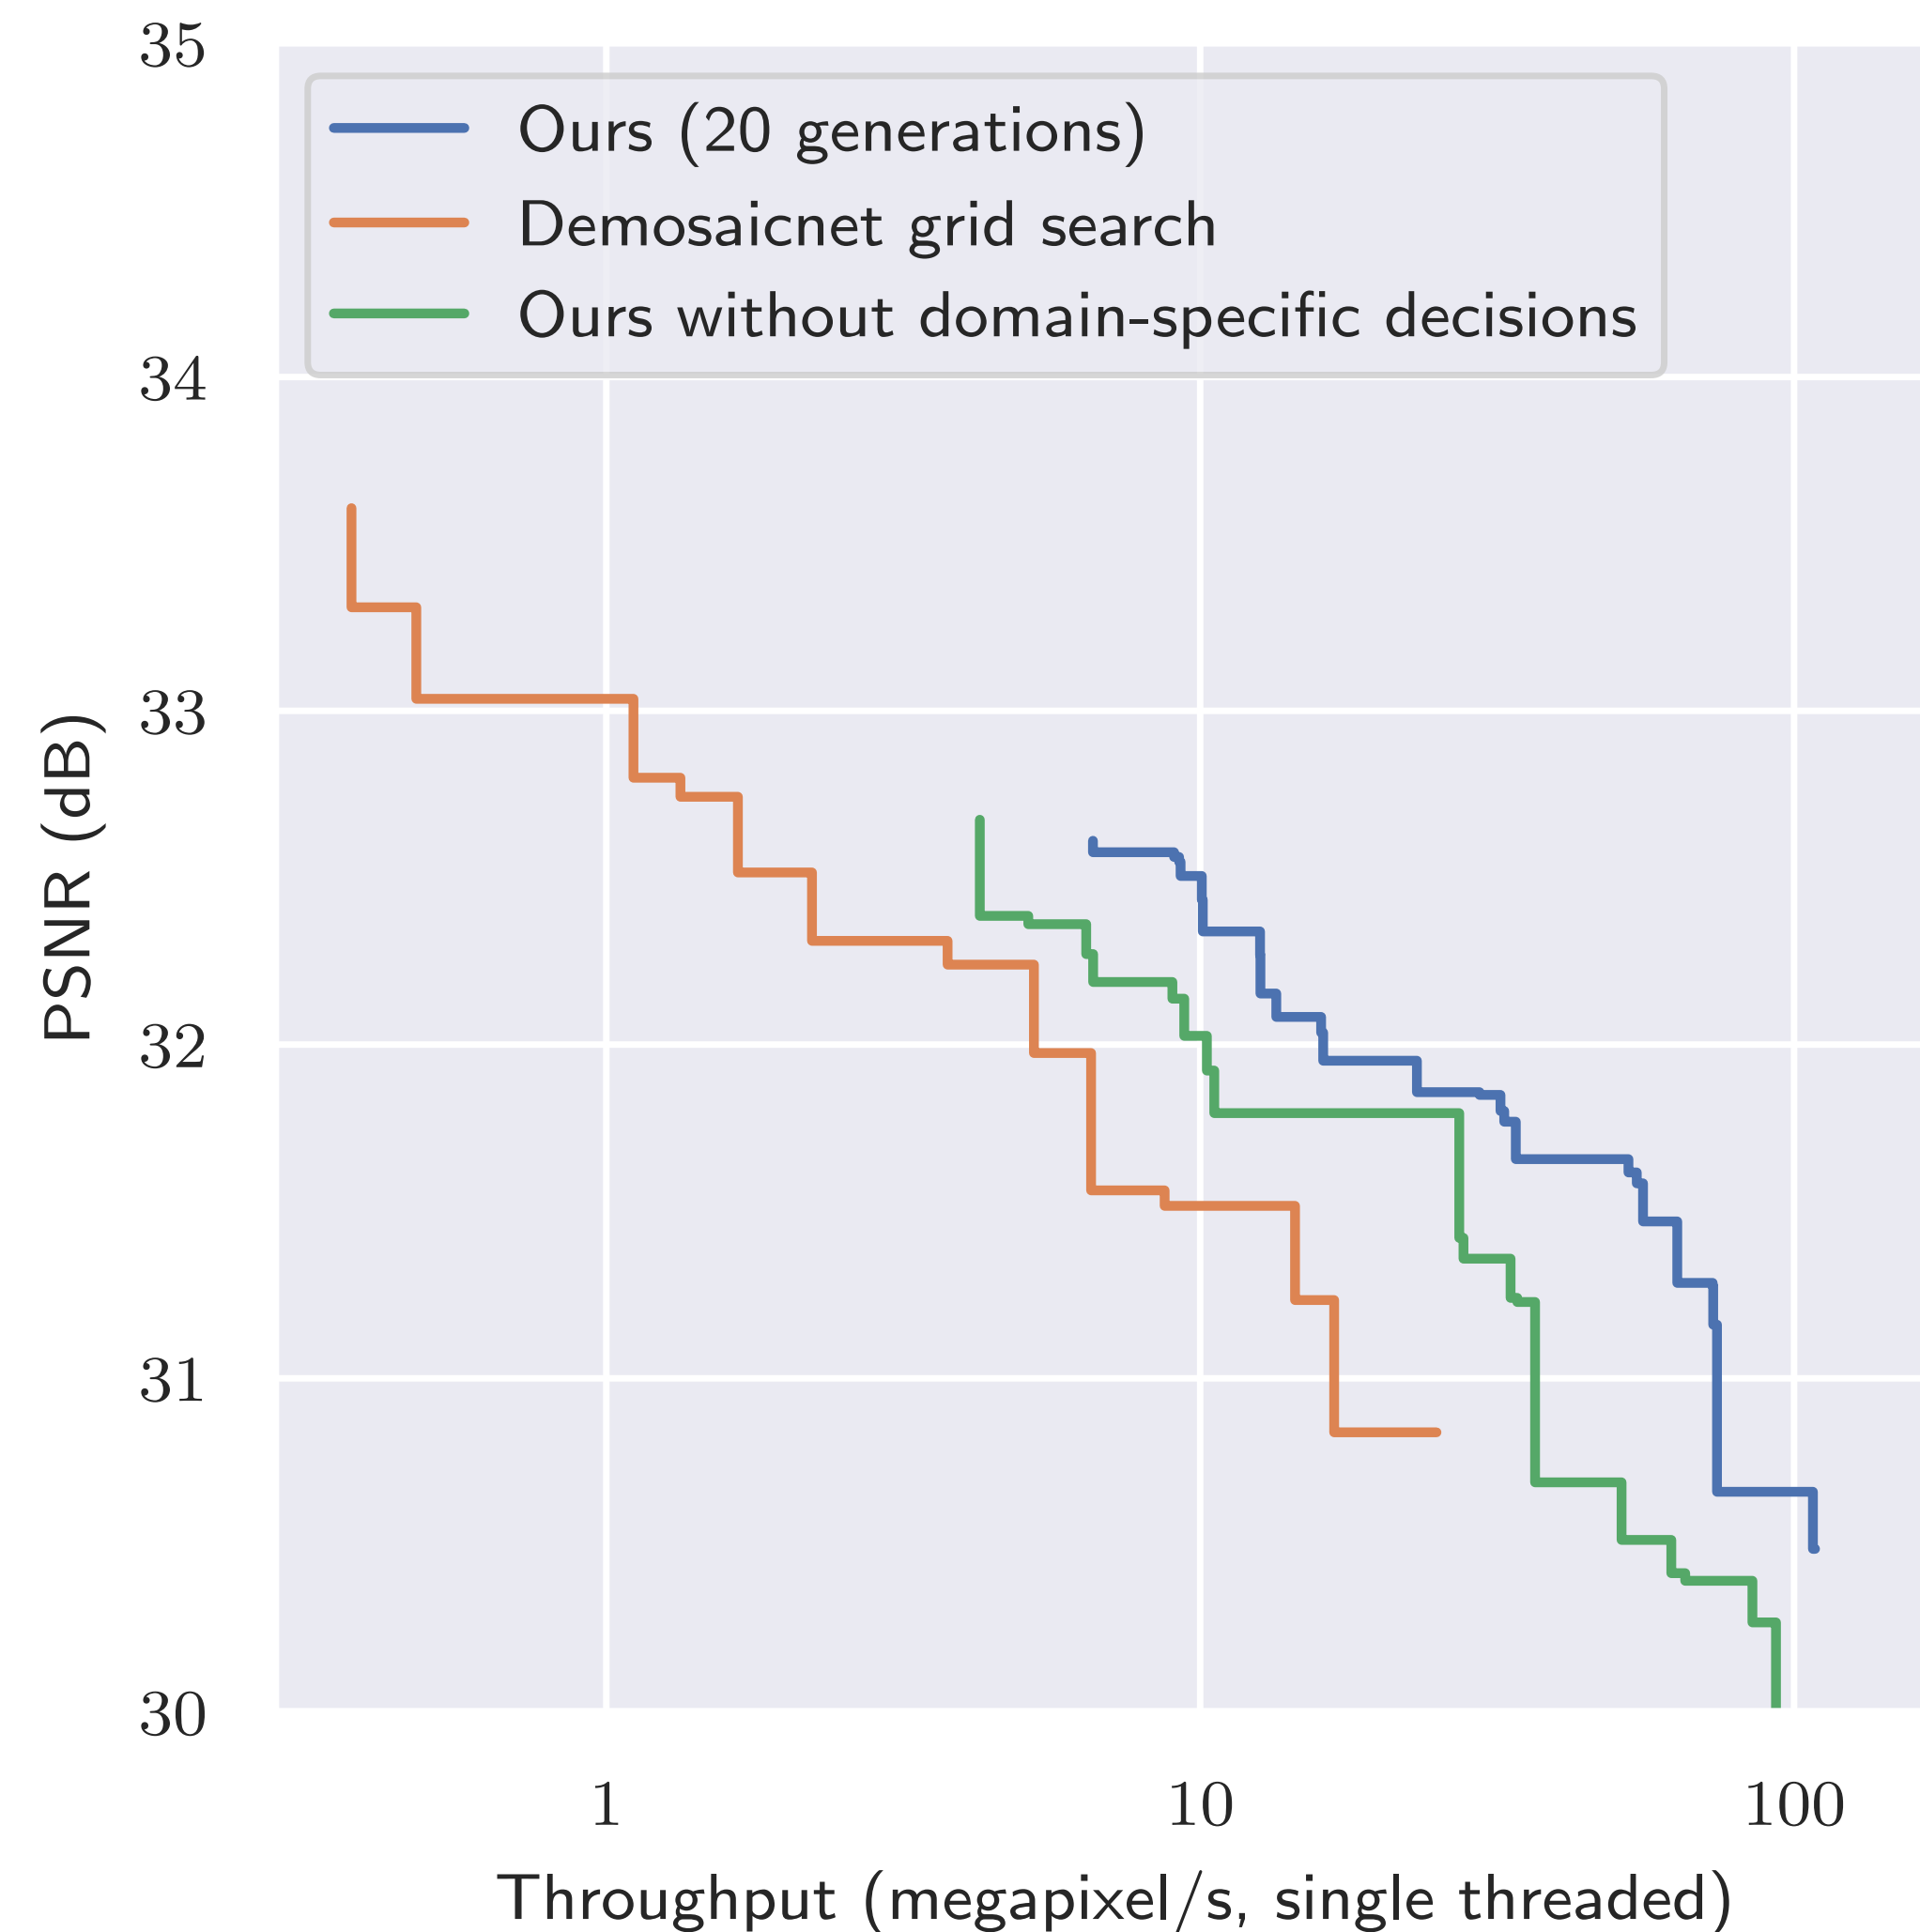



image demosaicing



work in submission

Challenge: deep learning frameworks are limited

new algorithms require new operators

2D convolution
`tf.conv2d`

warping images
`tfa.image.dense_image_warp`

downsample while taking maximum
`tf.max_pool2d`

Challenge: deep learning frameworks are limited

new algorithms require new operators

2D convolution
`tf.conv2d`

warping images
`tfa.image.dense_image_warp`

downsample while taking maximum
`tf.max_pool2d`

4D convolution?

spatially-varying convolution?

Lanzcos interpolation?

different boundary conditions?

progressive mesh?

ray tracing?

skinning?

material point method?

Deep learning frameworks are limited

tensorflow / tensorflow Watch 8,597

Code Issues 1,746 Pull requests 311 Projects 1 Insights

conv4d and higher dimension generalization #1661

Closed jerabaul29 opened this issue on Mar 26, 2016 17 comments

warping images

`tfa.image.dense_image_warp`

downsample while taking maximum

`tf.max_pool2d`

olution?

on?

ditions?

progressive mesh?

ray tracing?

skinning?

material point method?

Deep learning frameworks are limited

tensorflow / tensorflow

Watch 8,597

Code

Issues 1,746

Pull requests 311

Projects 1

Insights

conv4d and higher dimension generalization #1661

Closed

jerabaul29 opened this issue on Mar 26, 2016 · 17 comments

girving commented on Jun 16, 2017

Contributor



I think this is unlikely to happen in a way that makes anyone happy, so closing. All applications of 4D convs that I know of will almost certainly be forced to use special factorizations for efficiency.



girving closed this on Jun 16, 2017

Deep learning frameworks are limited

tensorflow / tensorflow

Watch 8,597

Code

Issues 1,746

Pull requests 311

Projects 1

Insights

conv4d and higher dimension generalization #1661

Closed


jerabaul29 opened this issue on Mar 26, 2016 · 17 comments

girving commented on Jun 16, 2017

Contributor



I think this is unlikely to happen in a way that makes anyone happy, so closing. All applications of 4D convs that I know of will almost certainly be forced to use special factorizations for efficiency.



girving closed this on Jun 16, 2017

Machine Learning Systems are Stuck in a Rut

Paul Barham
Google Brain


Michael Isard
Google Brain

Abstract

In this paper we argue that systems for numerical computing are stuck in a local basin of performance and programmability. Systems researchers are doing an excellent job improving the performance of 5-year-old benchmarks, but gradually making it harder to explore innovative machine learning research ideas.


Generate correct and efficient parallel derivative code is hard

- backpropagation *inverts* dependencies



parallel
read

→
backprop



race
condition

Generate correct and efficient parallel derivative code is hard


- backpropagation *inverts* dependencies
- optimized code is harder to analyze

$$y \leftarrow \sum x_i \qquad dx_i \leftarrow dy$$


Generate correct and efficient parallel derivative code is hard

- backpropagation *inverts* dependencies
- optimized code is harder to analyze

$$y \leftarrow \sum x_i$$




$$dx_i \leftarrow dy$$



Key idea: differentiate algorithm, not implementation

$$y = \sum x_i$$


high-level, domain-specific
algorithm



low-level
implementation

Key idea: differentiate algorithm, not implementation

- avoid race condition by transforming loops

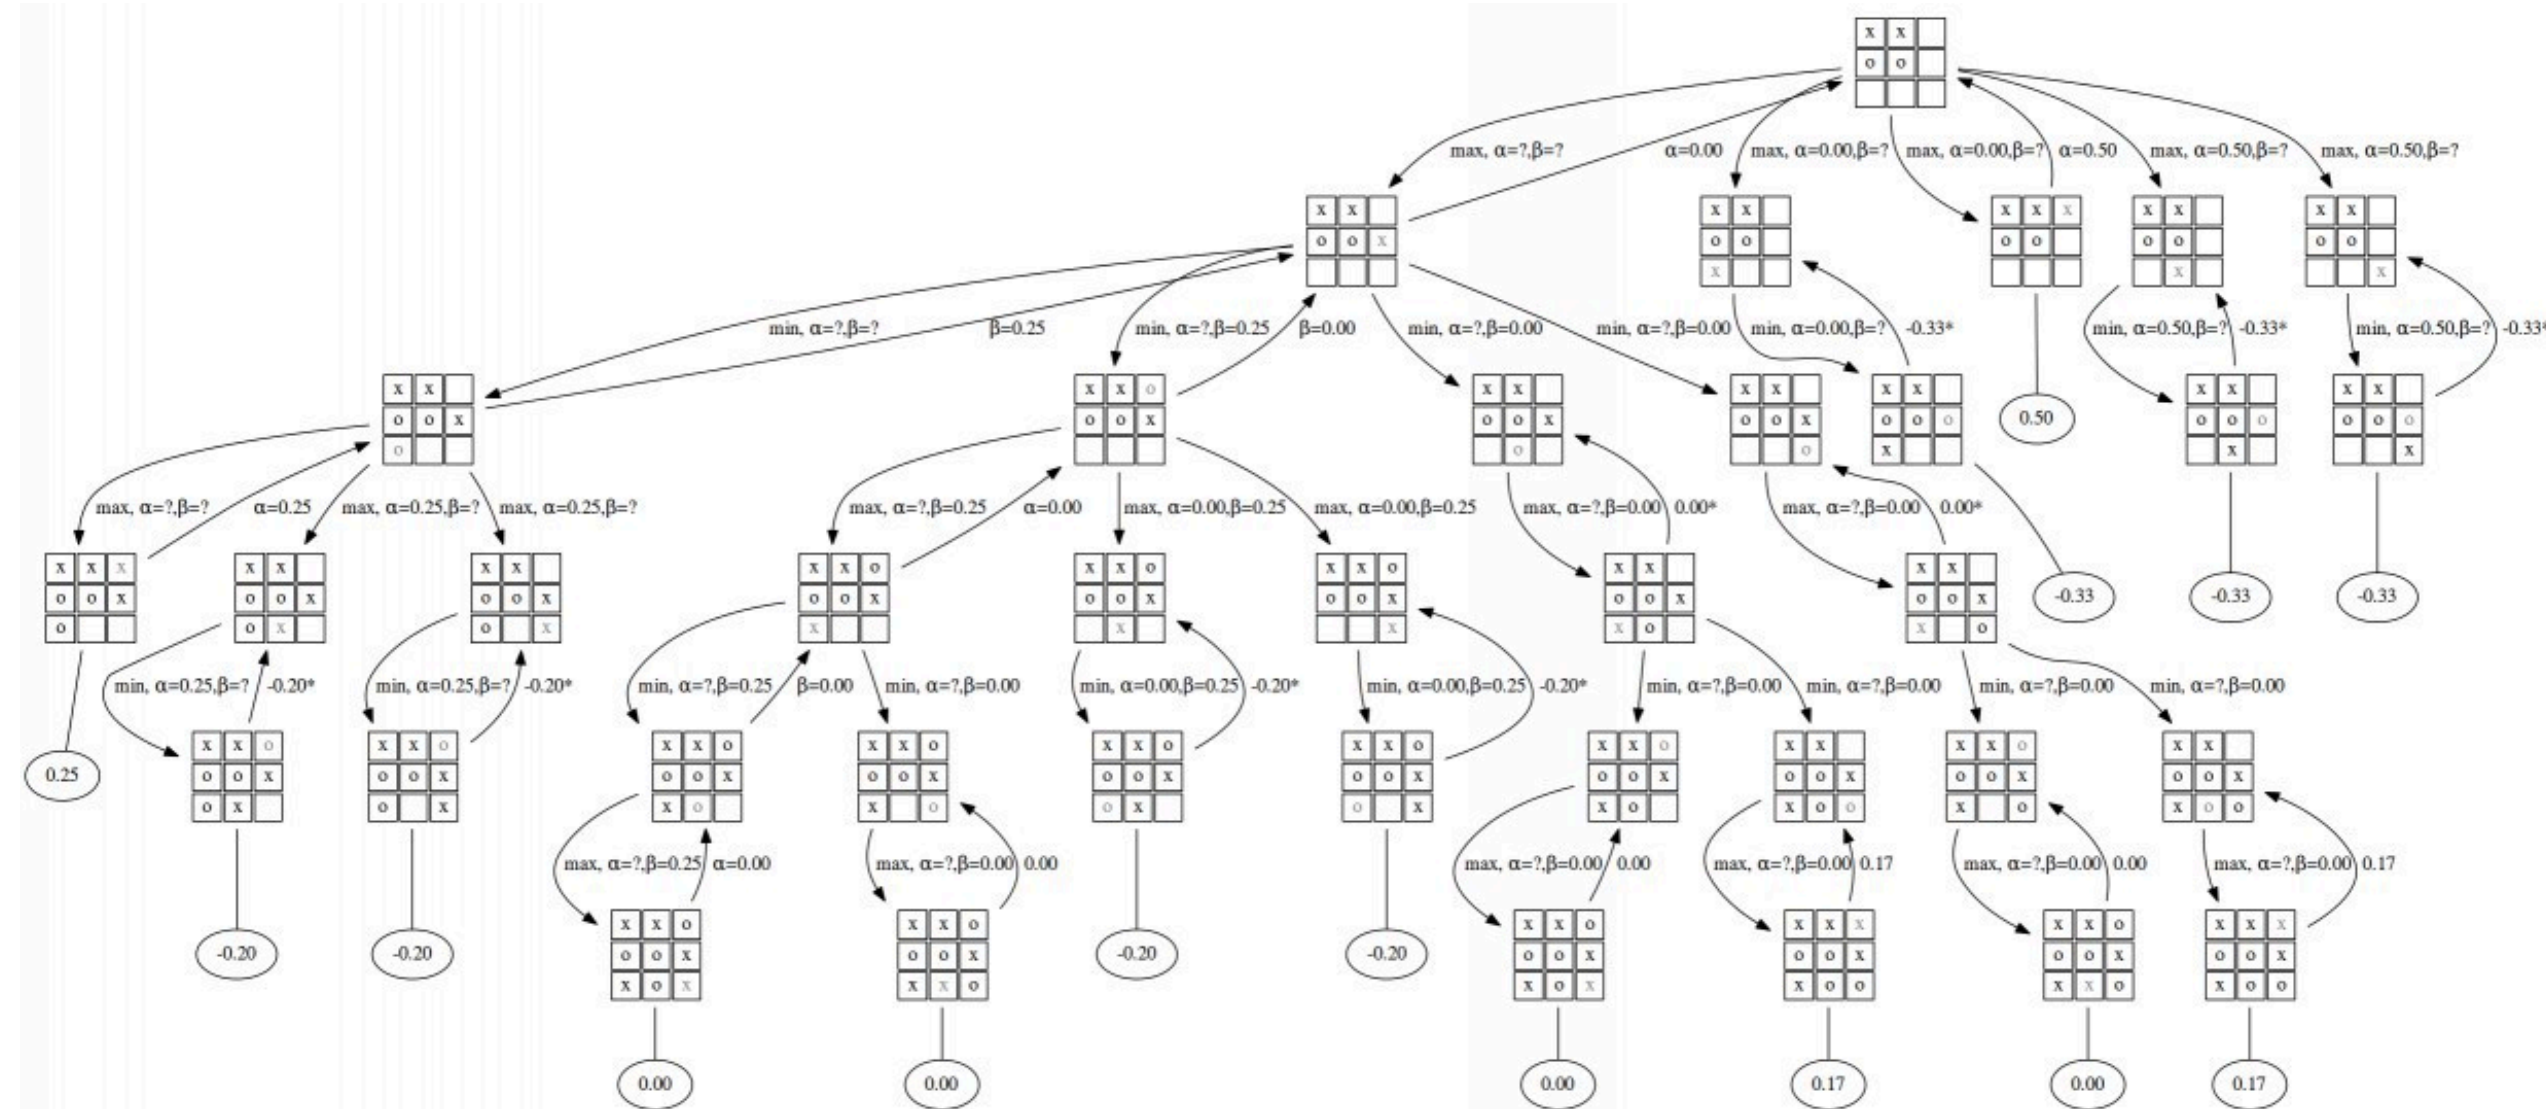


Key idea: differentiate algorithm, not implementation

- avoid race condition by transforming loops
- implementation can be defined by users or searched automatically (c.f. autoscheduling / program synthesis)

$$y = \sum x_i$$

algorithm



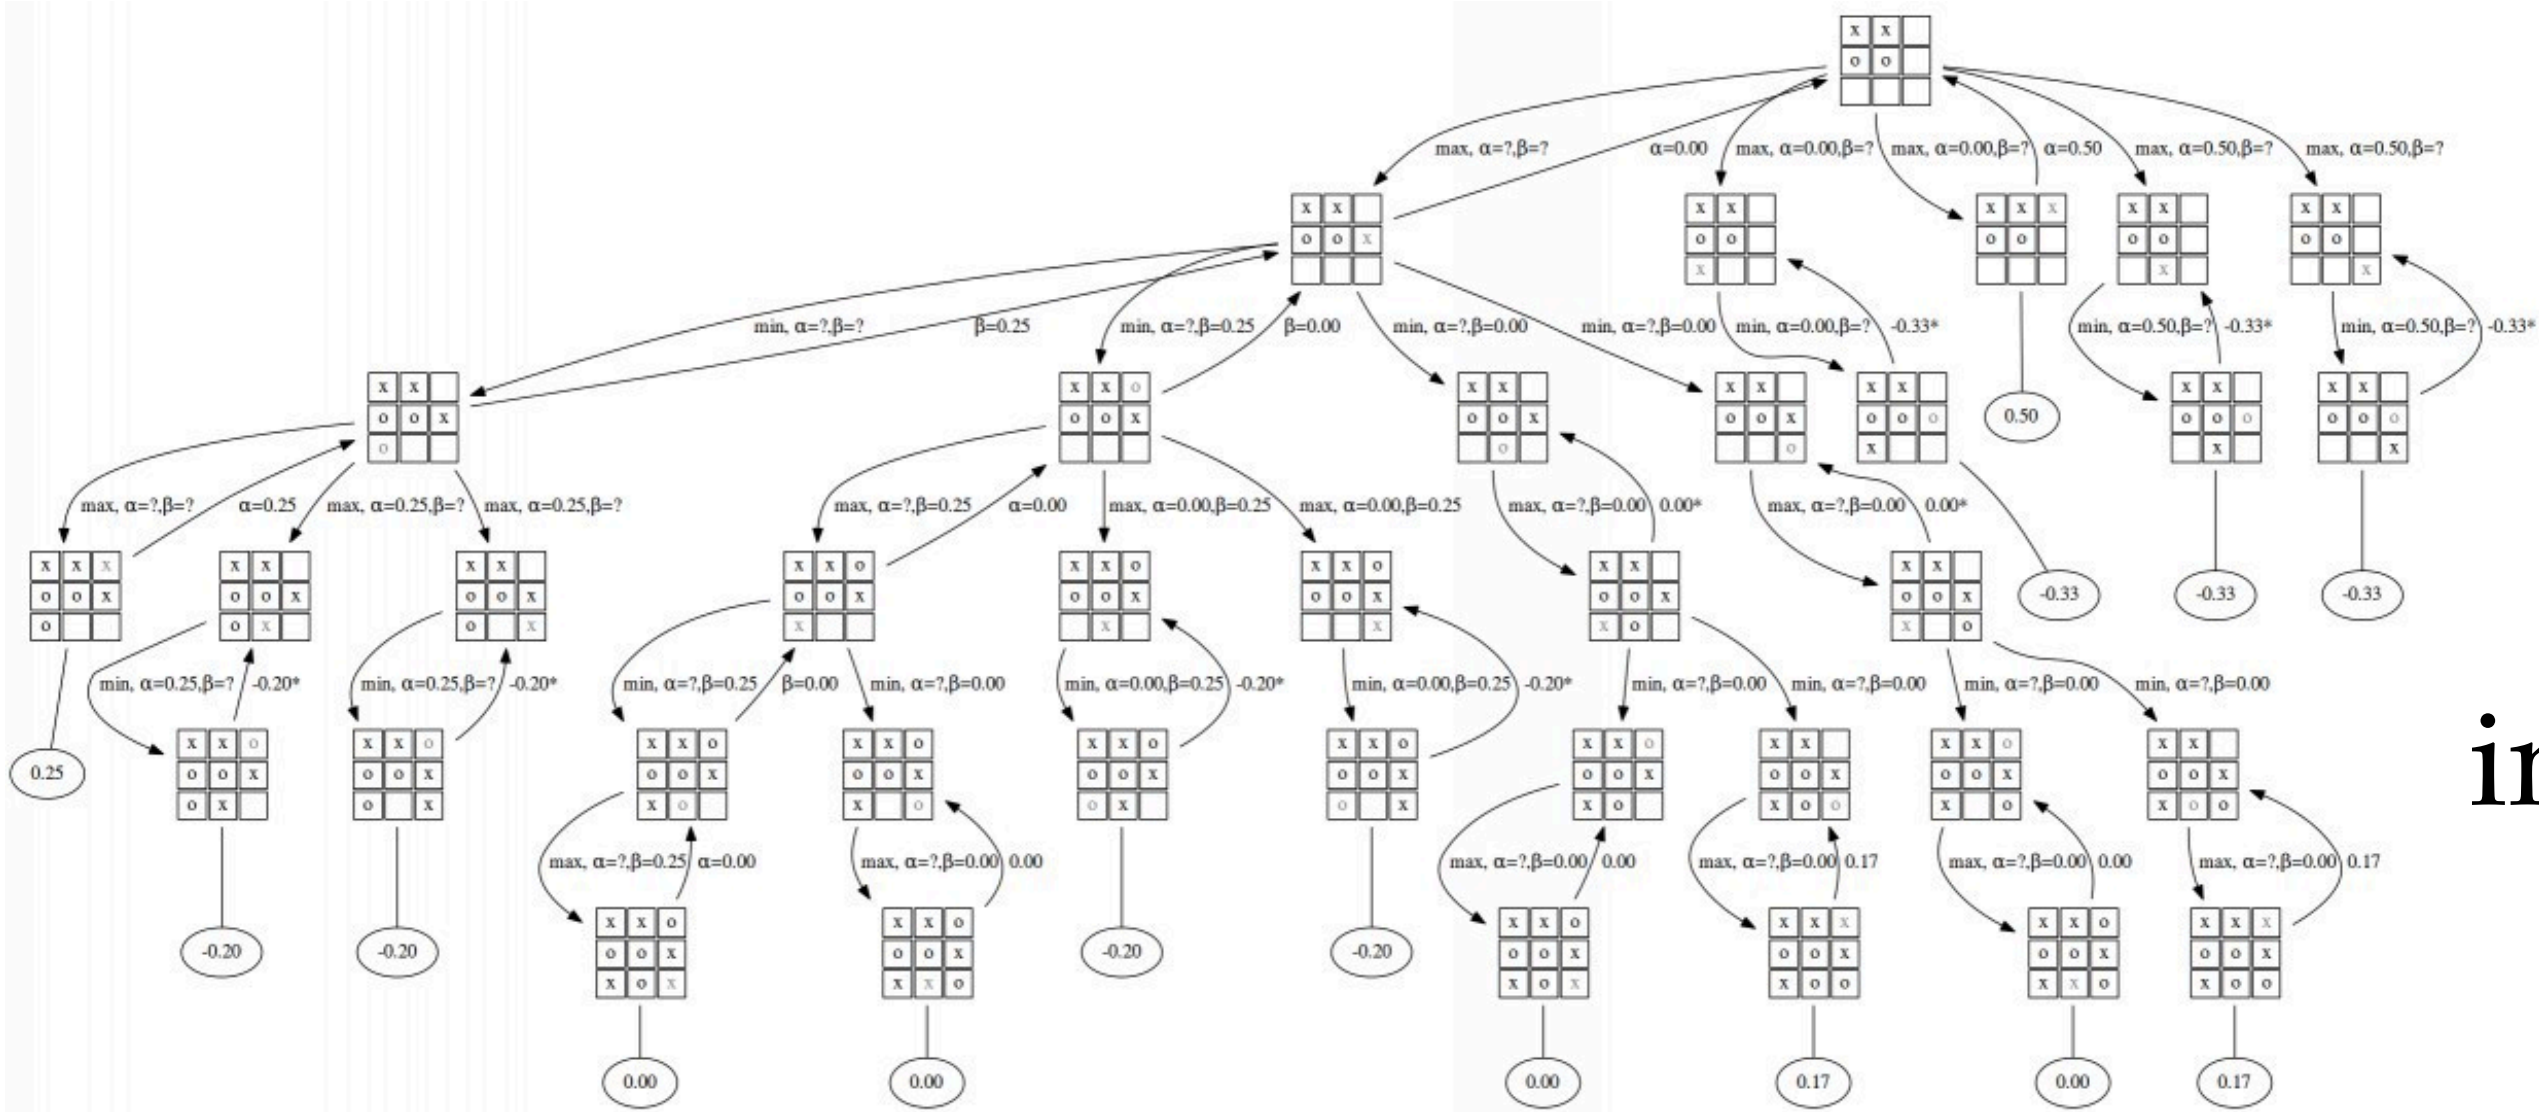
search for
implementation

Automatically searching for implementations

key idea: use a lightweight neural net to predict the performance

$$y = \sum x_i$$

algorithm



implementation

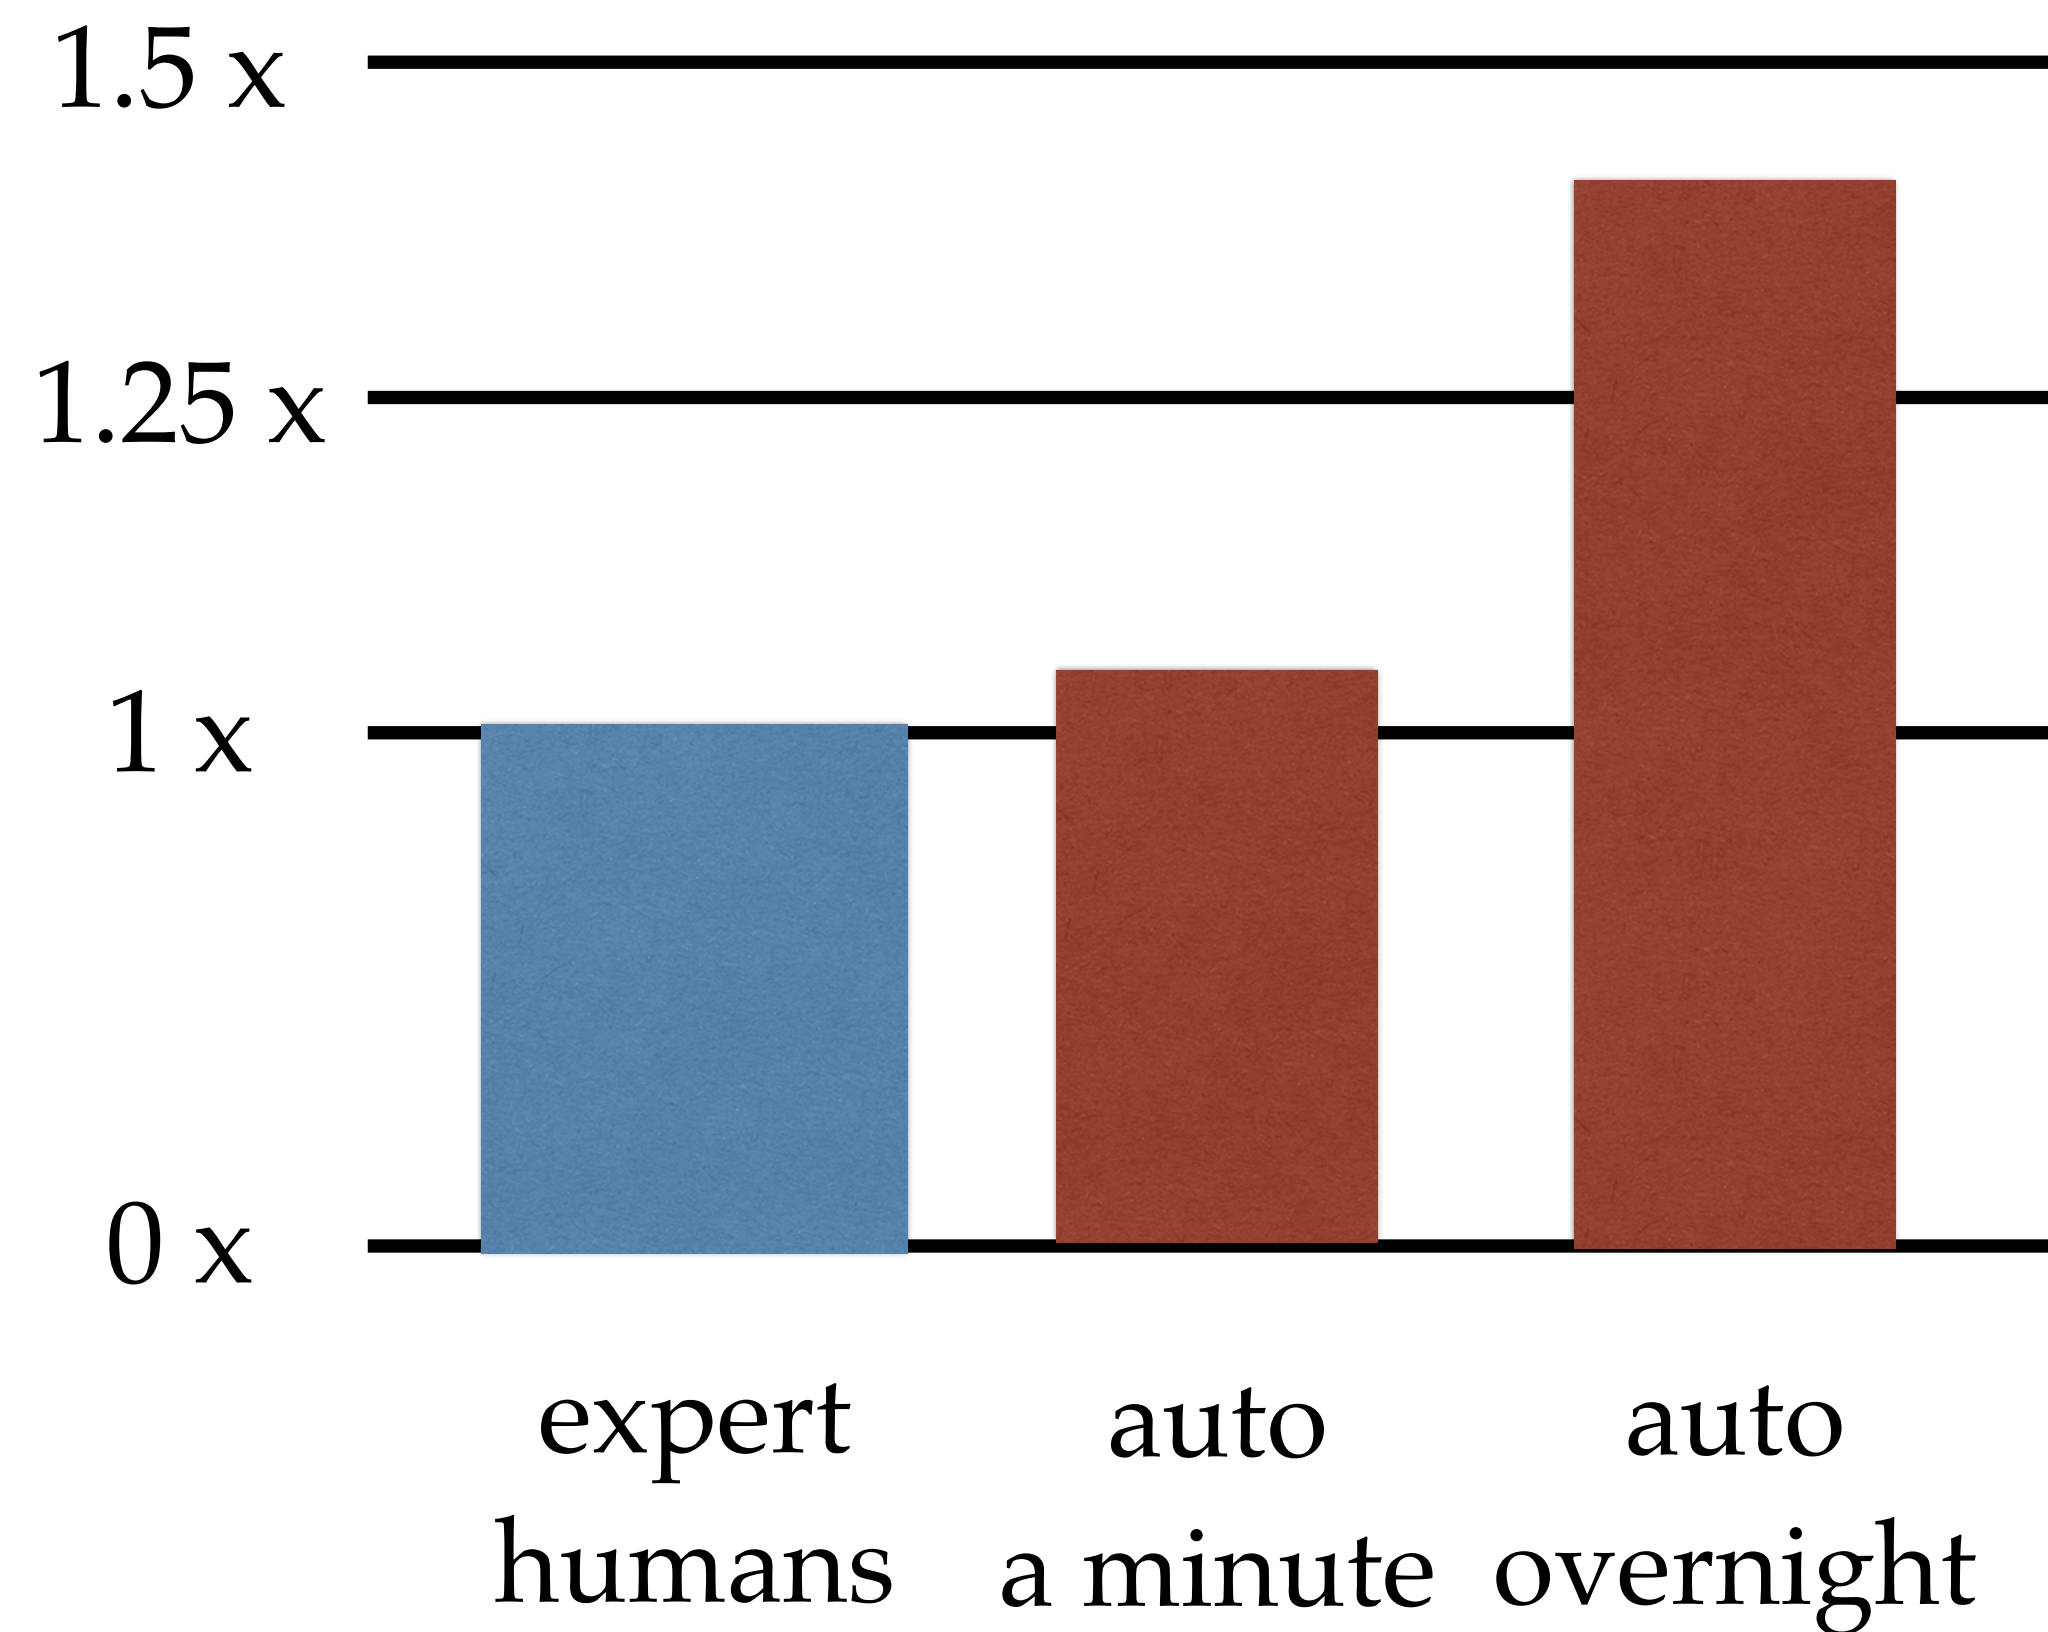
Adams, Ma, Anderson, Baghdadi, Li, Gharbi, ..., 2019

Anderson, Adams, Ma, Li, Ragan-Kelley, 2021

Benchmark: 16 imaging / ML pipelines

on a CPU

speed-up (higher = better)




e.g. local laplacian filter, resnet 50
matmul, non-local means



CAT

Also works for GPUs



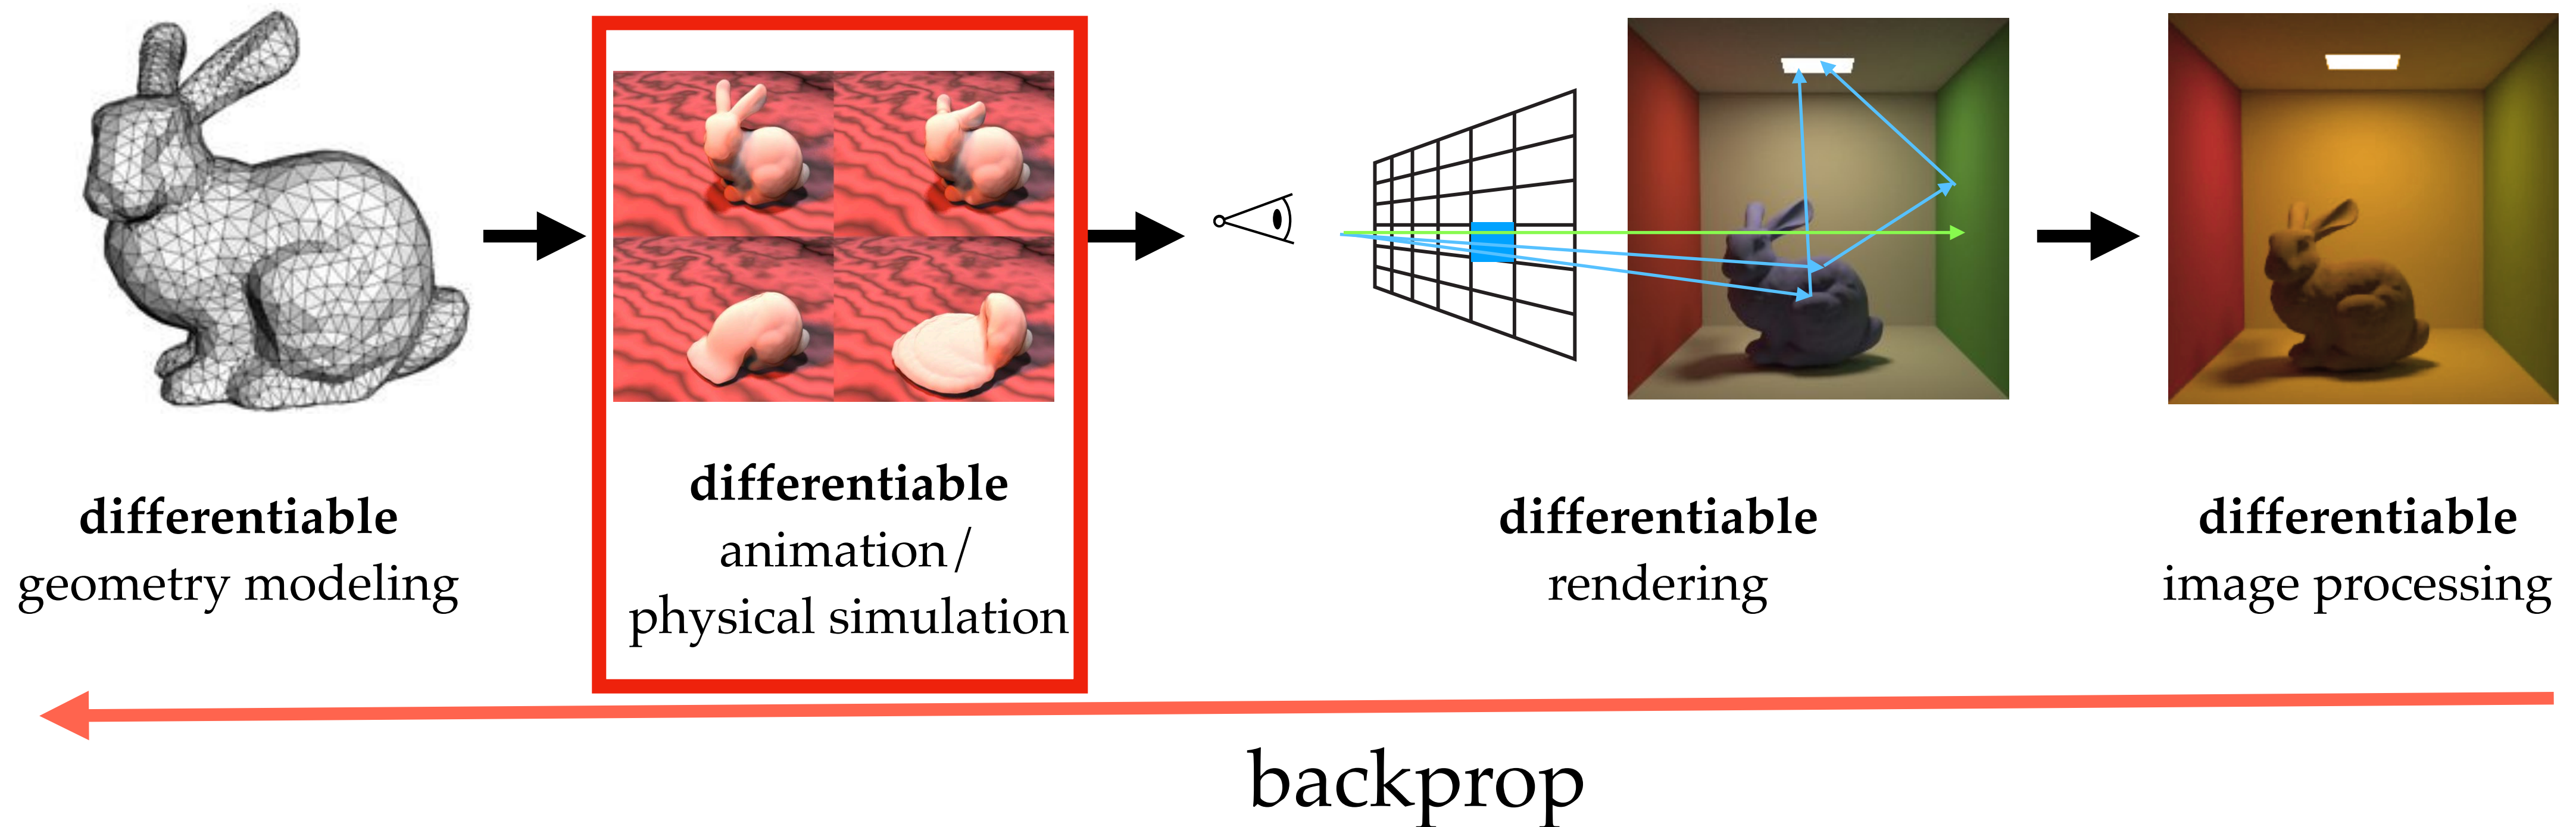
18 pipelines



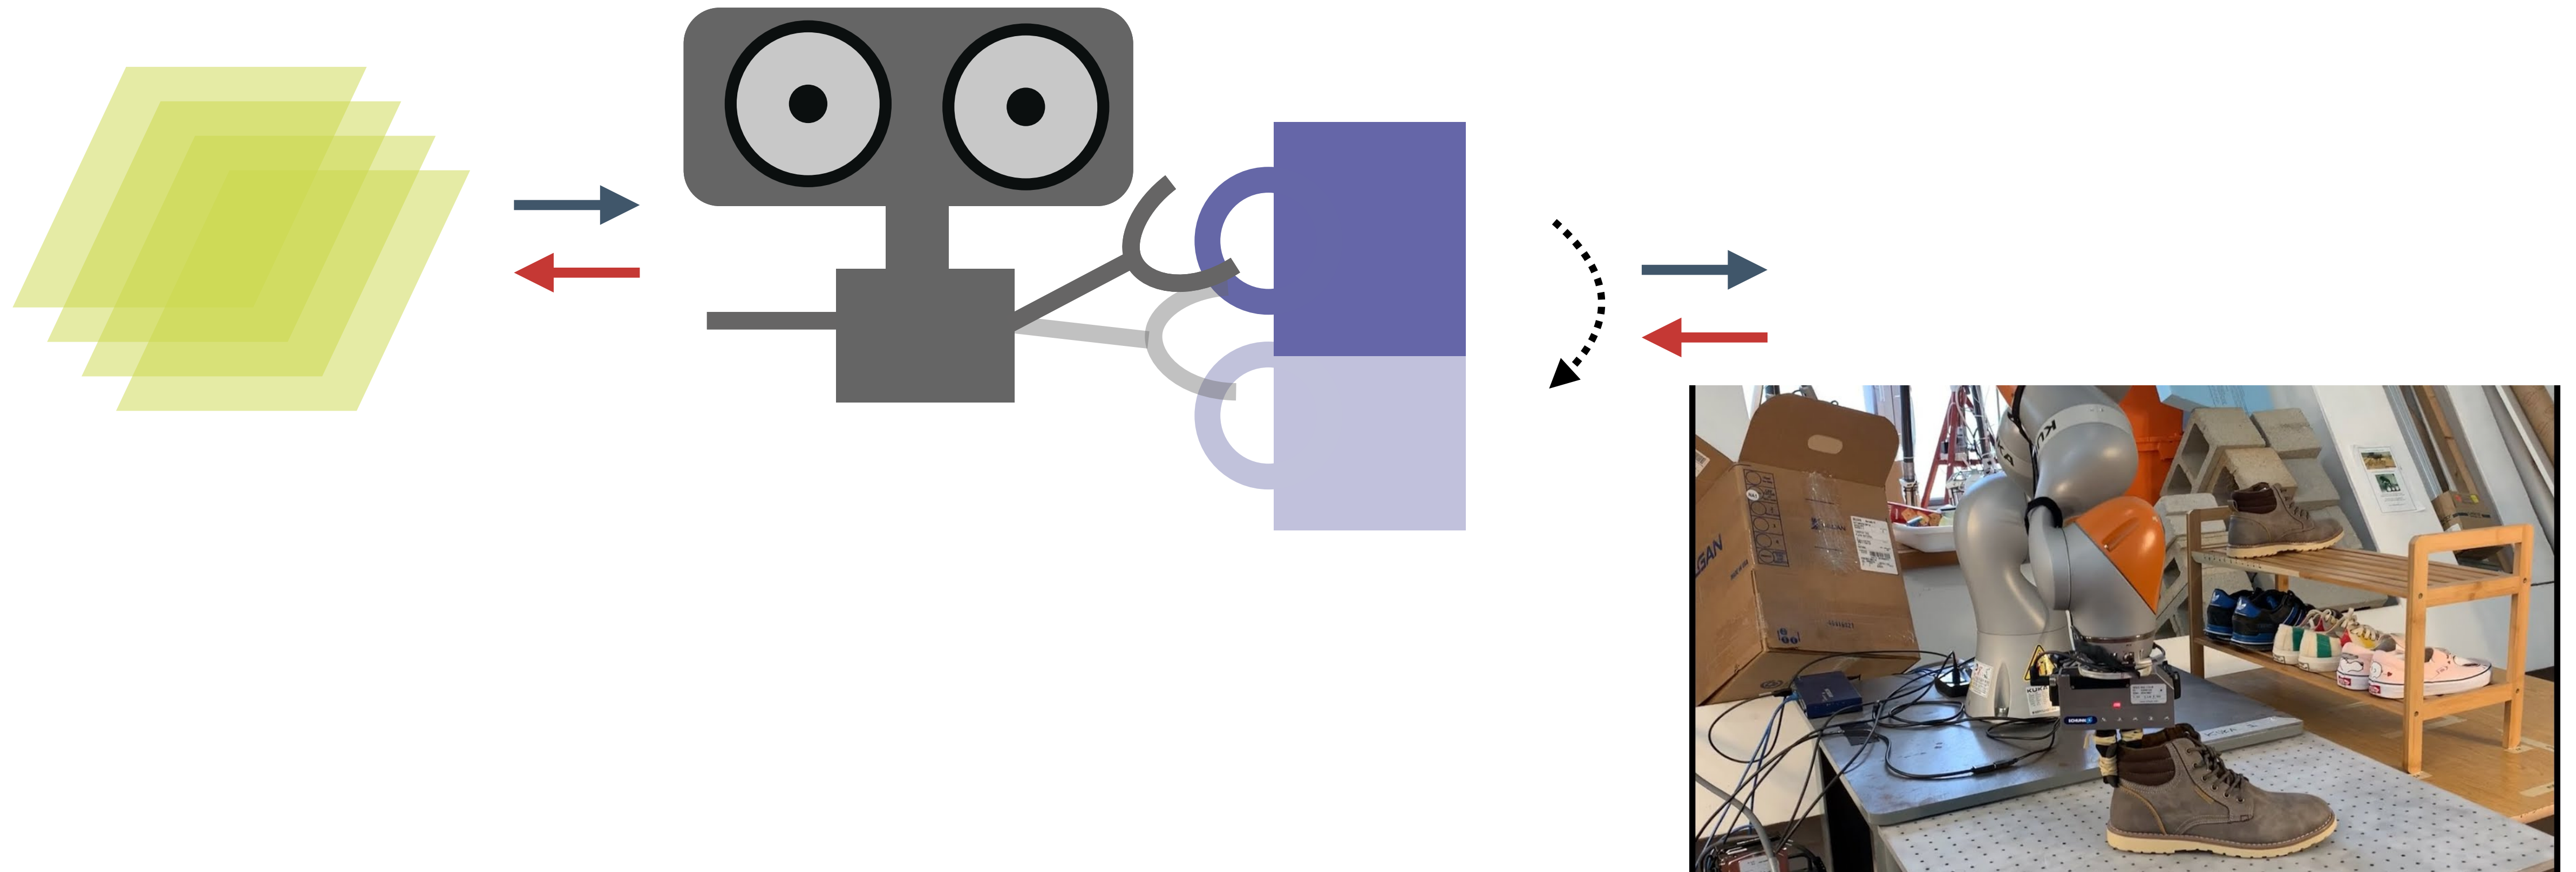
CAT

Differentiable graphics

connects classical graphics algorithms with modern data-driven methods **through derivatives**




Differentiating physics enables intelligent decision




(not my work!)

Challenges: 3D scalability & discontinuities

large-scale physical simulation
requires sparse data structure




512x512x512




3000x2400x1600

physical phenomenon such as
contact leads to discontinuities




Challenges: 3D scalability & discontinuities

large-scale physical simulation
requires sparse data structure




512x512x512




3000x2400x1600

physical phenomenon such as
contact leads to discontinuities



Hierarchical sparse array efficiently models sparsity



Taichi: a differentiating compiler for hierarchical sparse arrays

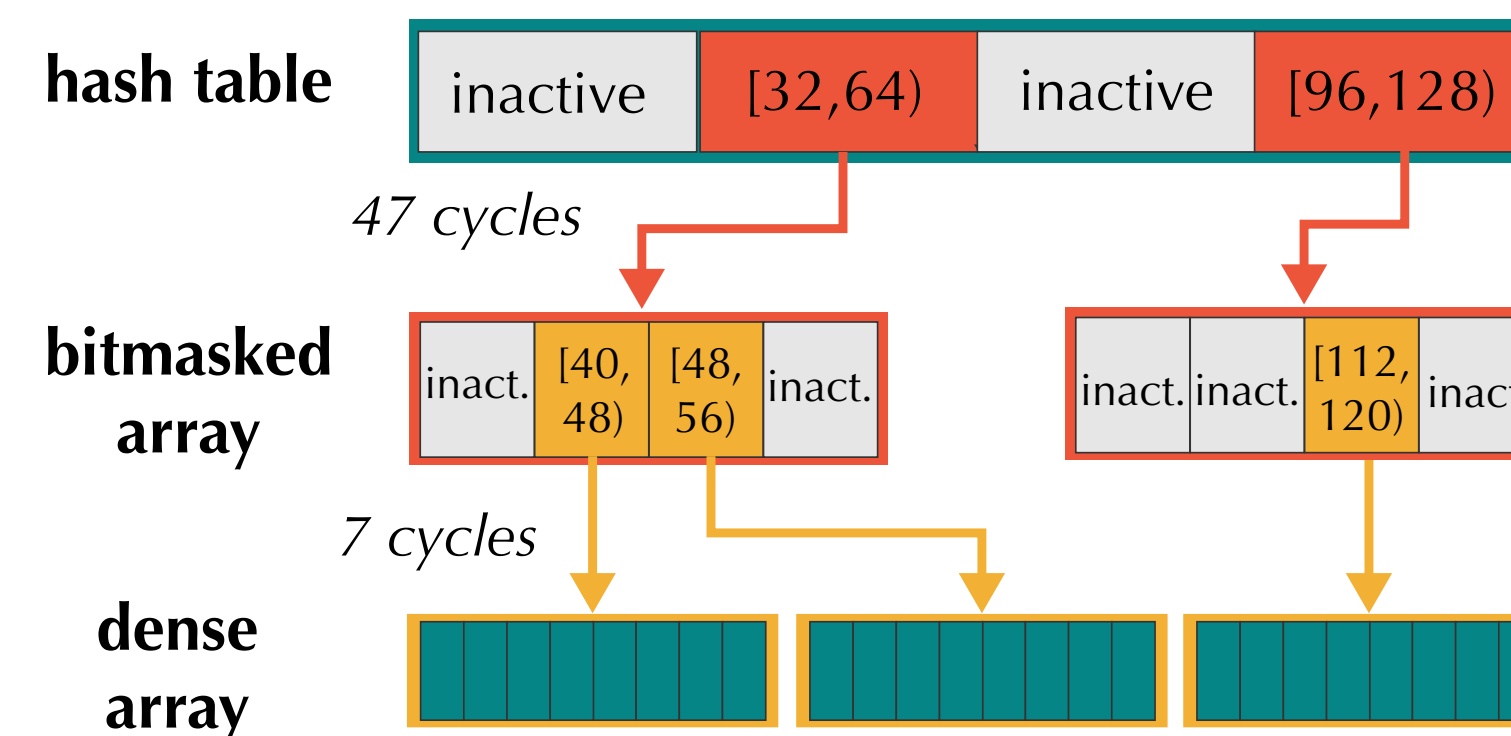
- decouple sparse data structure and access
- compiler outputs optimized code

$f[x, y, z]$

algorithm:

access like dense array

$f.layout =$




implementation:

specify layout separately

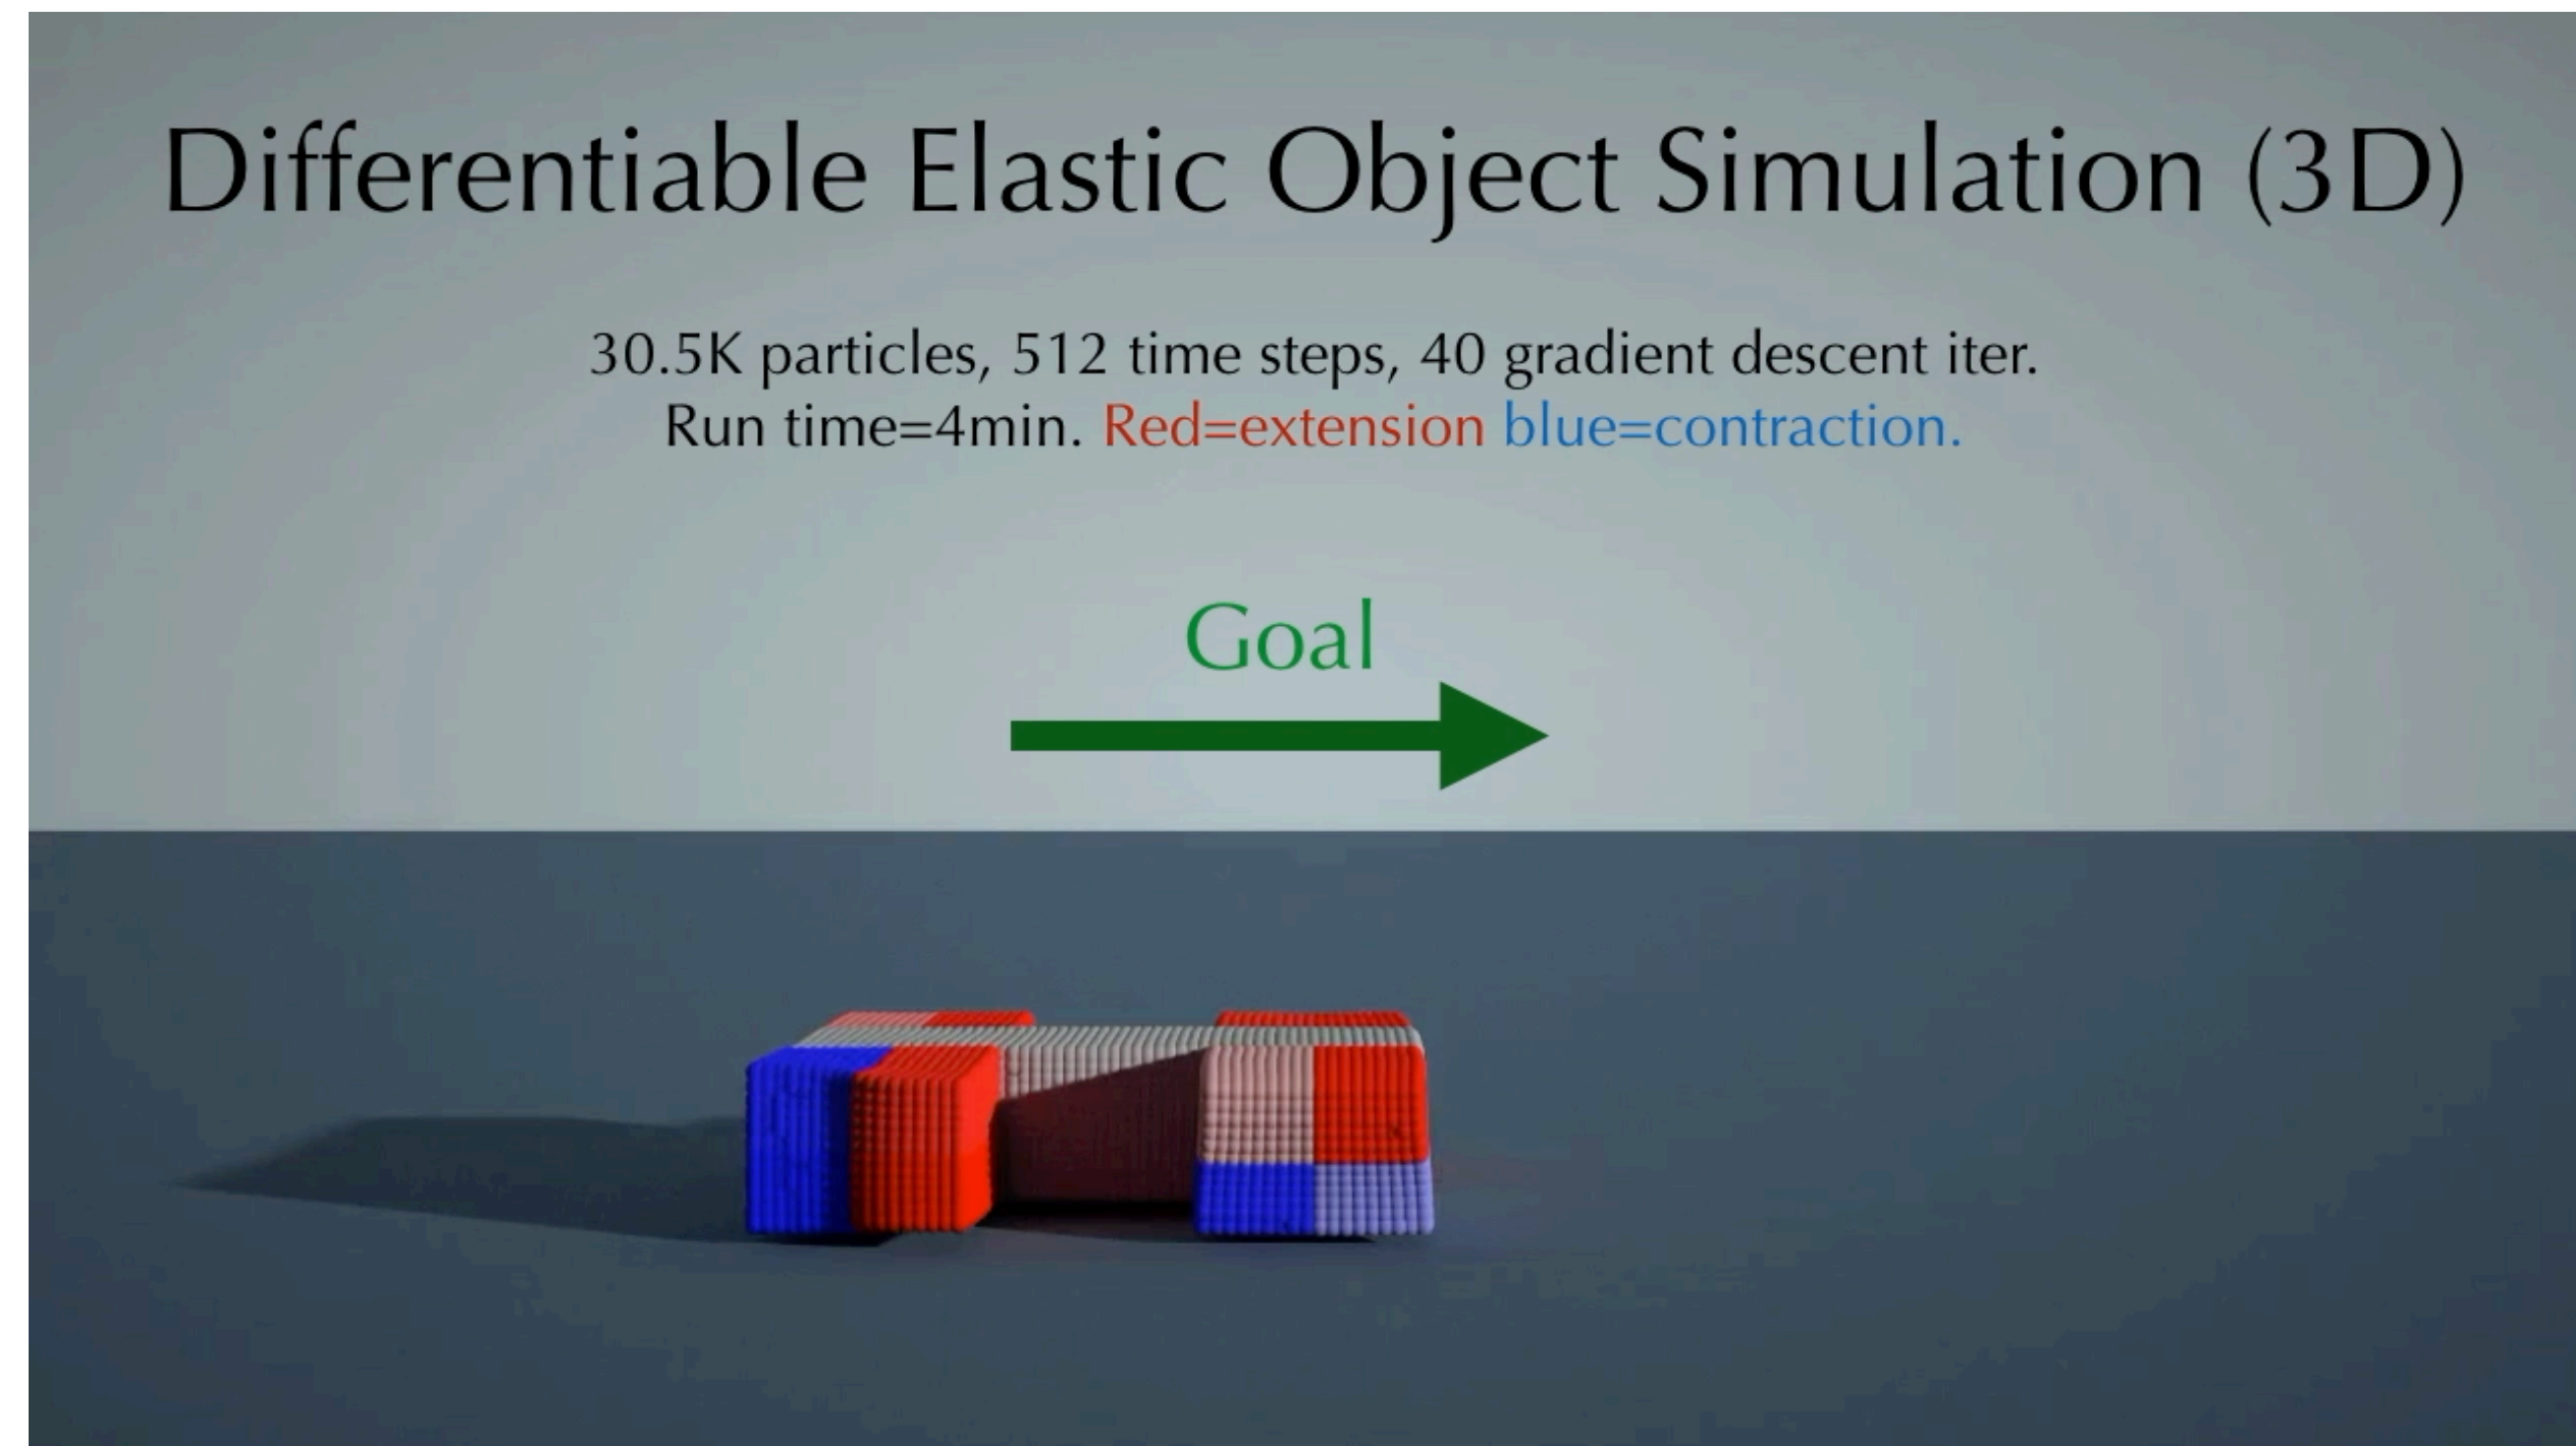
Taichi makes simulation code shorter and faster

- GPU variant of material point method [Gao 2018]
- code 13x shorter, 1.2x faster




First step towards large-scale model-based RL

propagate gradients through simulation to robot controllers




Challenges: 3D scalability & discontinuities

large-scale physical simulation
requires sparse data structure




512x512x512



3000x2400x1600

physical phenomenon such as
contact leads to discontinuities

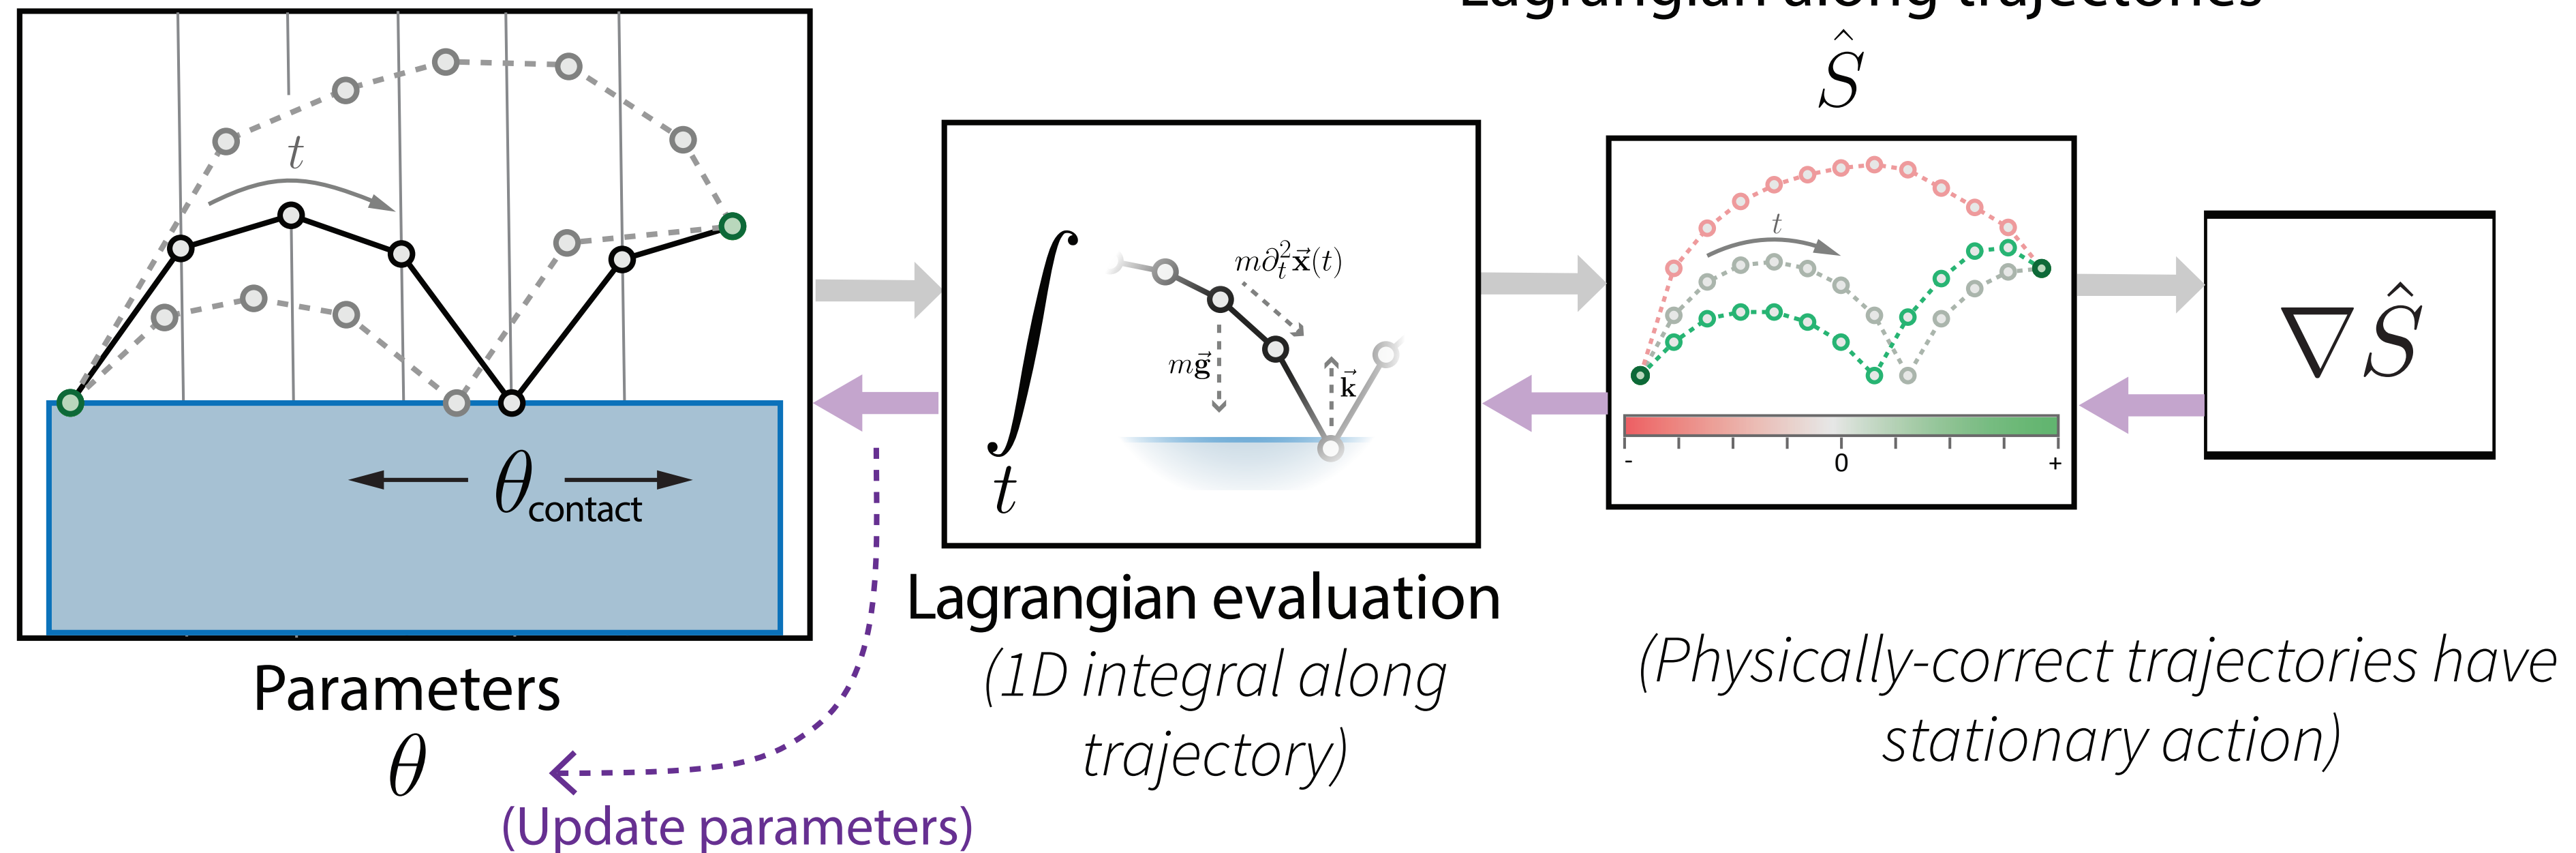


Many physics problems can be modeled as integrals

just like rendering!

$$S(\theta) = \int_{t=t_0}^{t_1} L(q, q_t, t)$$

$$L = [q.y > 0] * m * g * q.y - [q.y \leq 0] * m * f_c * q.y + \frac{1}{2} * m * q_t^2$$



Need a language for describing and differentiating integrals

- if statements = discontinuities

$$\frac{\partial}{\partial p} \int_{x=0}^{x=1}$$


```
def f(x, p):  
    if g(x, p) > 0:  
        return x * x  
    else:  
        return x
```

algorithm:
integrals & derivatives


$$\int \dots \rightarrow \Sigma \dots$$

implementation:
integral discretization

Application: animation design / motion planning

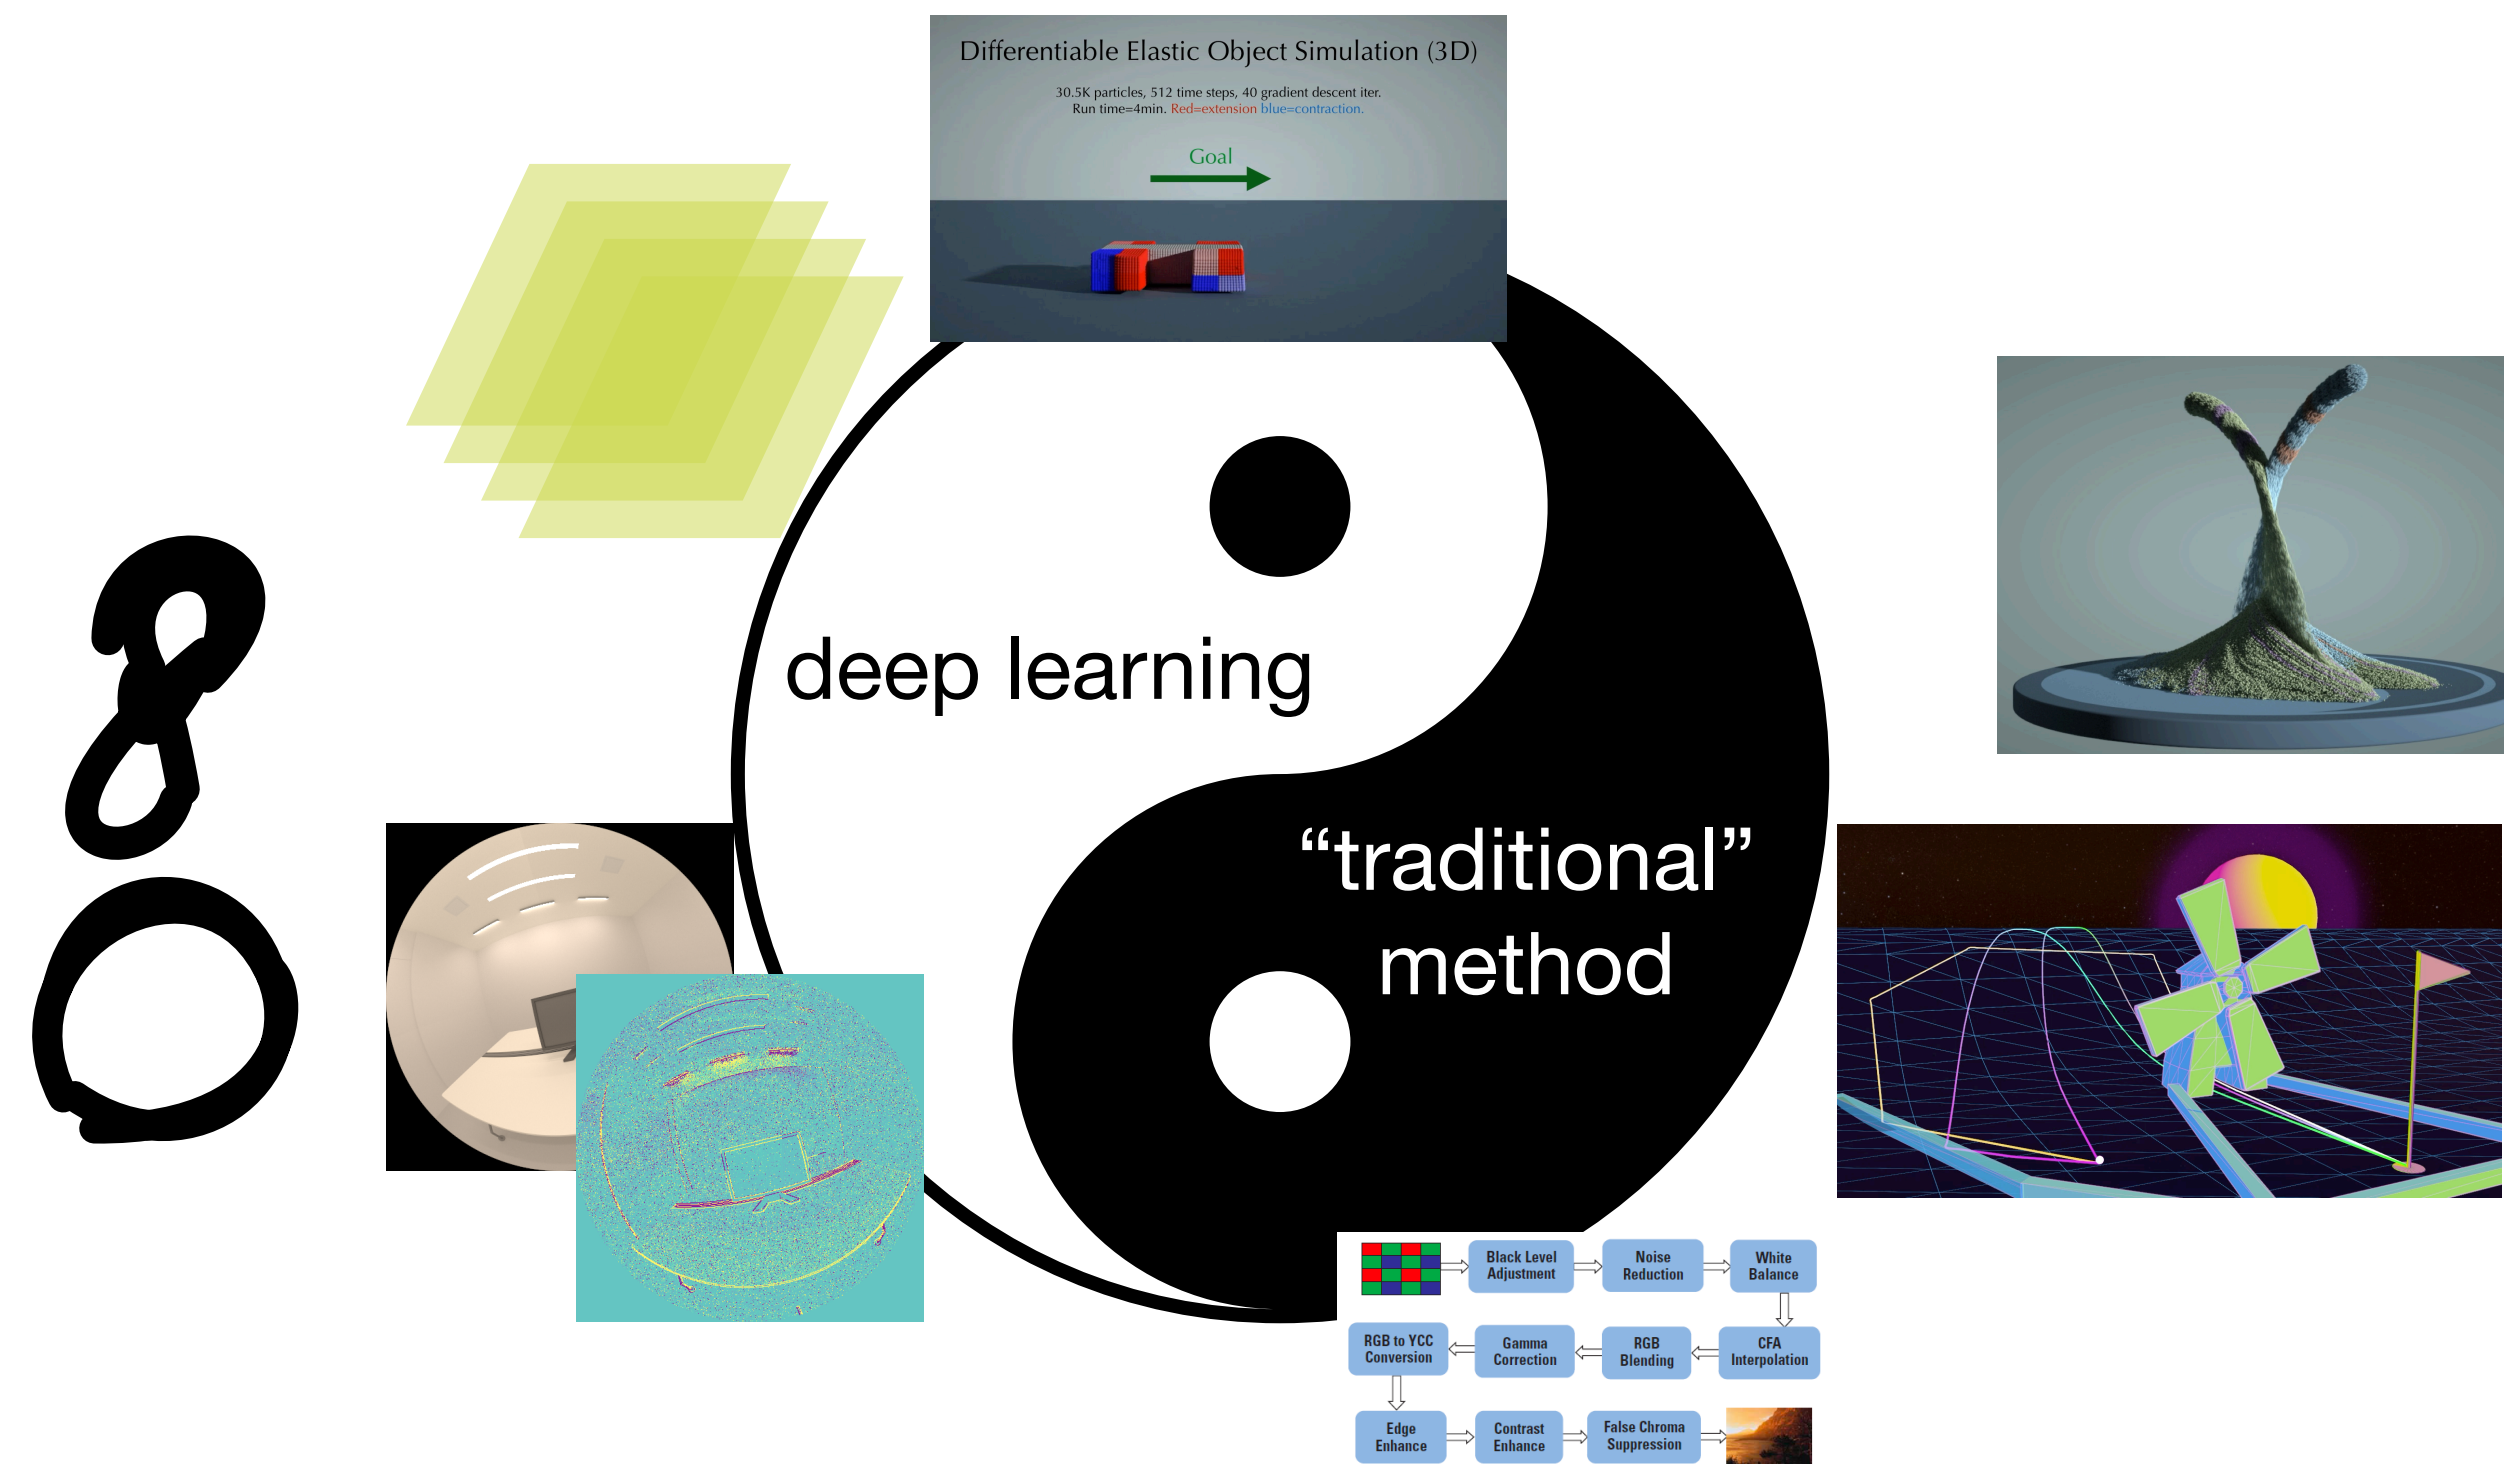


Application: animation design / motion planning



Vision

- differentiable programs bridge deep learning and traditional methods
- our work: extract domain knowledge encoded in graphics algorithms
- tools: differential calculus and PL/compiler



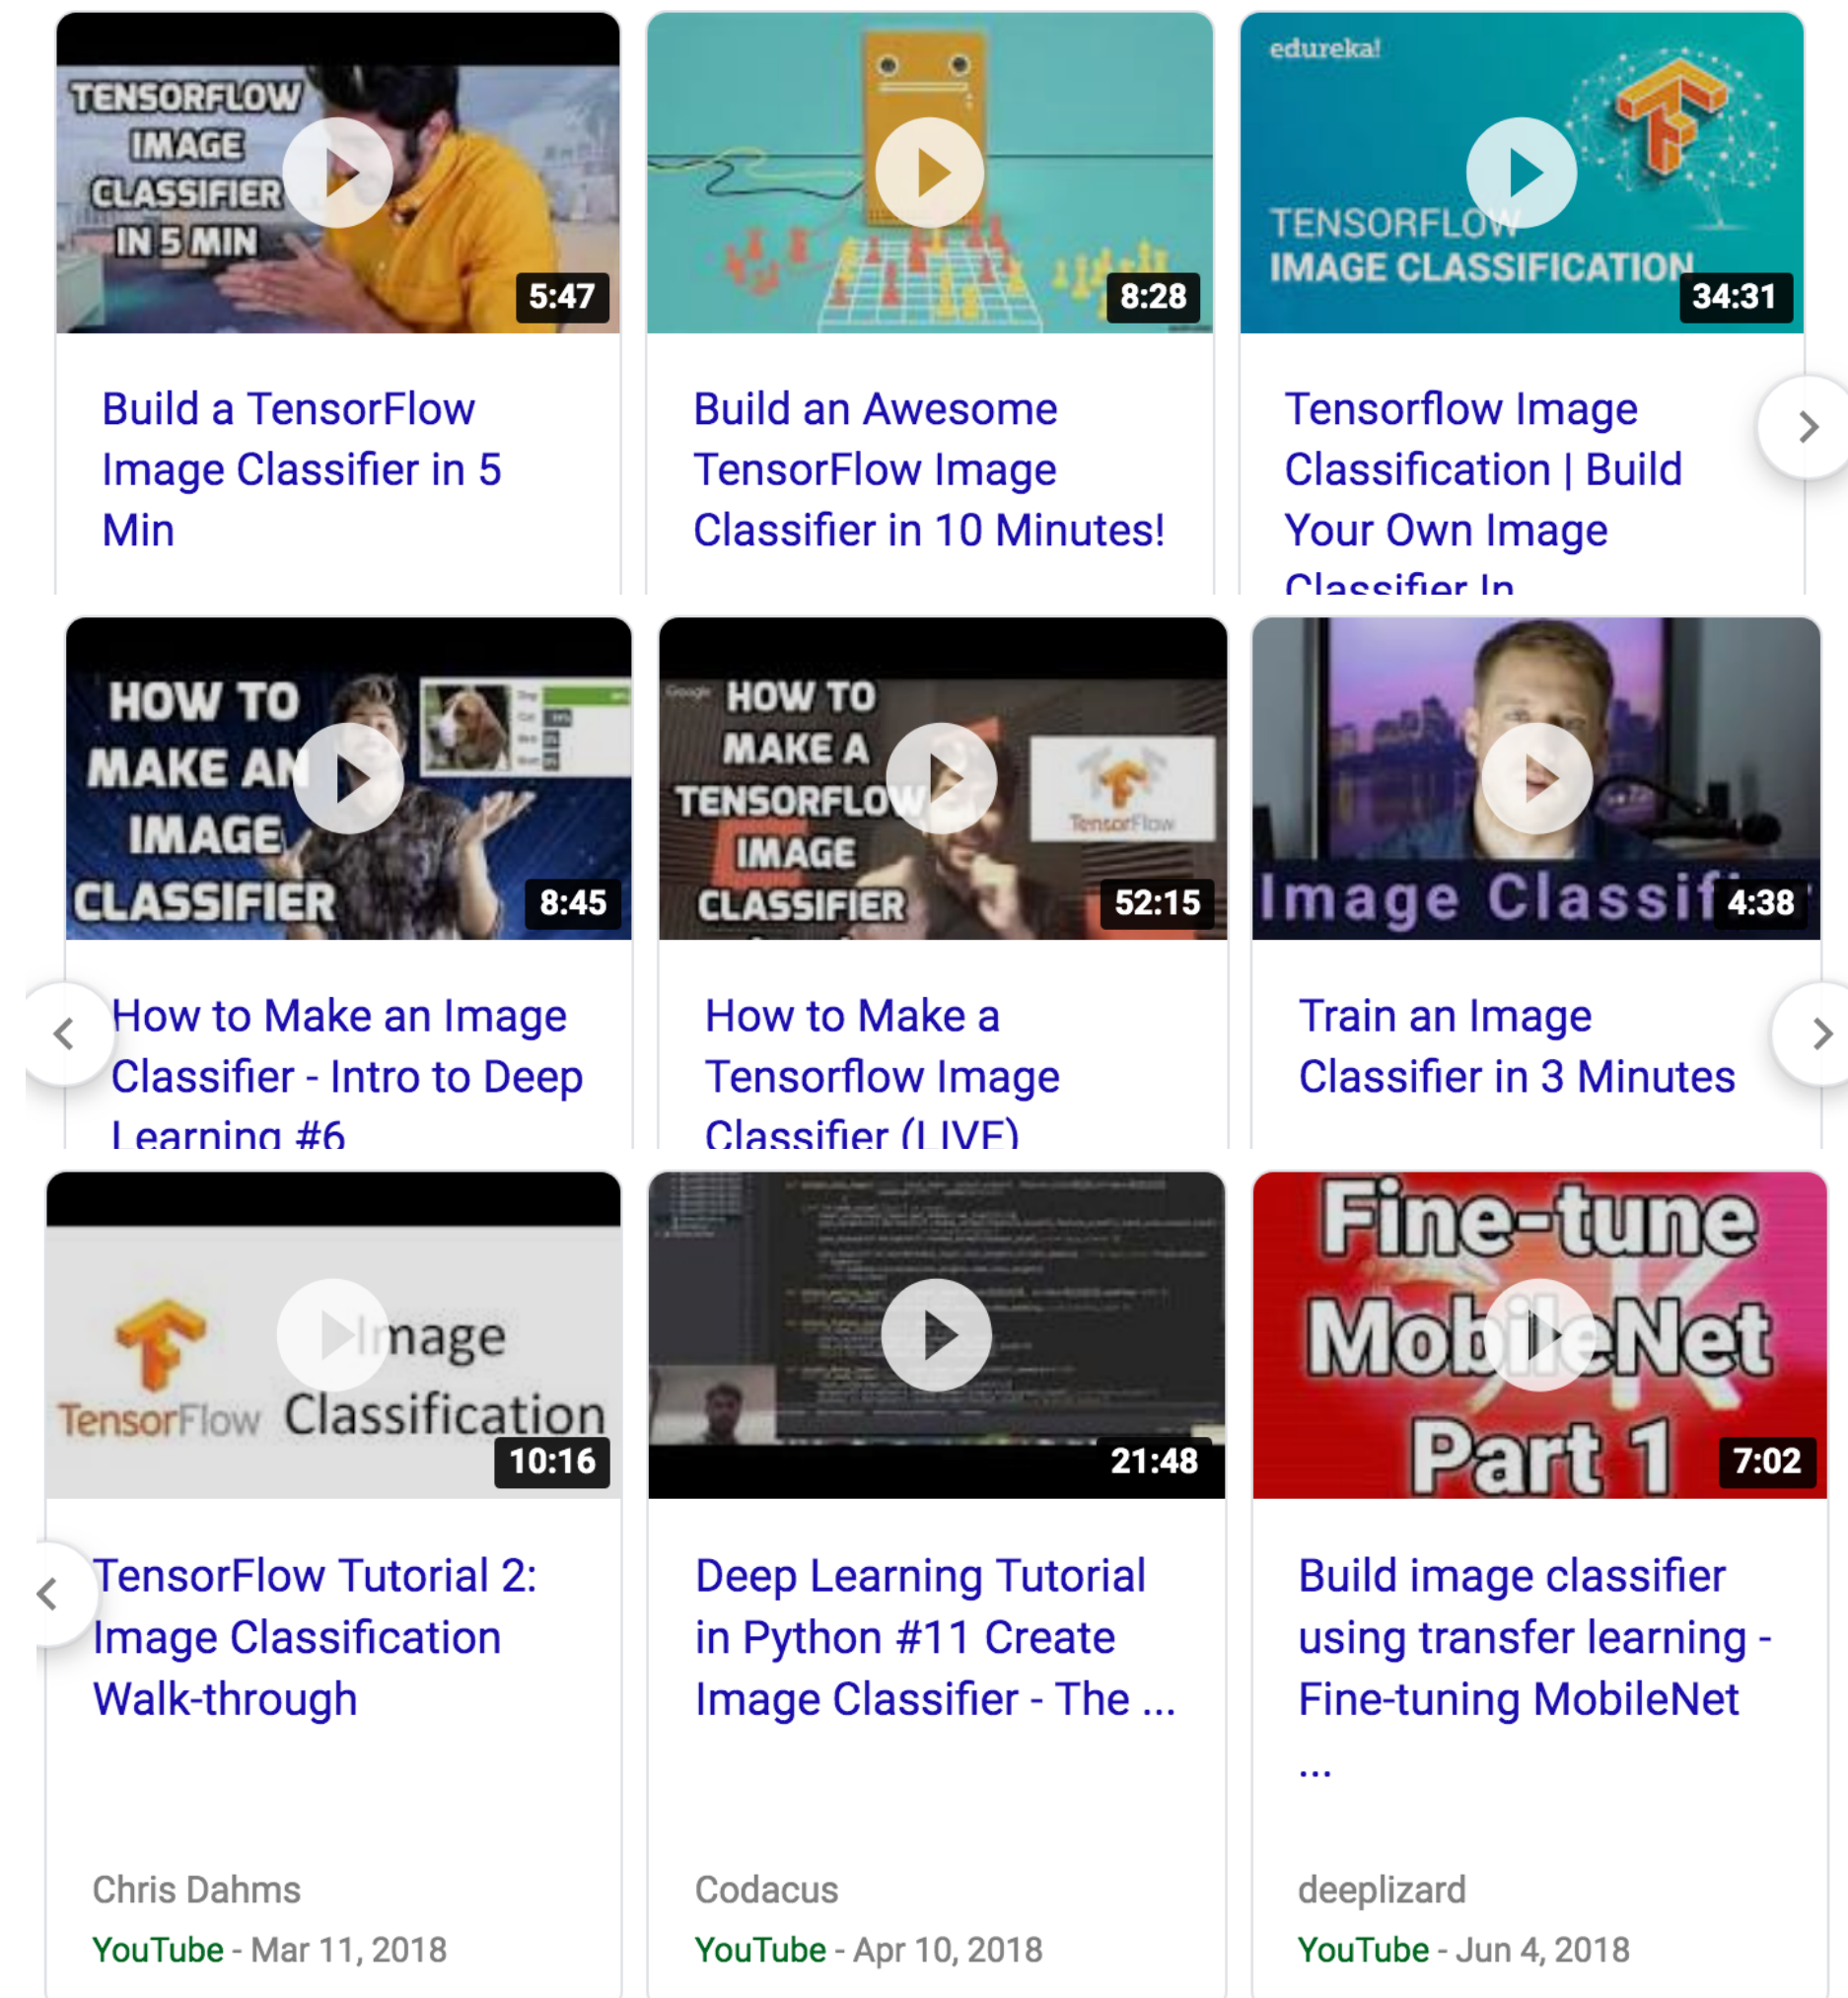
Why was deep learning successful?

Before

required expertise:

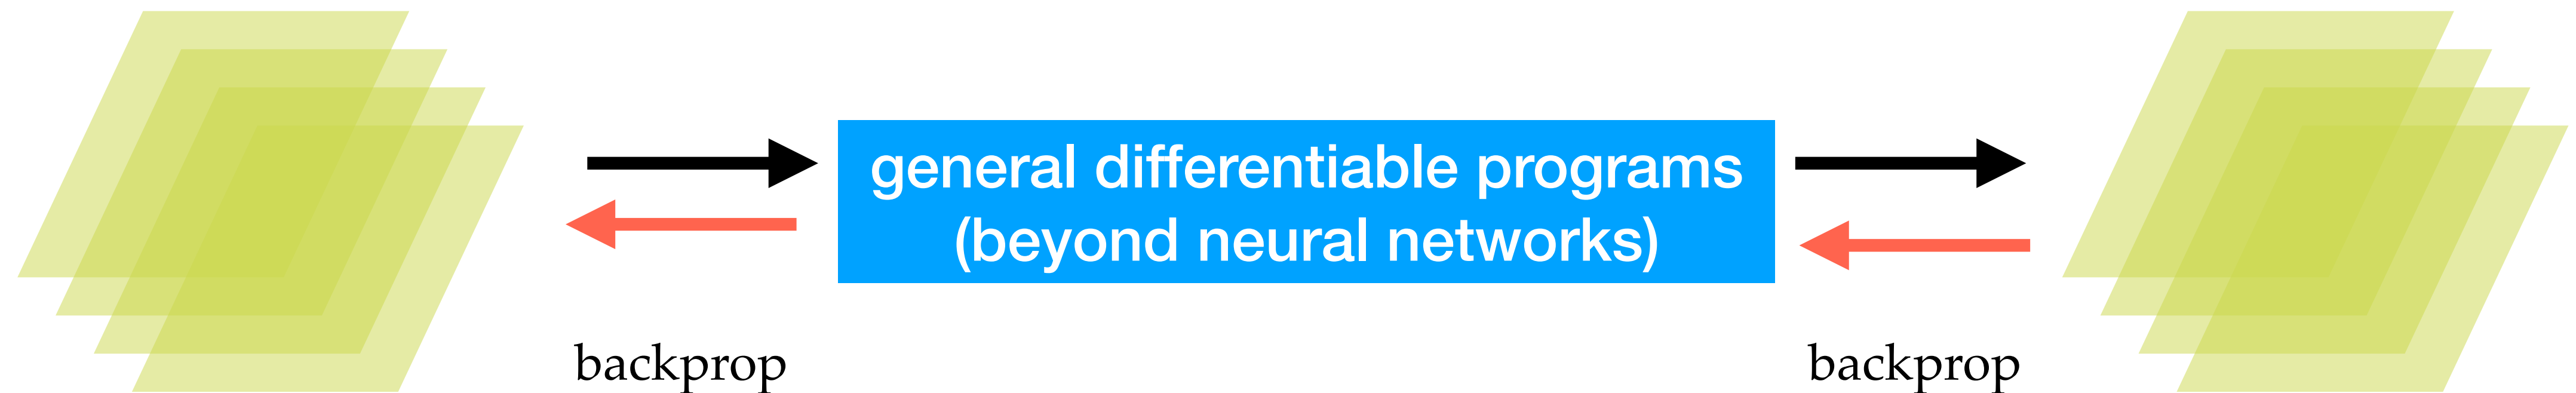
- program differentiation
- numerical computing
- GPU hacking
- optimization

After



Differentiable programming: next-gen deep learning

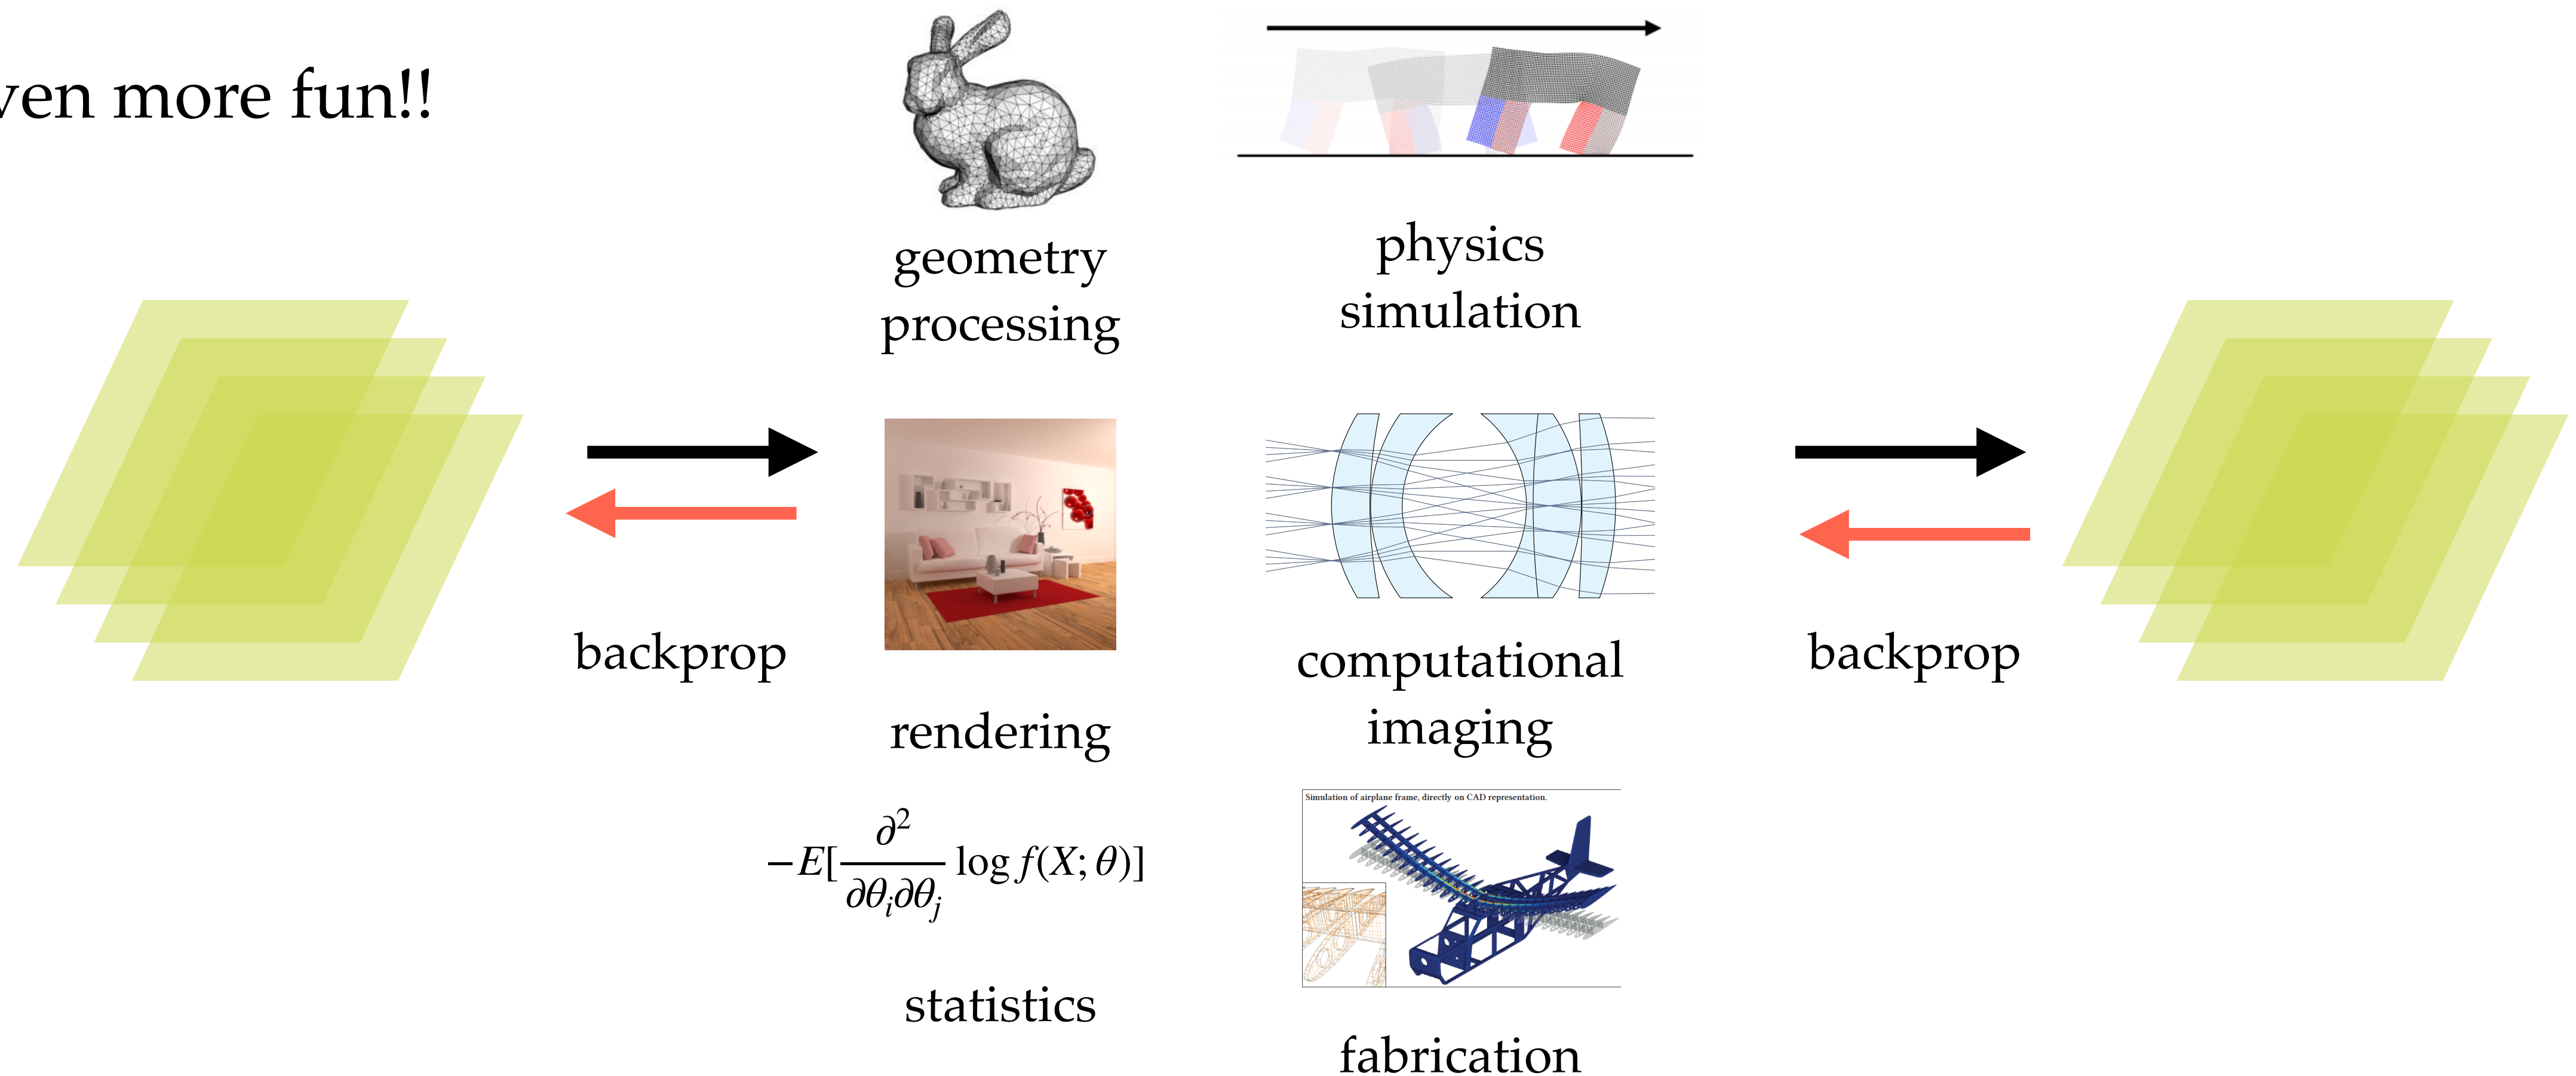
- interpretable, controllable, generalizable
- and fun!



Differentiable visual computing: next-gen graphics / vision

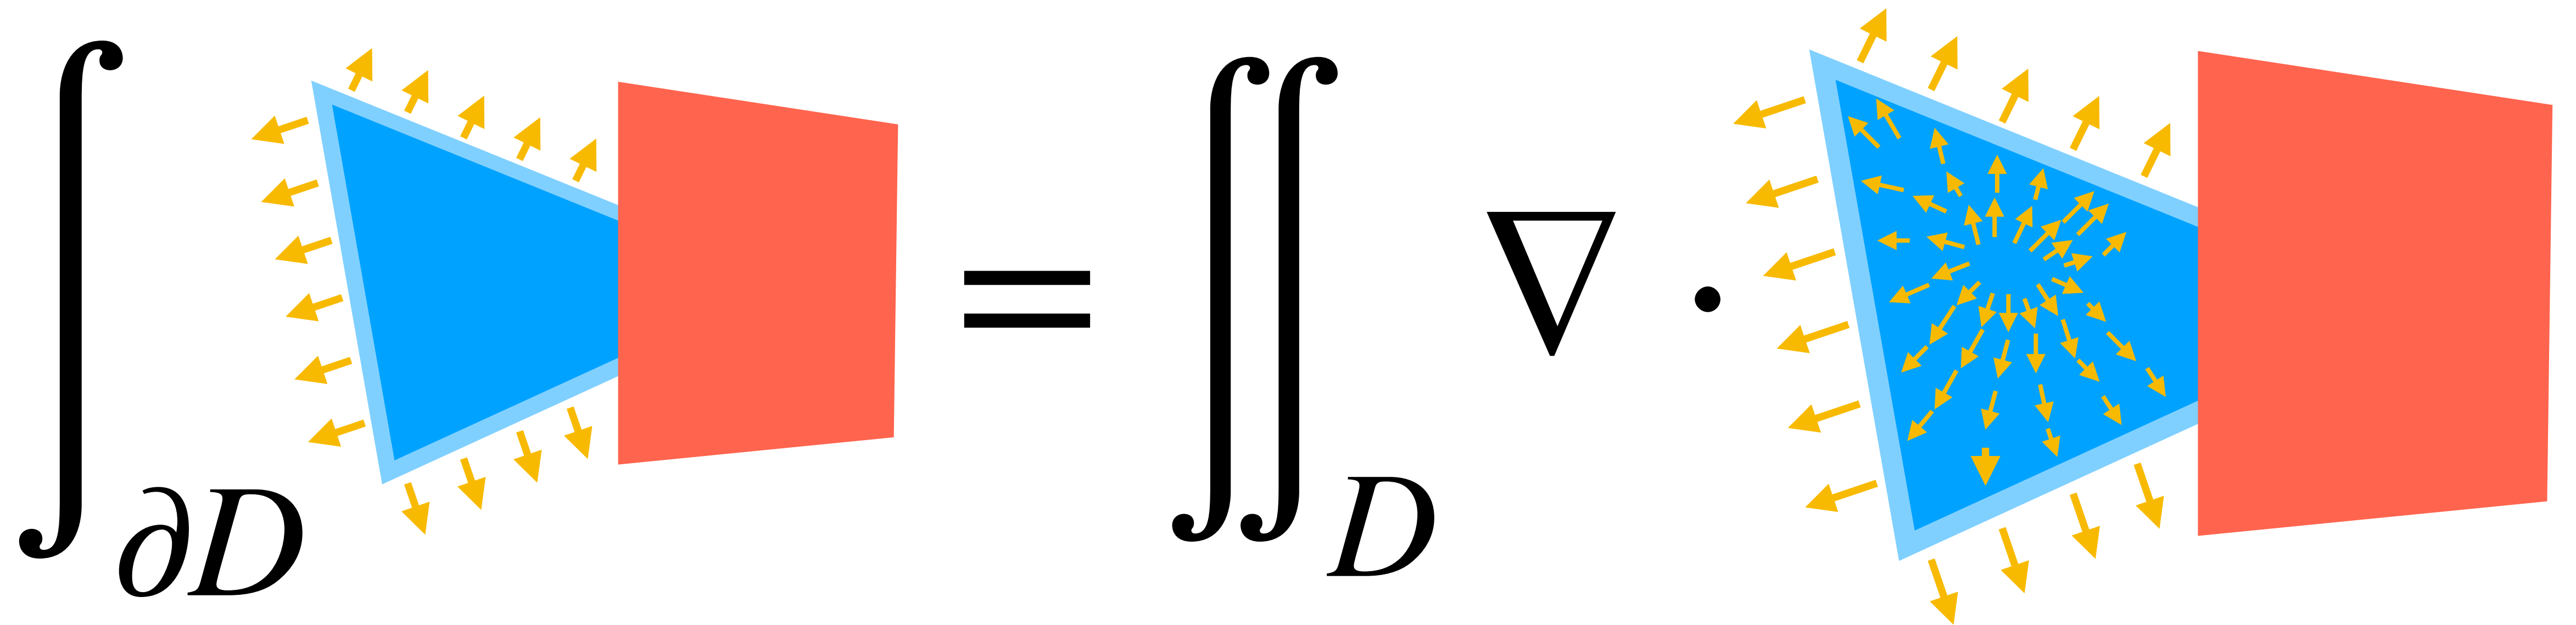
- brings physical inductive bias to graphics / vision systems

- even more fun!!



We need more theorems / algorithms

- for deriving & connecting different differential quantities
- for analyzing pros / cons of different estimators

$$\int_{\partial D} \mathbf{v} \cdot \mathbf{n} \, dS = \iint_D \nabla \cdot \mathbf{v} \, dV$$


We need better programming languages

- for describing programs that involve integrals and differentiation

$$\frac{\partial}{\partial p} \int_{x=0}^{x=1} \text{def } f(x, p):$$

```
    if g(x, p) > 0:
        return x * x
    else:
        return x
```


We need better compilers

- for efficiently mapping programs to hardware


$$\frac{\partial}{\partial p} \int_{x=0}^{x=1}$$

```
def f(x, p):  
    if g(x, p) > 0:  
        return x * x  
    else:  
        return x
```

→



We need our Tensorflow / PyTorch for differentiable visual computing




geometry processing




rendering

$$-E\left[\frac{\partial^2}{\partial\theta_i\partial\theta_j} \log f(X; \theta)\right]$$


statistics



physics simulation



computational imaging



fabrication

V.S.

Build a TensorFlow Image Classifier in 5 Min (5:47)

Build an Awesome TensorFlow Image Classifier in 10 Minutes! (8:28)

Tensorflow Image Classification | Build Your Own Image Classifier In (34:31)

How to Make an Image Classifier (8:45)

How to Make a TensorFlow Image Classifier (52:15)

Image Classifi (4:38)

How to Make an Image Classifier - Intro to Deep Learning #6

How to Make a Tensorflow Image Classifier (LIVE)

Train an Image Classifier in 3 Minutes

TensorFlow Classification (10:16)

Image Classification Walk-through

Deep Learning Tutorial in Python #11 Create Image Classifier - The ... (21:48)

Fine-tune MobileNet Part 1 (7:02)

Build image classifier using transfer learning - Fine-tuning MobileNet ...

Chris Dahms YouTube - Mar 11, 2018

Codacus YouTube - Apr 10, 2018

deeplizard YouTube - Jun 4, 2018