Realistic Image Synthesis

Philipp Slusallek
Karol Myszkowski
Gurprit Singh

Personnel

- Instructors:
  - Philipp Slusallek
    - http://graphics.cg.uni-saarland.de/slusallek/
  - Karol Myszkowski
    - http://www.mpi-inf.mpg.de/~karol/
  - Gurprit Singh
    - http://people.mpi-inf.mpg.de/~gsingh/
- Teaching Assistant:
  - Pascal Grittmann
    - https://graphics.cg.uni-saarland.de/people/grittmann.html
- Secretary:
  - Sabine Nermerich
    - https://graphics.cg.uni-saarland.de/people/nermerich.html

Administrative Information

- Type
  - Special lecture
  - Applied computer science (Praktische Informatik)
- ECTS
  - 9 credit points
- Prerequisites
  - Interest in mathematics, physics, some programming experience in C++
- Language
  - All lectures will be given in English
- Time and Location
  - Monday & Thursday, 10-12h, online via Zoom (maybe later also HS 03, E1.3)
- Web-Page
  - http://graphics.cg.uni-saarland.de/courses/ris-2020/
  - Schedule, slides as PDF, link to videos to watch again later
  - Literature, assignments, other information
- Mailing list
  - Up-to-date information, exercise updates, etc…
  - Sign up for the course on the Web page (if not done yet)
- Please also do not forget to sign up on LSF for the course

Grading

- Weekly assignments
  - Average of at least 50% of all assignments in the semester
  - Required for admission to final exam
  - Demonstrate your solution in exercise groups
  - Can be done in groups of up to two
- Practical assignments
  - Longer-term projects
  - Gradually building your own physically-based renderer
  - Can be done in groups of up to two
- Final grade
  - Assignments: 50%
  - Final oral exam: 50%

Textbooks

- Pharr & Humphreys, **Physically-Based Rendering: From Theory to Implementation**, Morgan Kaufmann, 3nd Edition (Dec 2016), now freely available at http://www.pbr-book.org/), also as e-book in CS library
- Dutre, Bekaert, Bala, **Advanced Global Illumination**, A.K. Peters, 2006, 2nd Edition.
- Jensen, **Realistic Image Synthesis Using Photon Mapping**, A.K. Peters, 2005, 2nd Edition, also see http://graphics.ucsd.edu/~henrik/papers/book
- Shirley & Morley, **Realistic Ray Tracing**, A.K. Peters, 2003, 2nd Ed.
- Reinhard, Ward, Pattanaik, Debevec, Heidrich, Myszkowski**, High Dynamic Range Imaging**, Morgan Kaufmann Publish.,2010, 2nd Ed.

- Cohen & Wallace, **Radiosity and Realistic Image Synthesis**, Academic Press, 1993.
- Apodaca & Gritz, **Advanced Renderman: Creating CGI for the Motion Pictures**, Morgan Kaufmann, 1999.
- Glassner, **Principles of Digital Image Synthesis**, 2 volumes, Morgan Kaufman, 1995.

- Iliyan Georgiev, **Path Sampling Techniques for Efficient Light Transport Simulation**, PhD Thesis, Saarland University, 2015

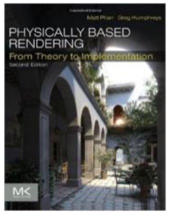

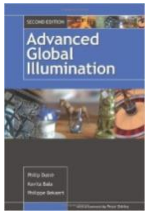

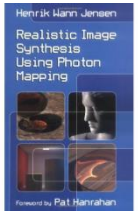

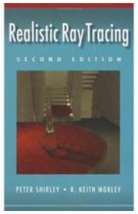

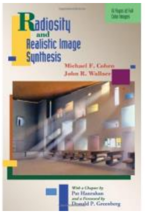

Ingredients for Realistic Images

- *Shape* (Geometry)
  - Objects in our scene: surfaces, volumes, points, ...
- *Material* of surfaces & volumes
  - Places of interaction of light with matter
    - Reflection, refraction, scattering, absorption, ...
  - Applied to shapes (“shaders”)
- *Light sources*
  - Sources of light
    - Positions, color, directional characteristics, …
  - Applied to shapes or independent (“light shaders”)
- *Camera*
  - Sensor that captures the light from the scene
    - Lenses, shutter & film; also surfaces can be sensors: e.g. light maps
- **Simulation of Light Propagation**
  - Computing the distribution of light at the sensor (and thus in scene)

Motivation

- Goal: Create images on the computer that are
  - Indistinguishable from reality typically for a human (but also for sensors!)
    - “(Photo-)Realistic rendering” or “Predictive rendering”
    - Must understand human perception (or sensor characteristics)
  - That convey specific information
    - “Visualization” or “non-photorealistic rendering (NPR)”
- Applications
  - Industrial design
  - Movies and games
  - Architecture and 3D geospatial data
  - Cultural heritage
- Holy Grail: “Digital Reality”
  - Provide simulated reality that feels “real” – for humans & machines
  - All optical (acoustic, haptic, …) features one would perceive in reality
  - Truly convincing real-time simulated reality (aka “Holo-Deck”)
  - Simulation of these models can be used to train computers (AI) to understand and act in the world around us

Applications

- Entertainment Industry: Special effects for motion pictures

Applications

- Entertainment Industry: Animated films

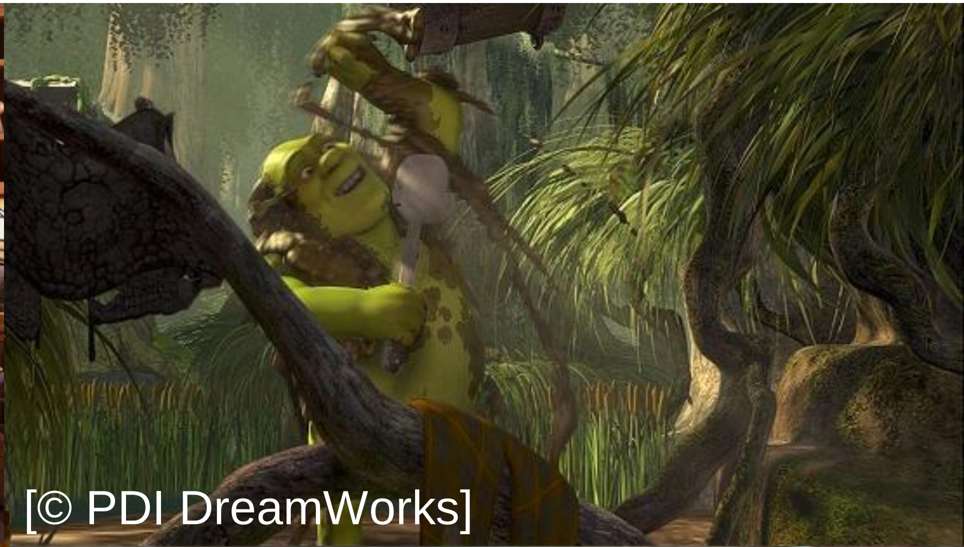

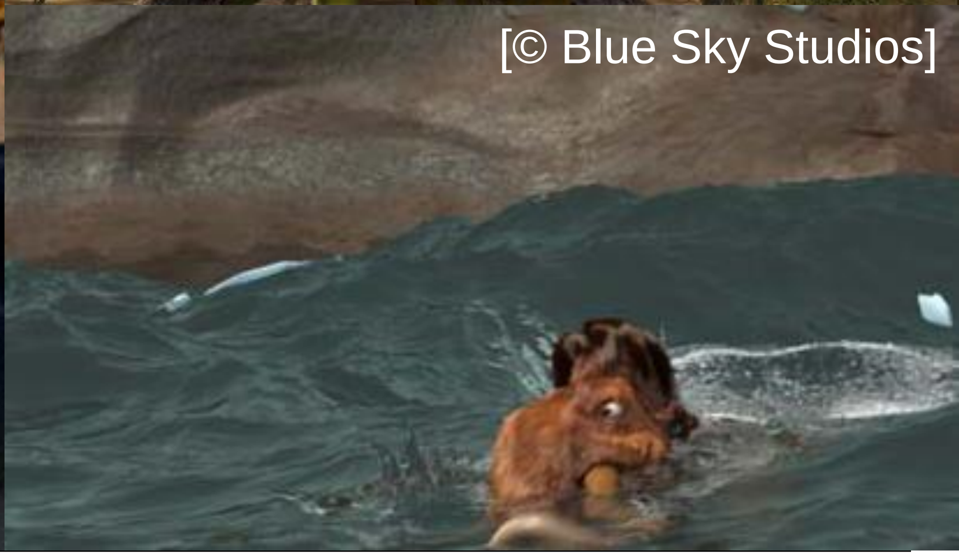

Applications

- Entertainment Industry: Video games

[© Valve]

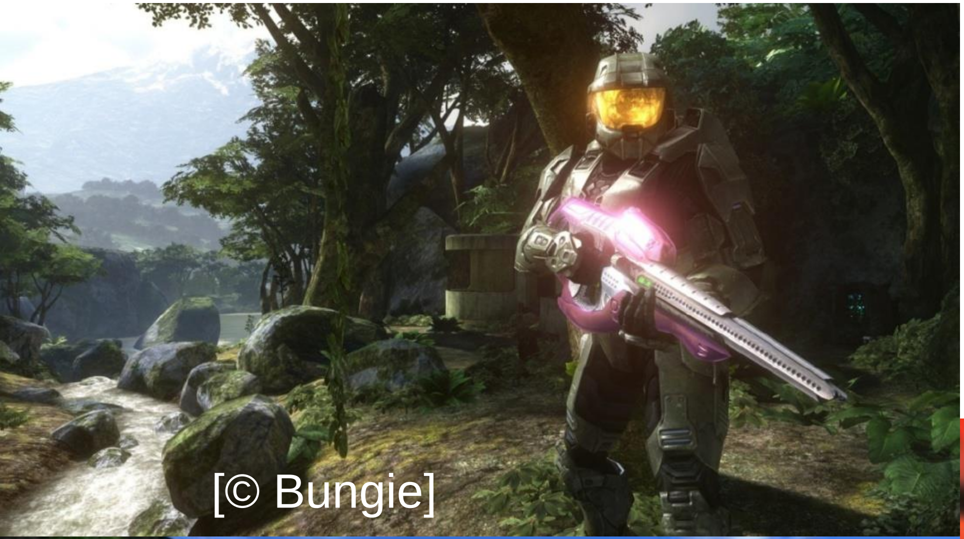

[© Bungie]

[© Crytek]

[© Blizzard Entertainment]

Applications

- Simulation & Augmented Reality

[© NASA]

[© Renault]

[© ENIB]

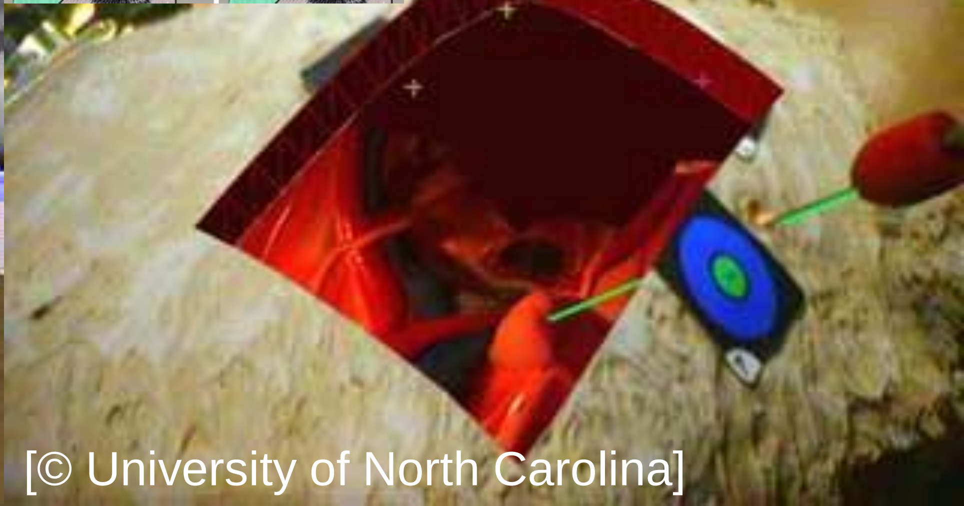

[© University of North Carolina]

Applications

- Industrial Design & Engineering: Automotive / Aerospatial

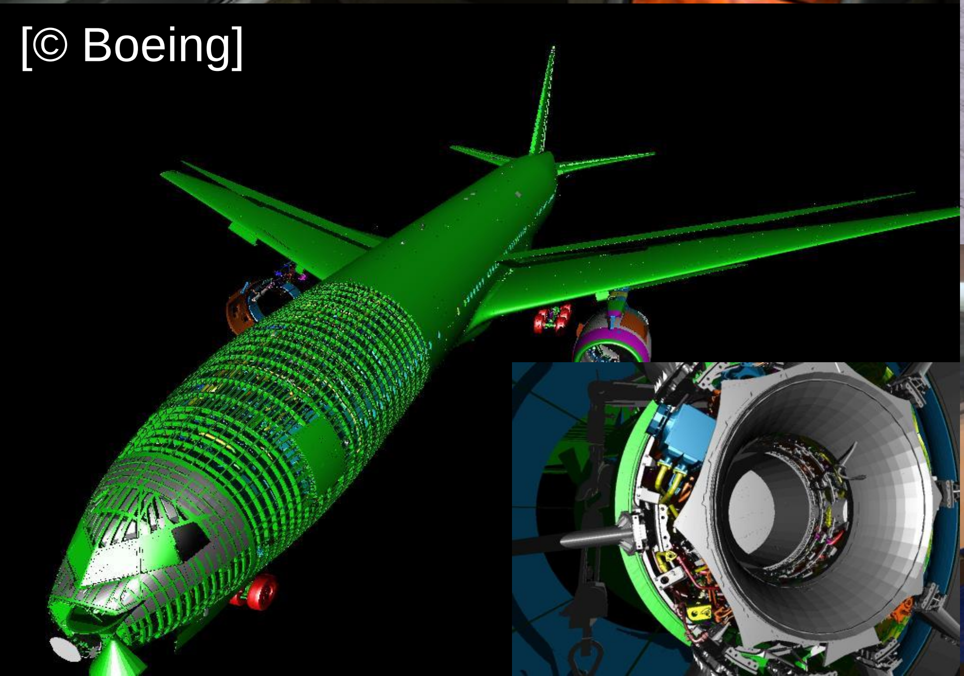

Applications

- Architectural / Interior Design
- Landscape / Urban Planning
- Archeological Reconstruction


[© Saarland University]


[© PBRT]

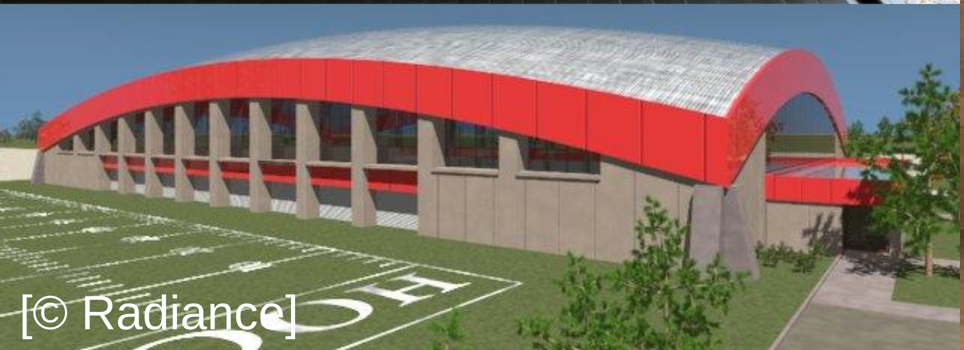
[© Radiance]

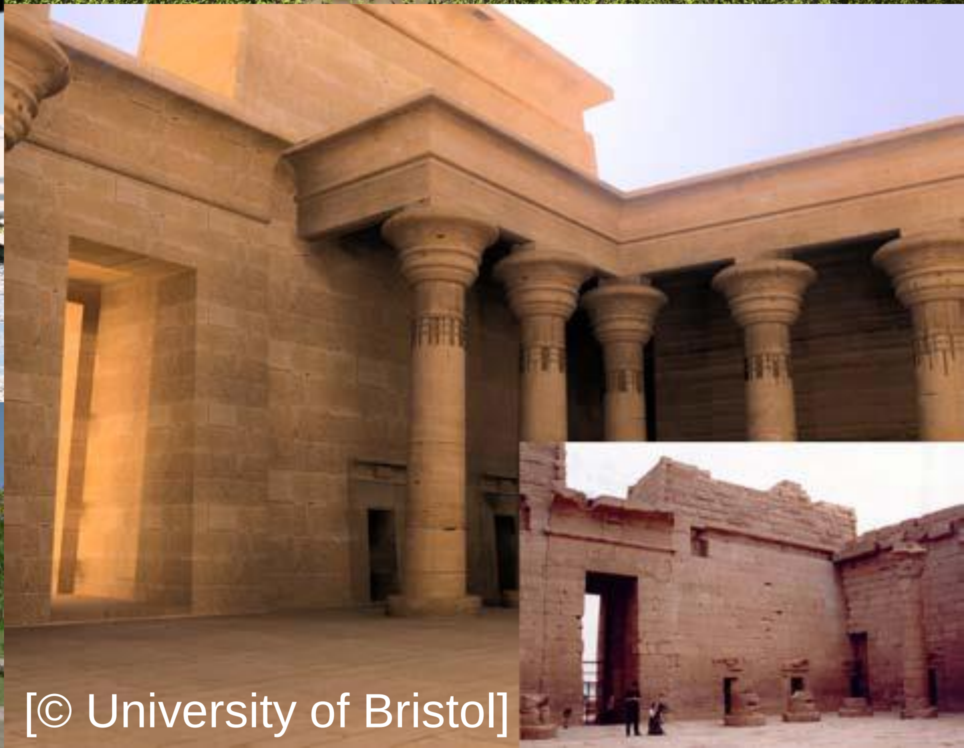
[© University of Bristol]

Digital Reality

- Training and Validation in Reality
  - E.g. driving millions of miles to gather data
  - Difficult, costly, and non-scalable
  - Even millions of miles does not get you a reliable AI system
    - Issue of long-tail distributions (critical scenarios)

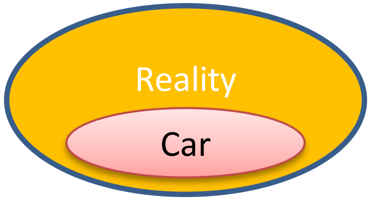

Digital Reality

- Training and Validation in the Digital Reality
  - Arbitrarily scalable (given the right platform)
  - But: Where to get the models and the training data from?

Digital Reality: AI to Certify AI

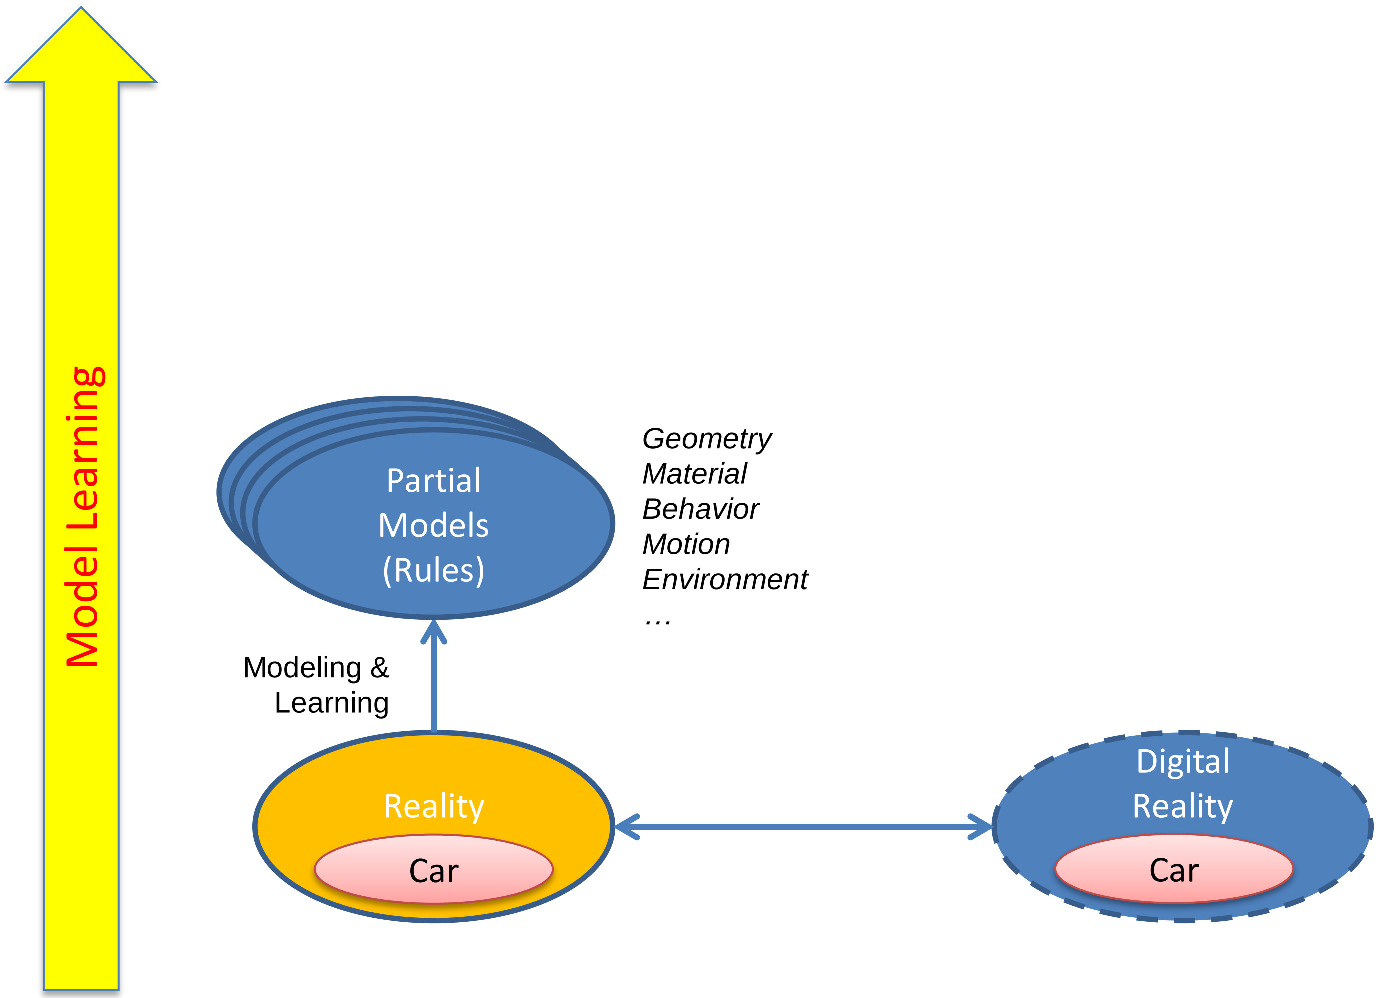

Digital Reality: AI to Certify AI

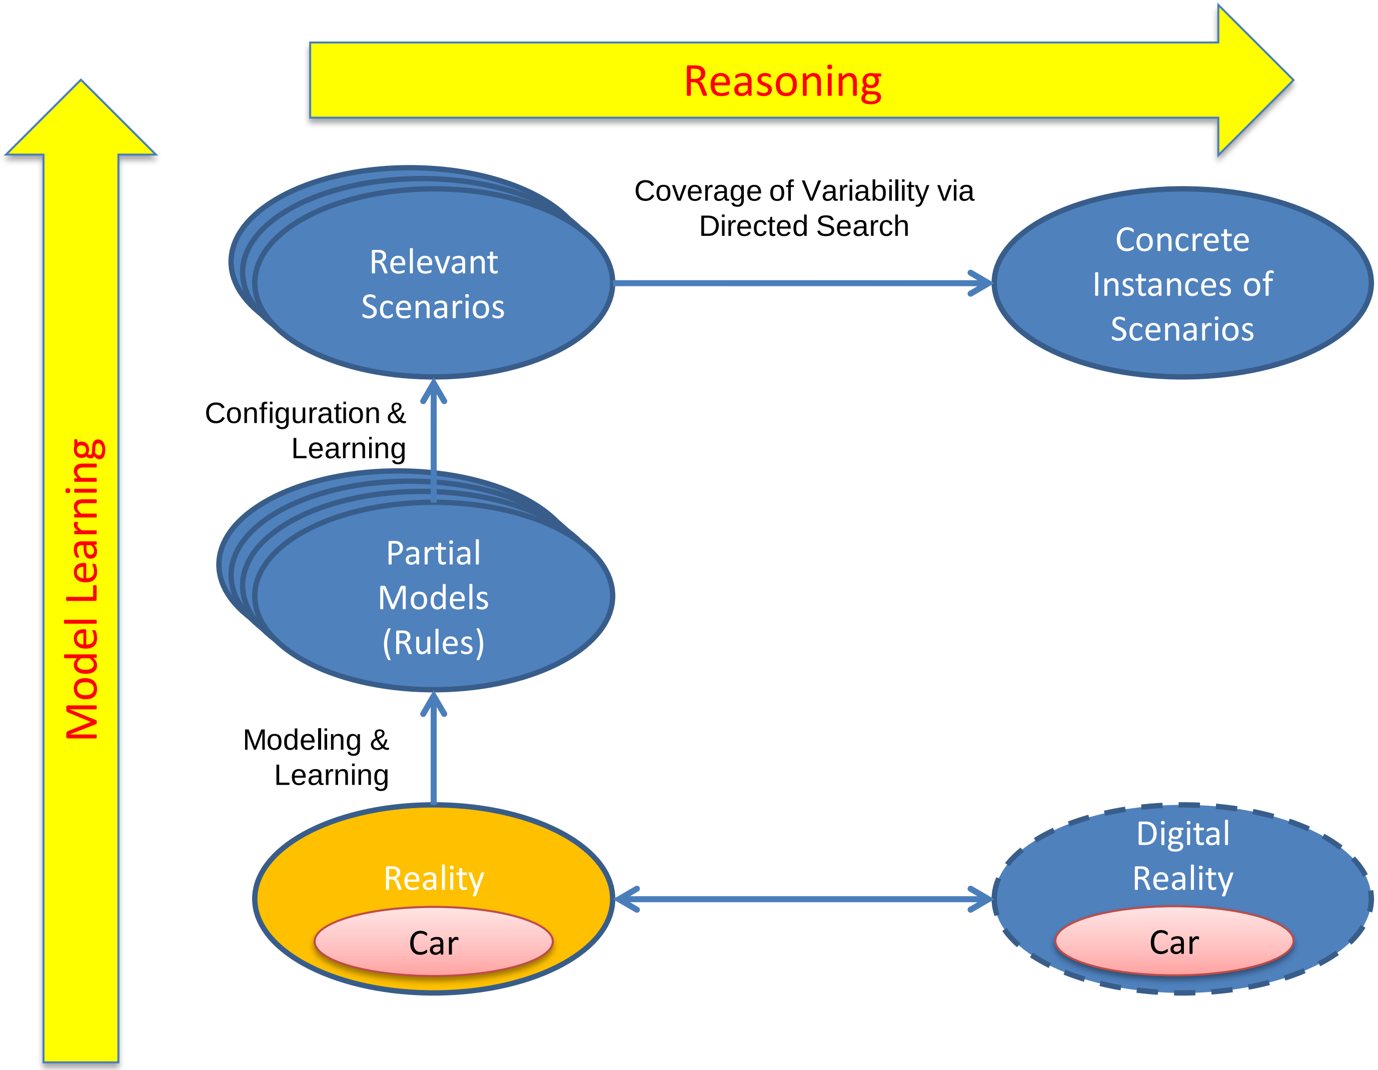

Digital Reality: AI to Certify AI

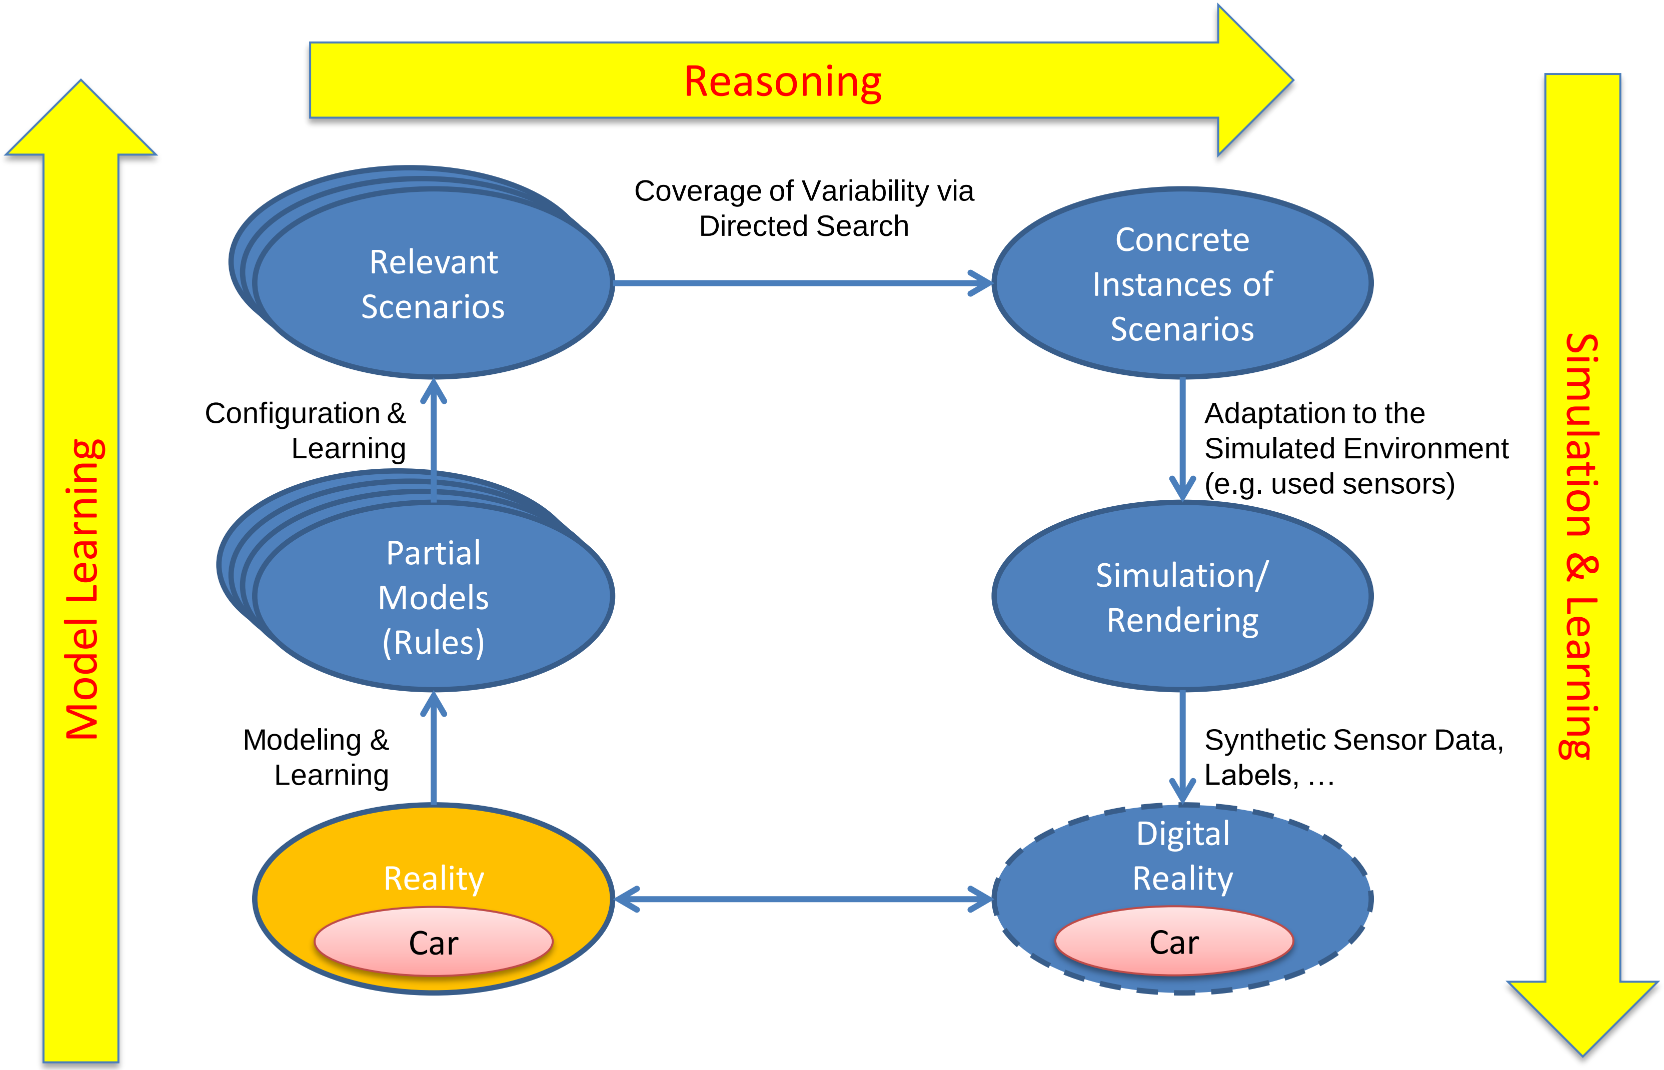

Digital Reality: AI to Certify AI

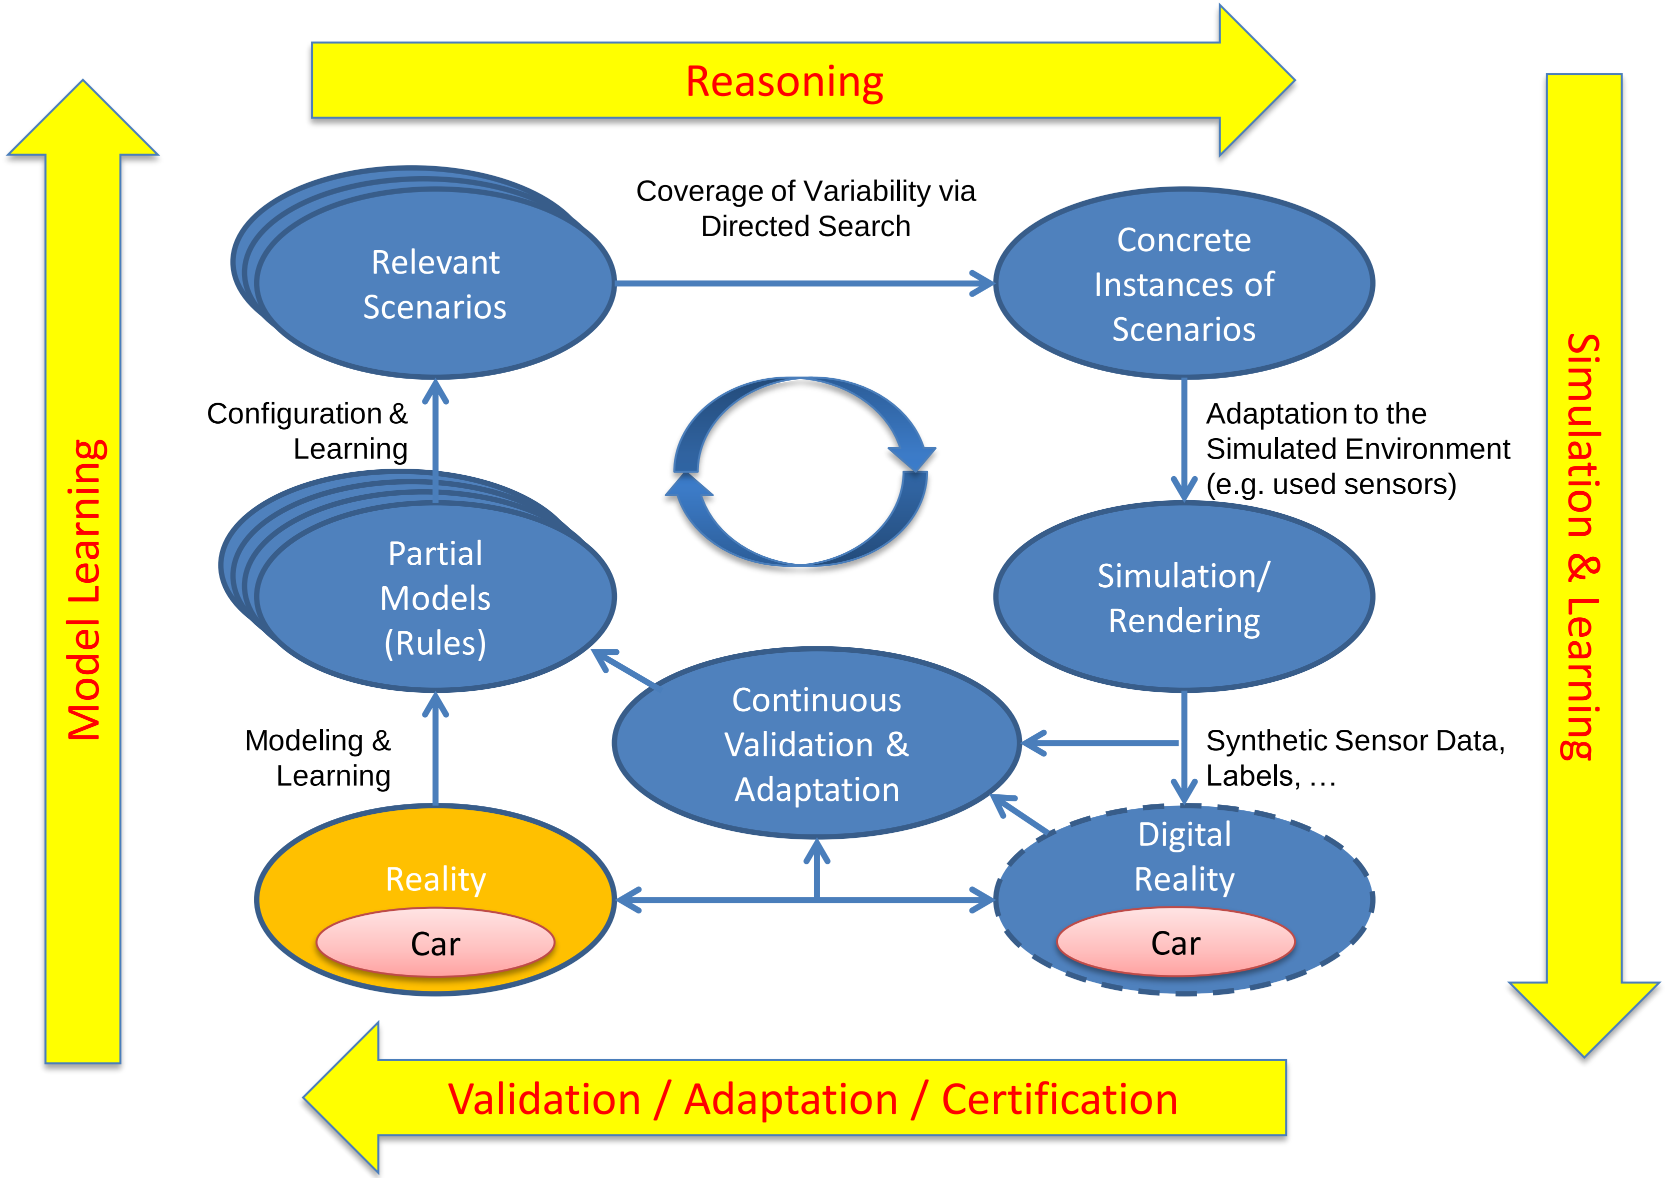

Digital Reality: AI to Certify AI

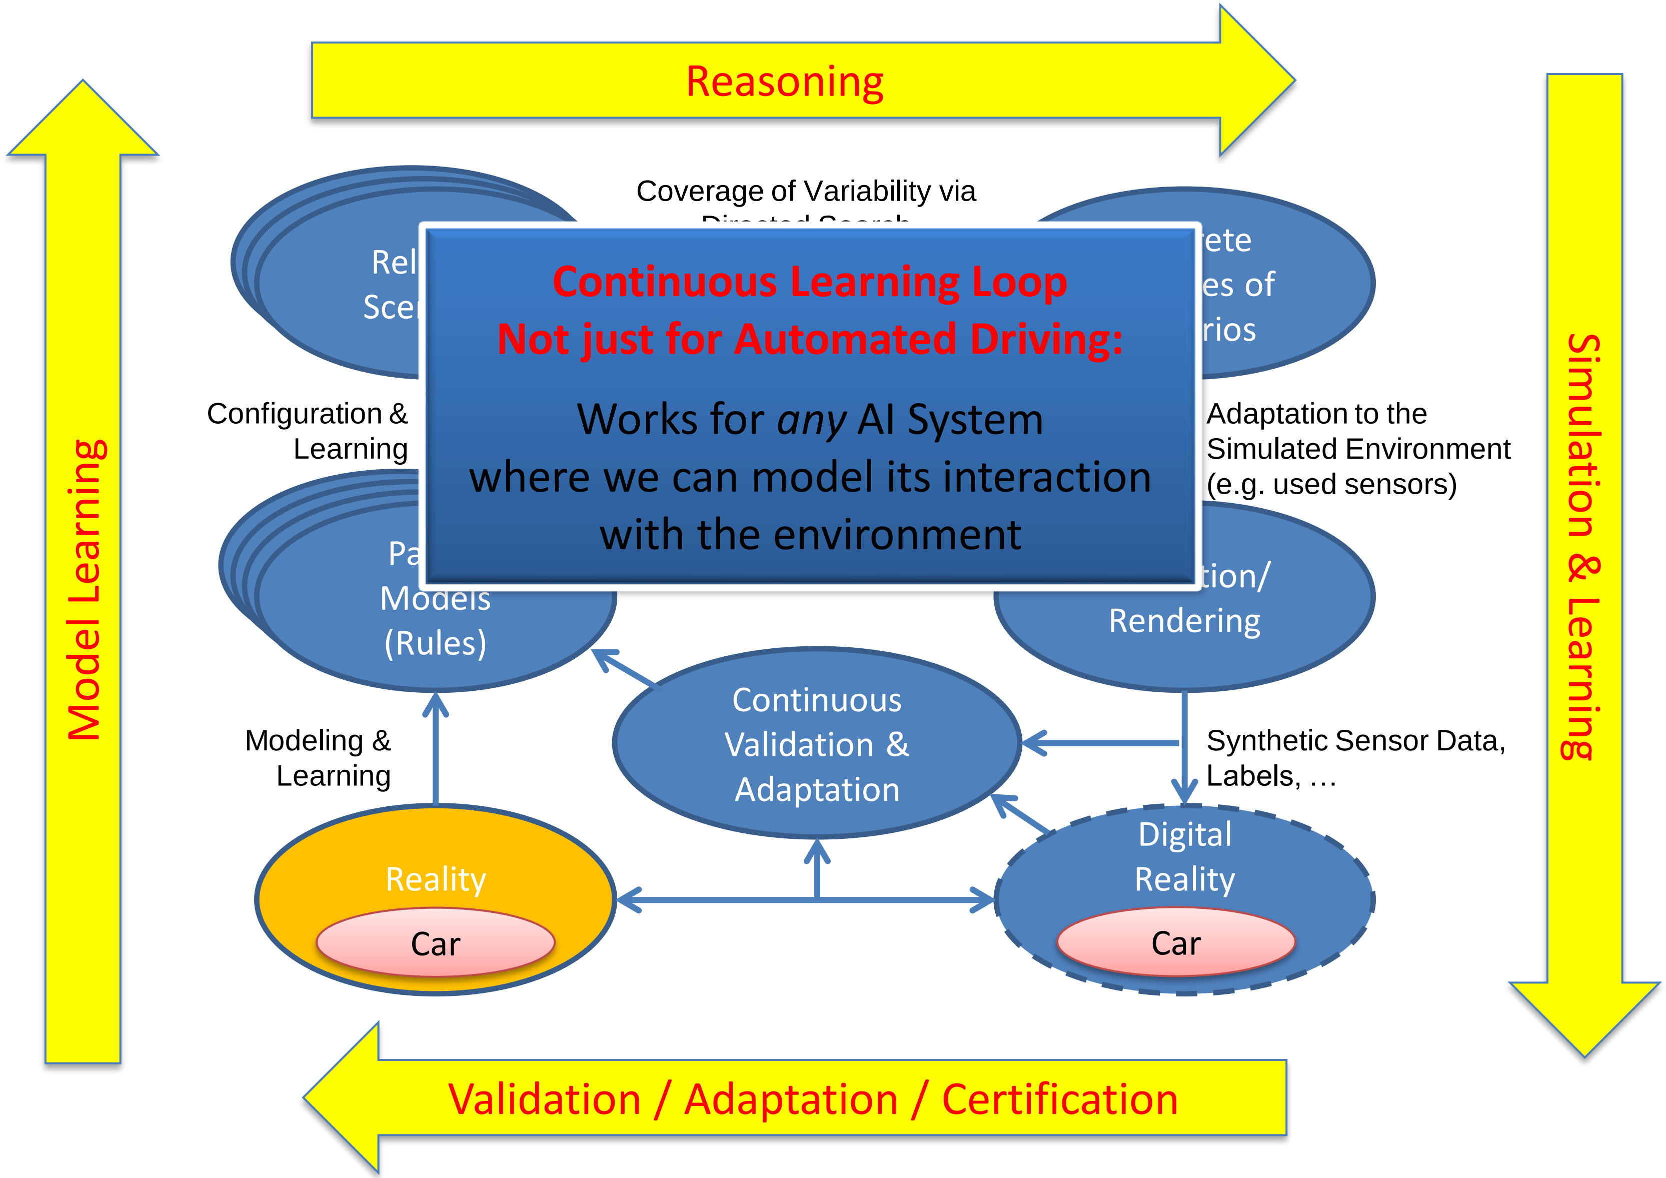

Syllabus

- Rendering Equation
- Finite Elements/Radiosity
- Perception, HDR Imaging, Tone Mapping
- Perception-based Rendering & Display Limitations
- Probability Theory & Monte-Carlo (MC) Integration
- BRDF & Path Tracing
- Density Estimation, Photon Mapping, Merge with MC
- Spatio-Temporal Sampling, Temporal Filtering
- Sampling & Reconstruction
- BiDir Tracing & MCMC
- Volume Techniques
- Interactive GI & HW-Support for Rendering and Lighting

Research From Saarbrücken

- Some examples

Reflection & Refraction

- Visualization of a car headlight
  - It reflects and refracts light almost entirely from the environment. Up to 50 rays per path are needed to render this image faithfully (800k triangles).

Instant Global Illumination

- Real-time simulation of indirect lighting ("many-light method")

Real-Time Photon Mapping

- Real-time performance with procedural textures and density estimation. Interleaved sampling allows to reduce computation by a factor of 10.

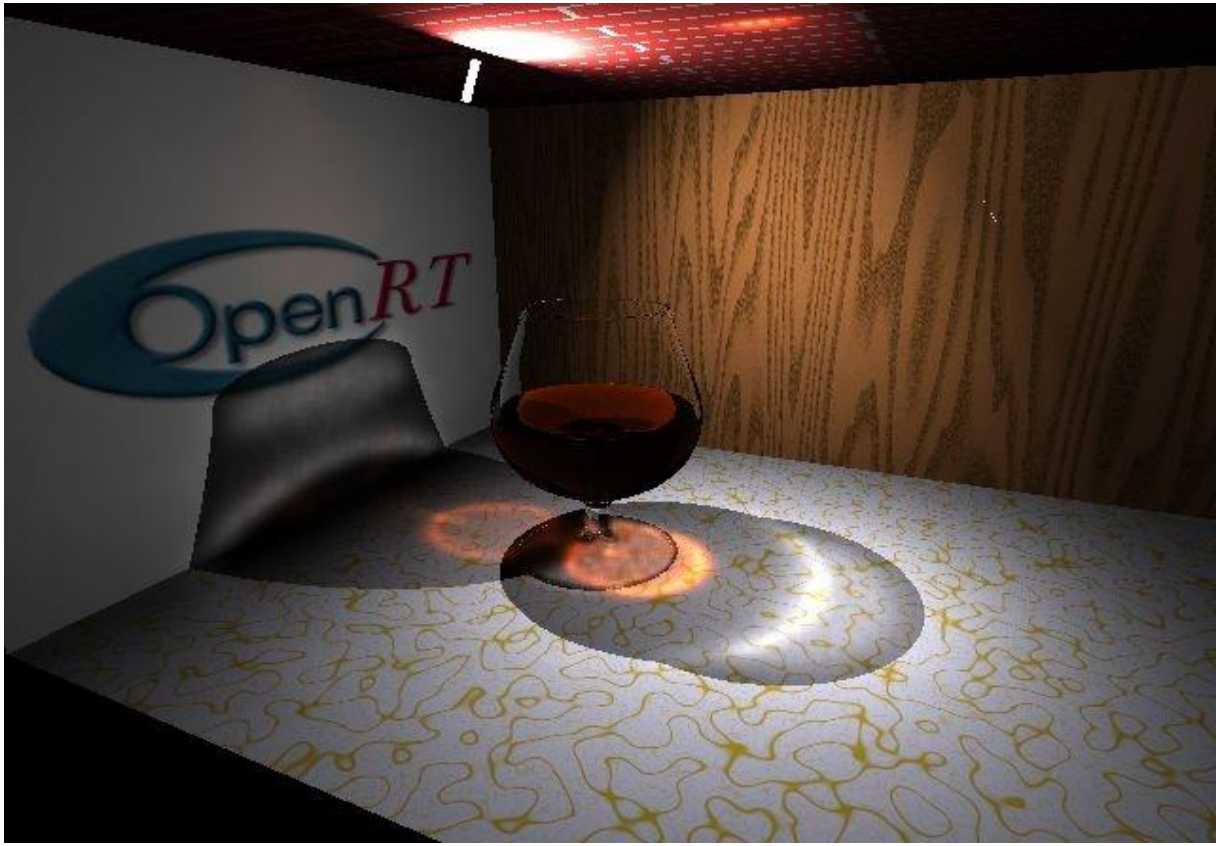

Photon Mapping

- Car headlight used as a light source
  - Photons are emitted and traced until they hit a wall. Density estimation is used to reconstruct the illumination. The results run at 3 FPS with 250k photons on a cluster of 25 cores (in 2004). Visualization without running the simulation achieves even 11 FPS (lower center) and compare well to a real photograph (lower right).

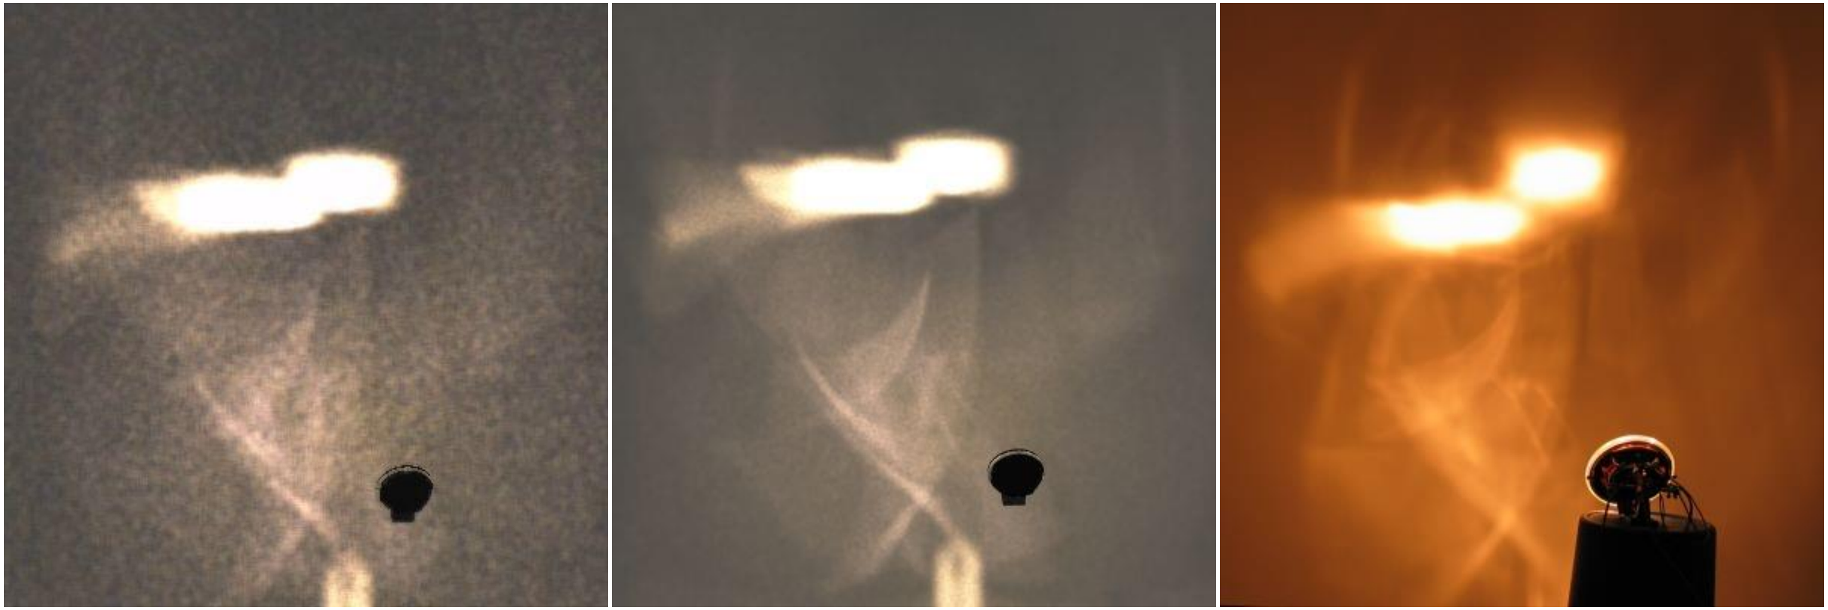

Advanced Materials

- Application to a real car using spline surfaces, realistic paint shaders, BTF shaders in the interior, and realistic environment lighting.

Advanced Materials

- The use of BTF for realistic materials with optical effects on the meso-scale (e.g. shadows in bumps and creases).

Light Transport Simulation

- Volkswagen’s large Corporate Visualization Center in Wolfsburg using using ray tracing technology developed in Saarbrücken (Spin-off “inTrace”).

Massive Models

- The original CAD model of a Boeing 777 consisting of 365 million polygons (30 GB). Ray tracing was the first method to allow real-time visualization of such models.

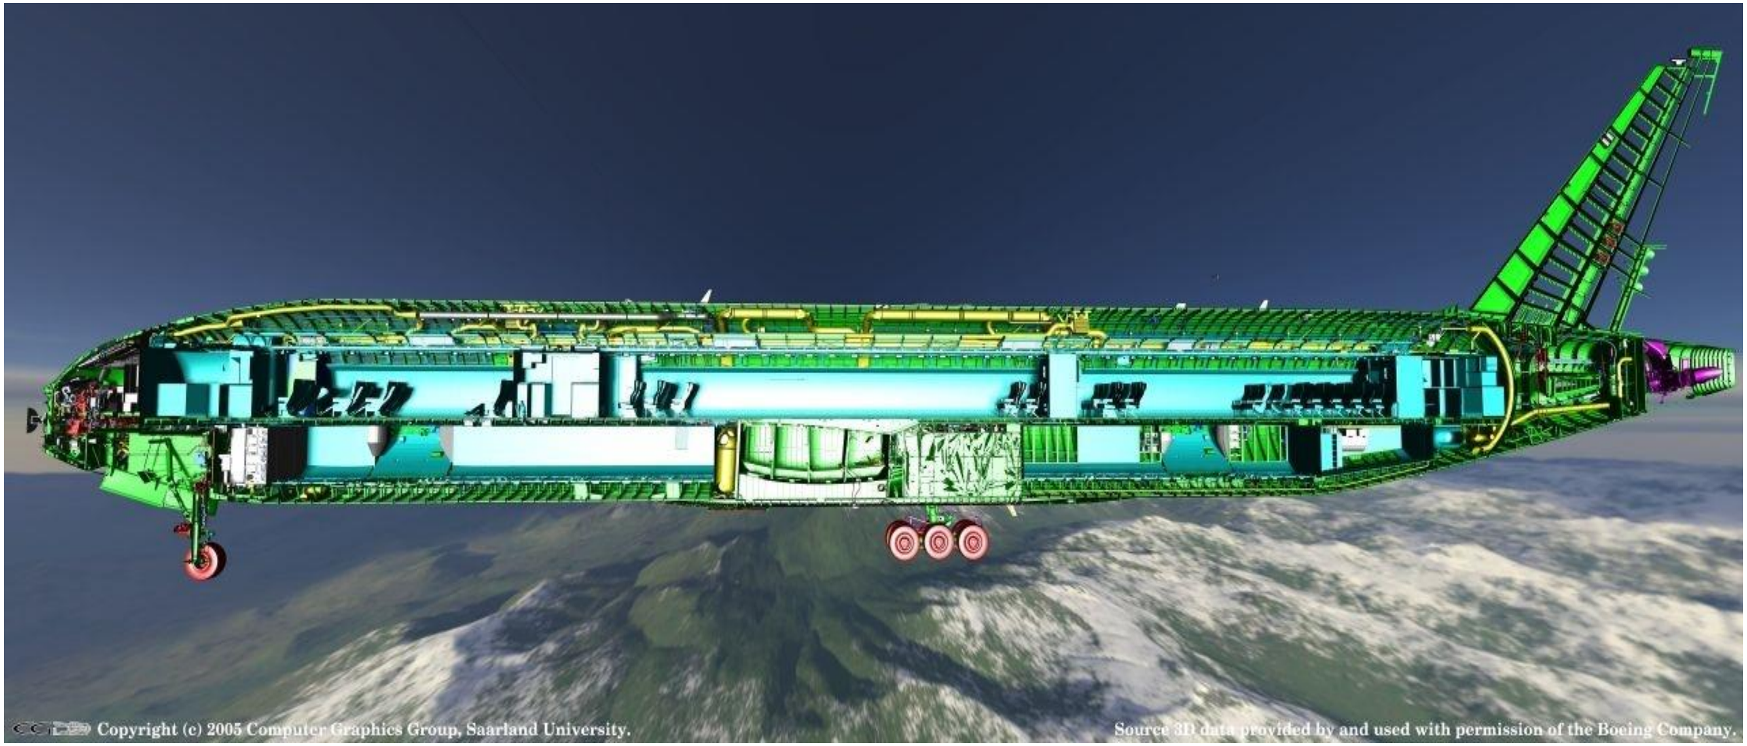

Massive Models

- Visualization of large outdoor scenes (300x300m^2) with 365k plants and several billion triangles.

Massive Models

- Much larger outdoor scene (80x80 km^2) with realistic lighting and full vegetation ($90*10^{12}$ triangles)

Volume Rendering

- Global illumination of iso-surfaces.

Multiple Iso-Surfaces

- Ray tracing allows easy integration of multiple modalities into a single rendering framework.

High-Performance Simulation

- Advanced rendering techniques in games

Importance Caching

- Iliyan Georgiev, et al. [Eurographics 2012]
  - Reuse samples based on probability

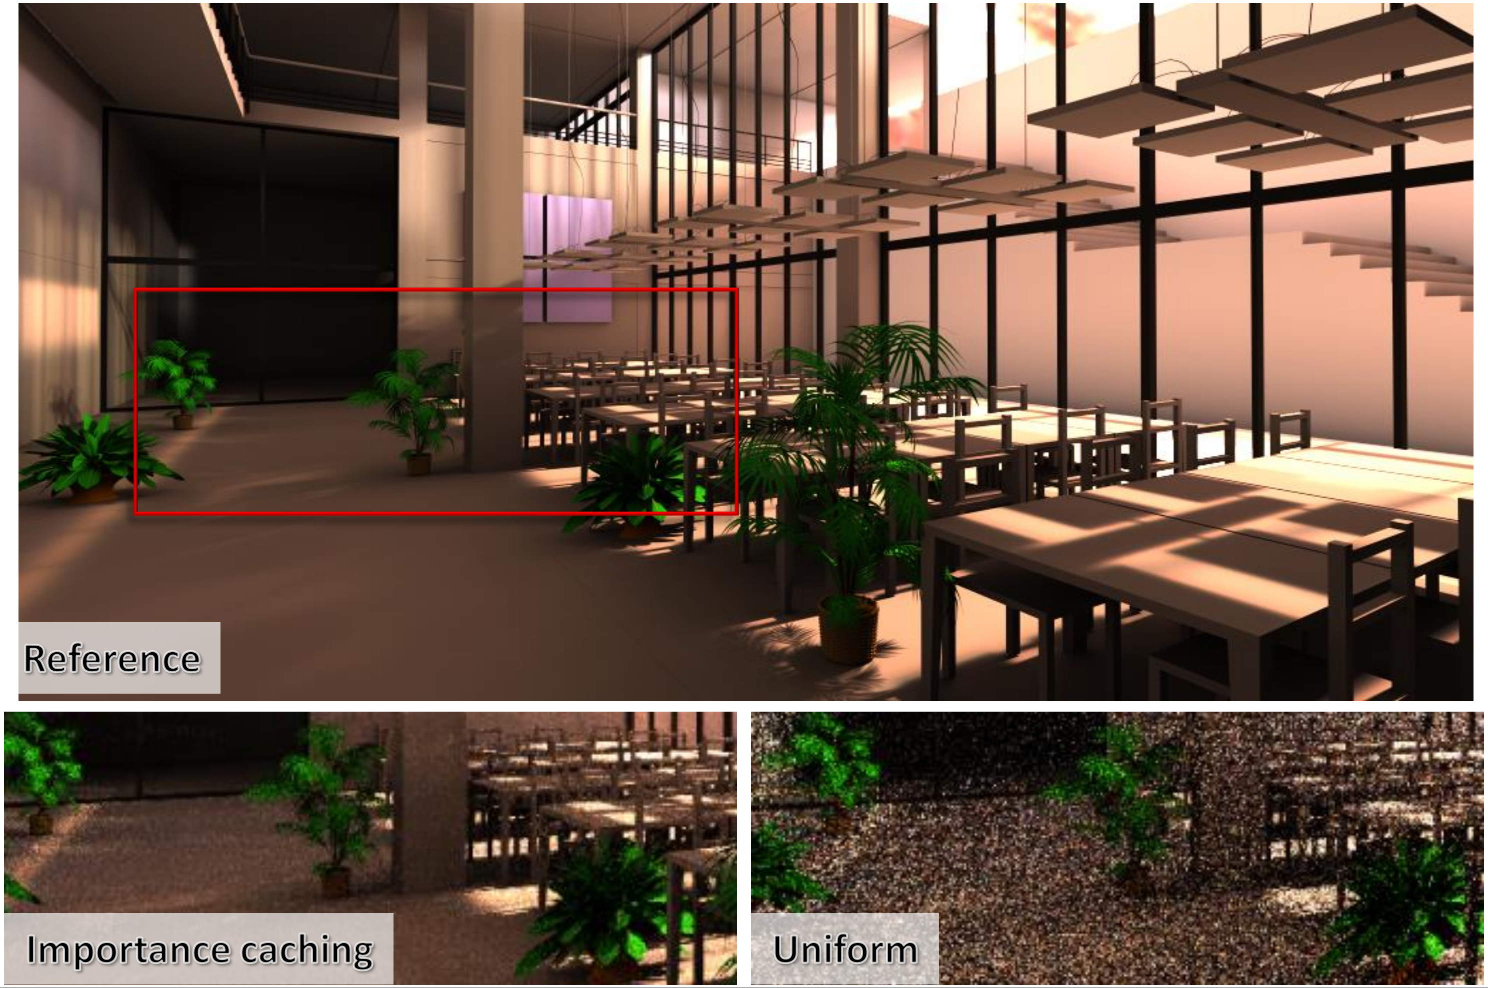

Monte-Carlo vs Density Estimation

- Vertex Connection & Merging, Ilijan Georgiev [SiggraphAsia´12]
  - Formulating Density Estimation algorithms as a Monte-Carlo (MC) techniques
  - Allows for direct combination with other MC techniques via MIS

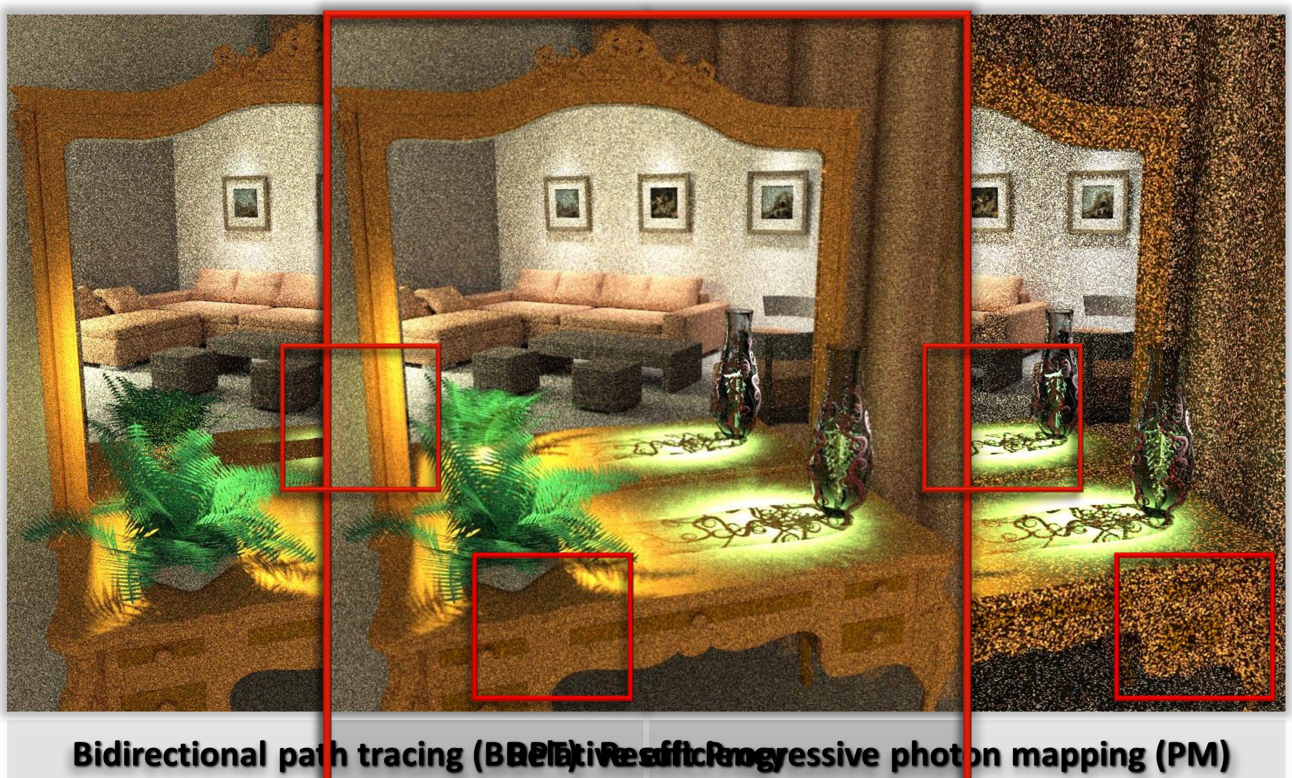

Same time (1 minute)

Monte-Carlo vs Density Estimation

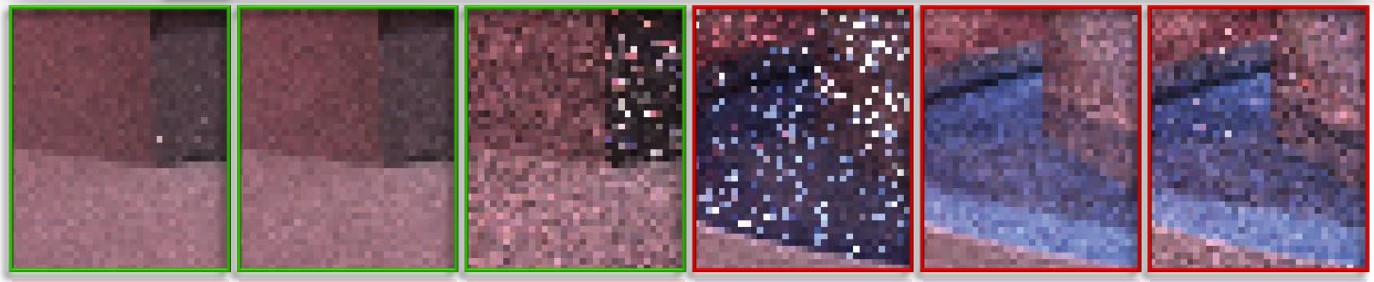

BDPT	BDPT+VM	PPM	BDPT	BDPT+VM	PPM

Same time (3 minutes)

Order of Convergence

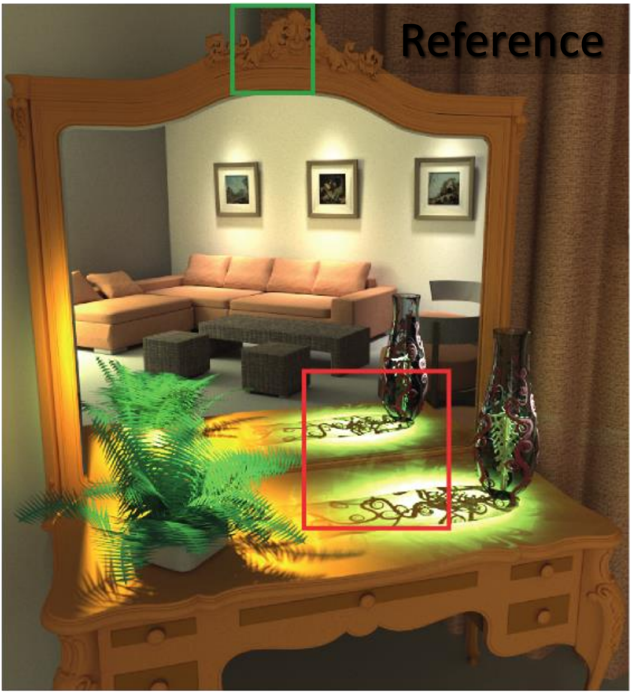

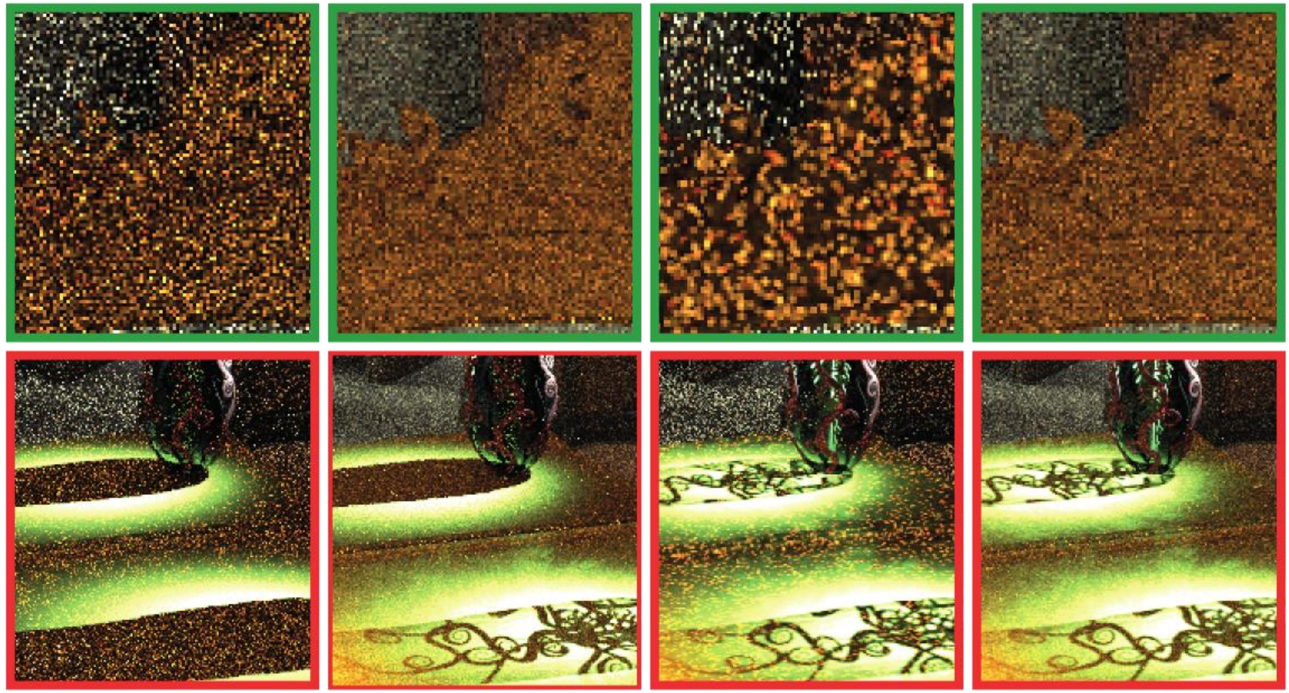

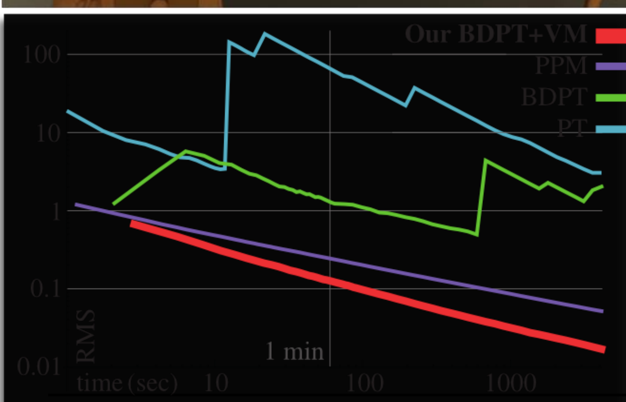

PT	BDPT	PPM	BDPT+VM
$O(N^{-0.5})$	$O(N^{-0.5})$	$O(N^{-0.33})$	$O(N^{-0.5})$

Same time (1 minute)

Joint Path Sampling

- Iliyan Georgiev, et al. [SiggraphAsia 2013]
  - Joint sampling of set of next events

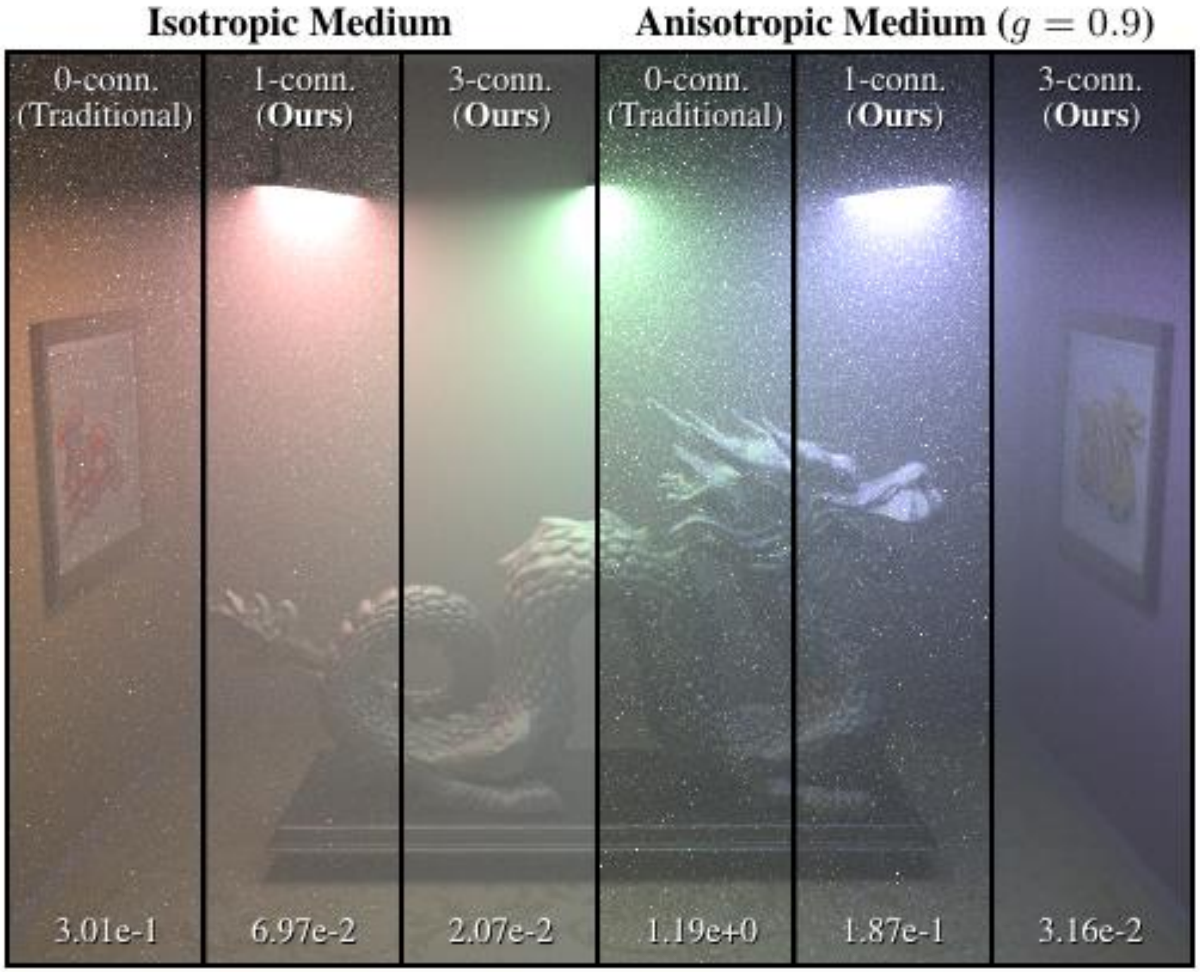

Emission Guiding

- Pascal Grittmann [EGSR´18]

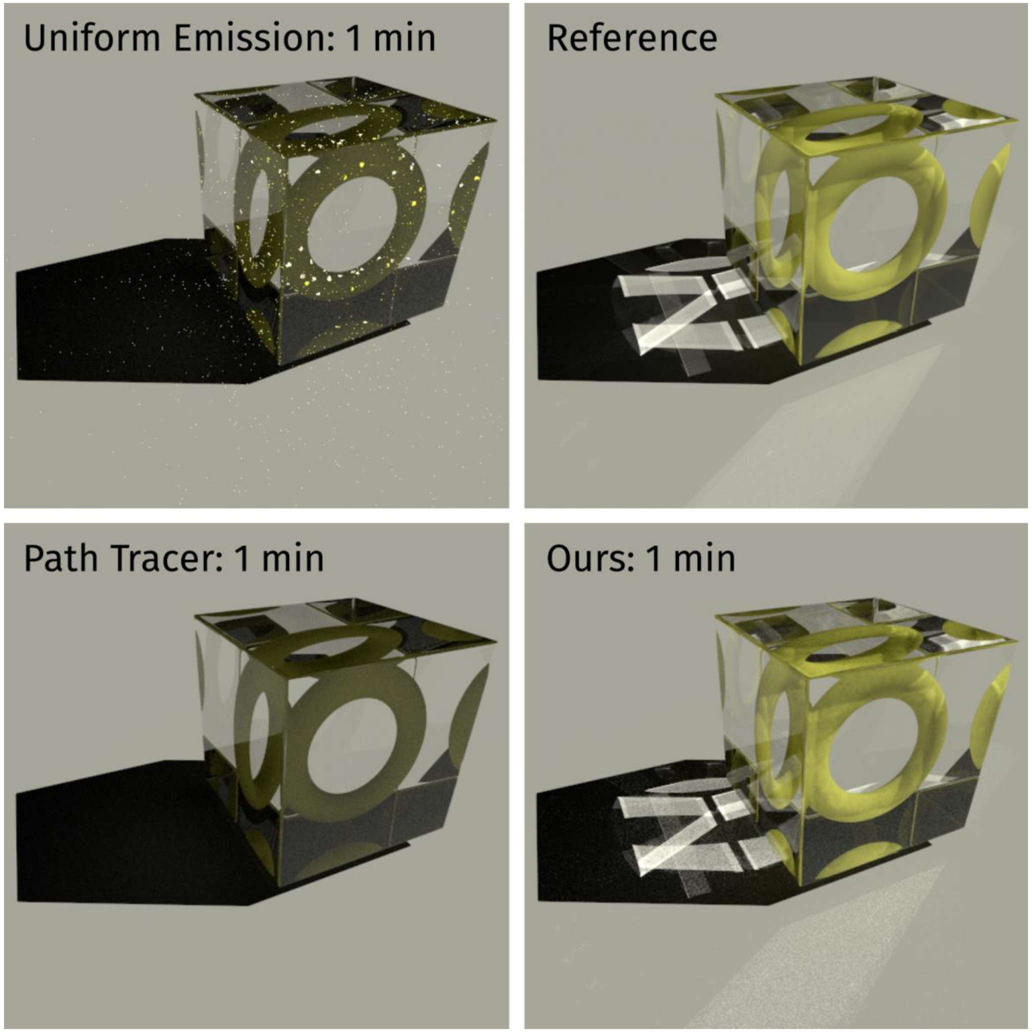

Emission Guiding

- Using Photon Mapping only where it is useful

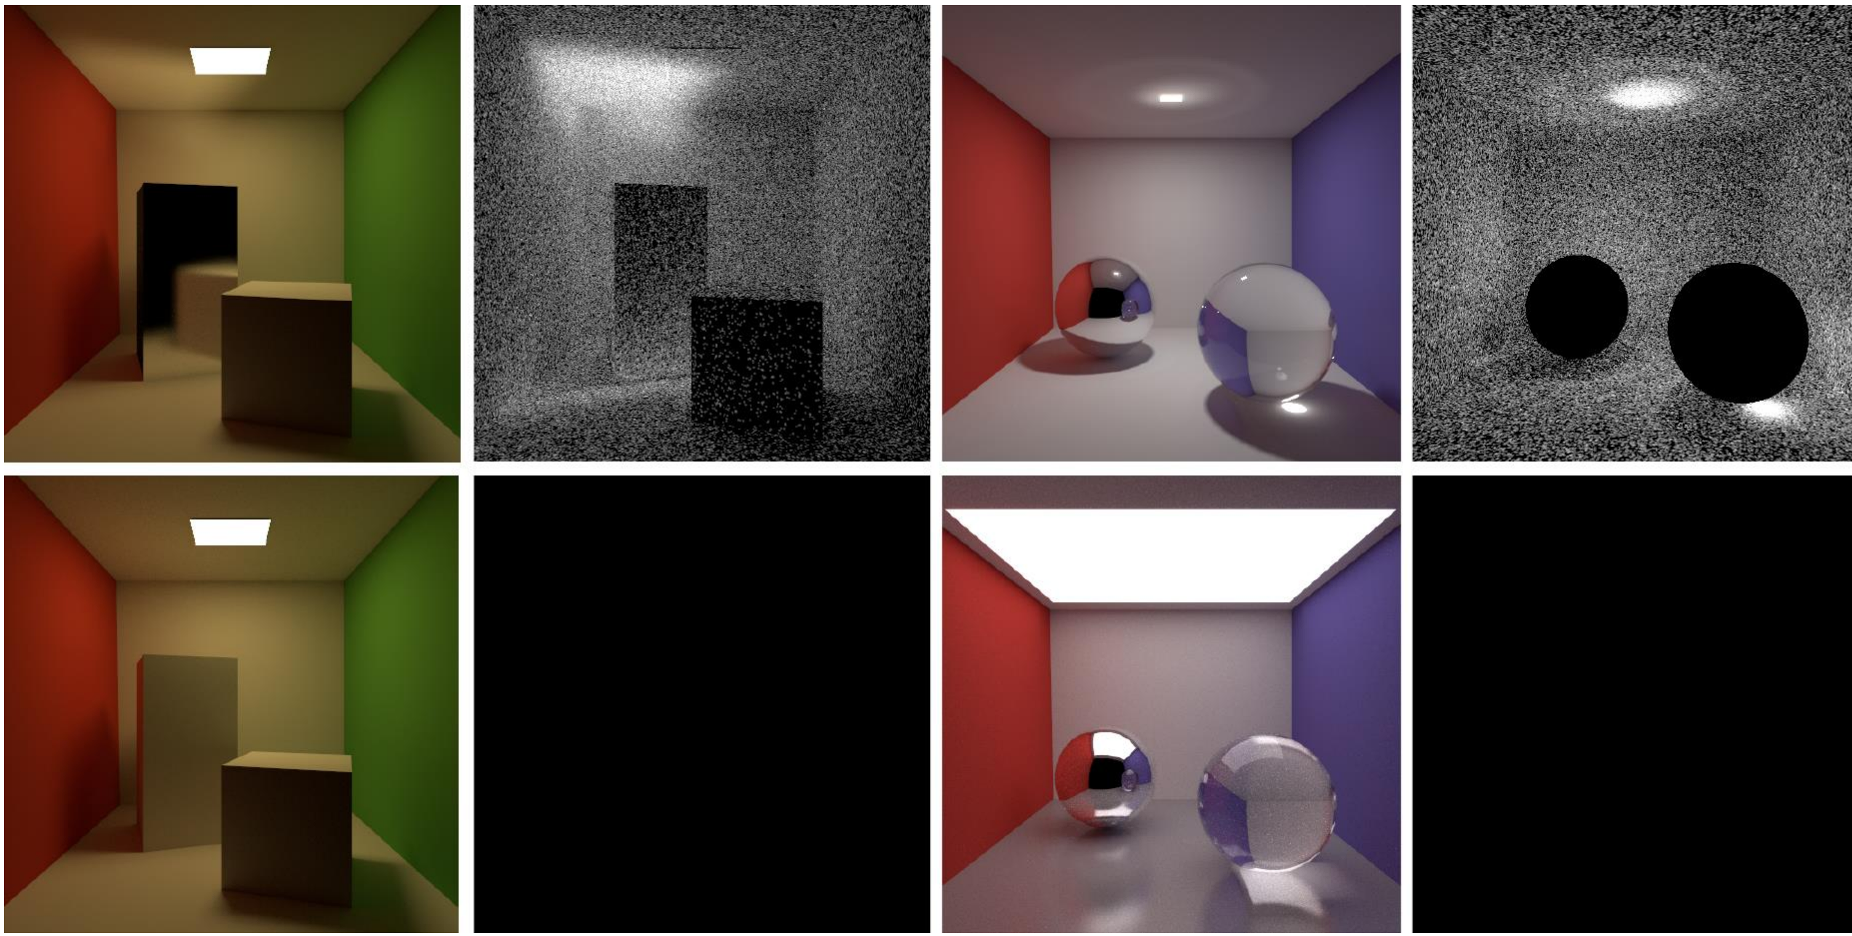

Optimal MIS

- Pascal Grittmann, et al. [Siggraph´19]
  - Multiple Importance Sampling (MIS) should optimally combine multiple estimators (i.e. sampling strategies) via suitable weights
  - Unfortunately, original technique made too specific assumptions
    - Finally fixed (24 years later!!) – but quite costly

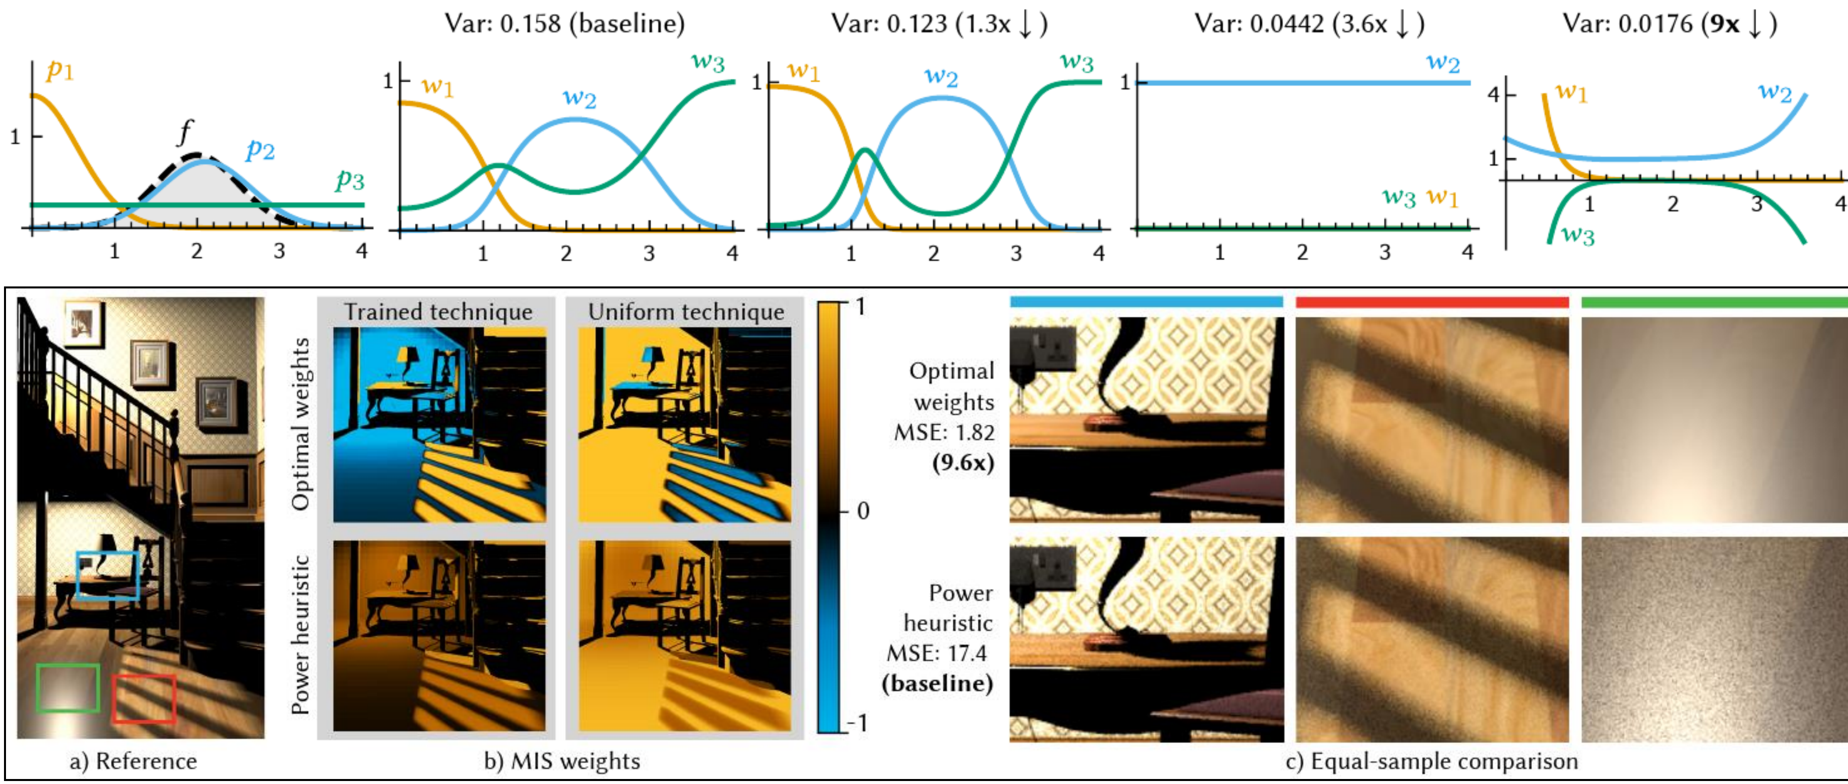

Variance-Aware MIS

- Pascal Grittmann et al. [Siggraph Asia´20]
  - MIS should provide better estimator than individual estimators
    - This is not always true :-(
    - E.g. the effects of stratification are not taken into account
  - Solved by injecting variance estimates for each individual technique
  - Essentially cost-free !!!

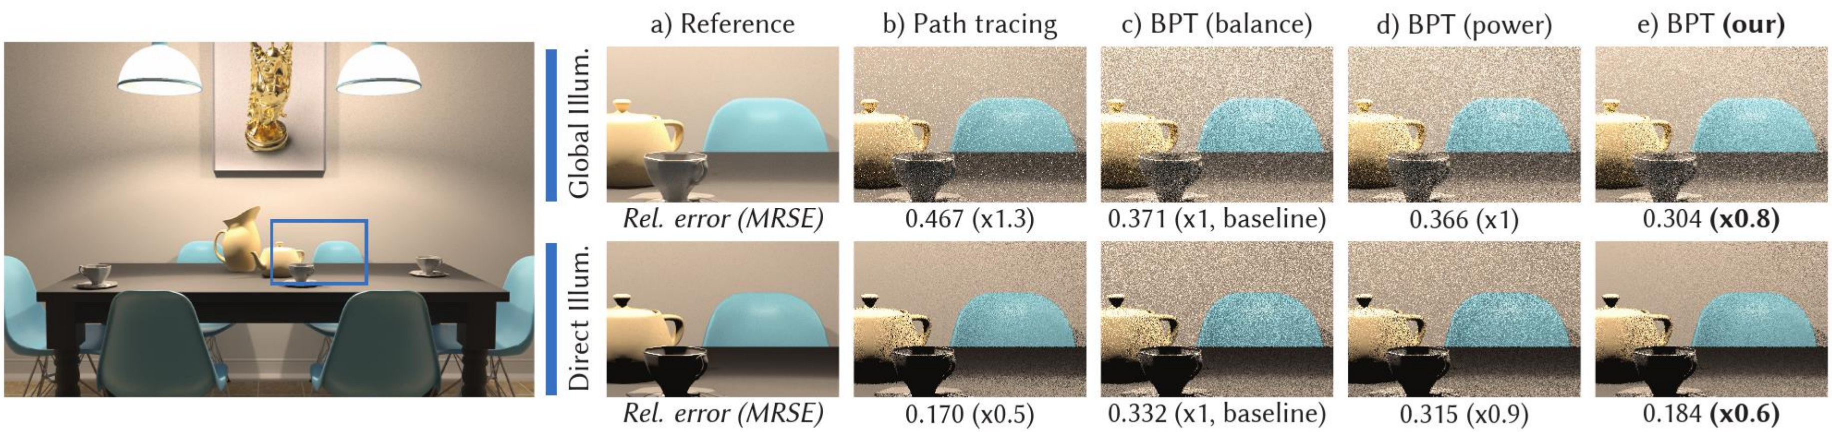

Optimal Target Densities for Guiding

- Alexander Rath, PascalGrittmann, et al. [Siggraph´20]
  - Need better estimate where to trace photons to
  - Assume that decisions are not perfect and take BRDF into account
  - Derive theoretically optimal target densities for local path guiding

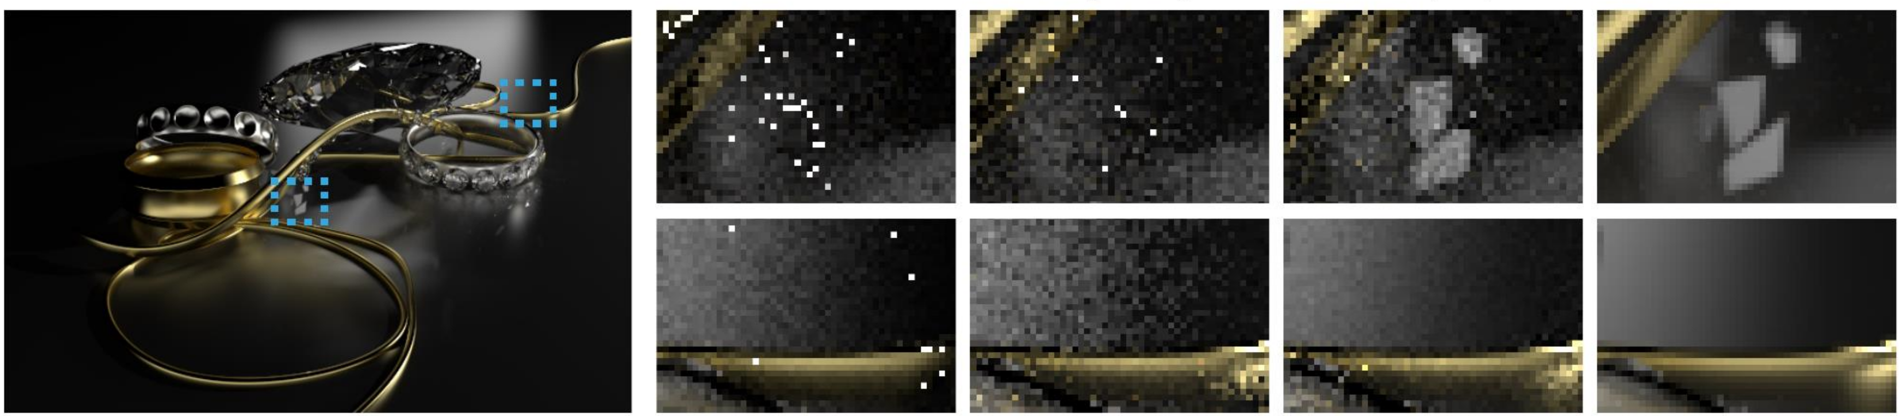

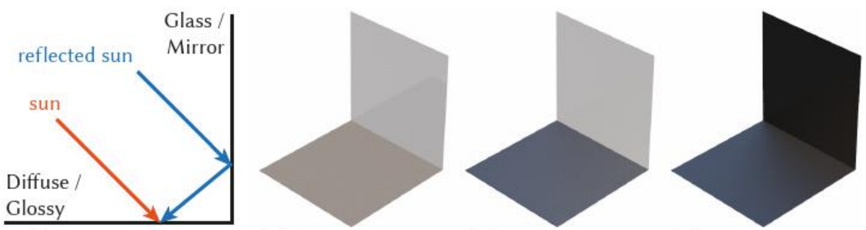

(a) Cross-section (b) Diffuse / Glass (c) Glossy / Glass (d) Glossy / Mirror

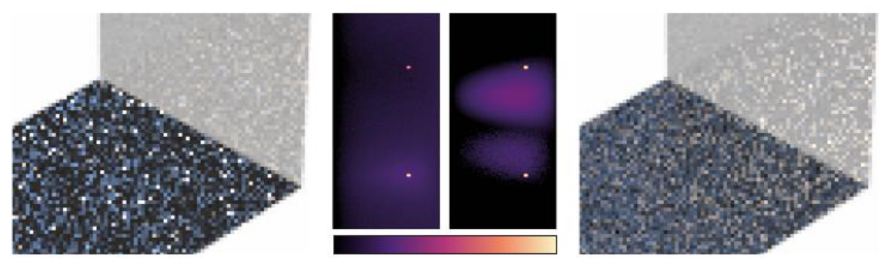
(a) Radiance-based (b) Densities (c) Our target function

Ultimate Goal

- Reality check
  - Can we render real-time video of such scenes ?

Lecture Related Research at MPII

- High Dynamic Range Imaging
  - Tone mapping, image quality metrics
- Apparent display quality improvement
  - Cornsweet effect, glare simulation, resolution enhancement
- Rendering for modern displays
  - Multi-focal plane displays, deformable beam splitters
  - Light-fields, focal stacks
- Foveated rendering
  - Latency reduction - saccade landing point prediction
  - Image-content aware foveation
- Advanced rendering
  - Intelligent sample point selection
  - Machine learning solutions for denoising and rendering

High Dynamic Range (HDR) Imaging

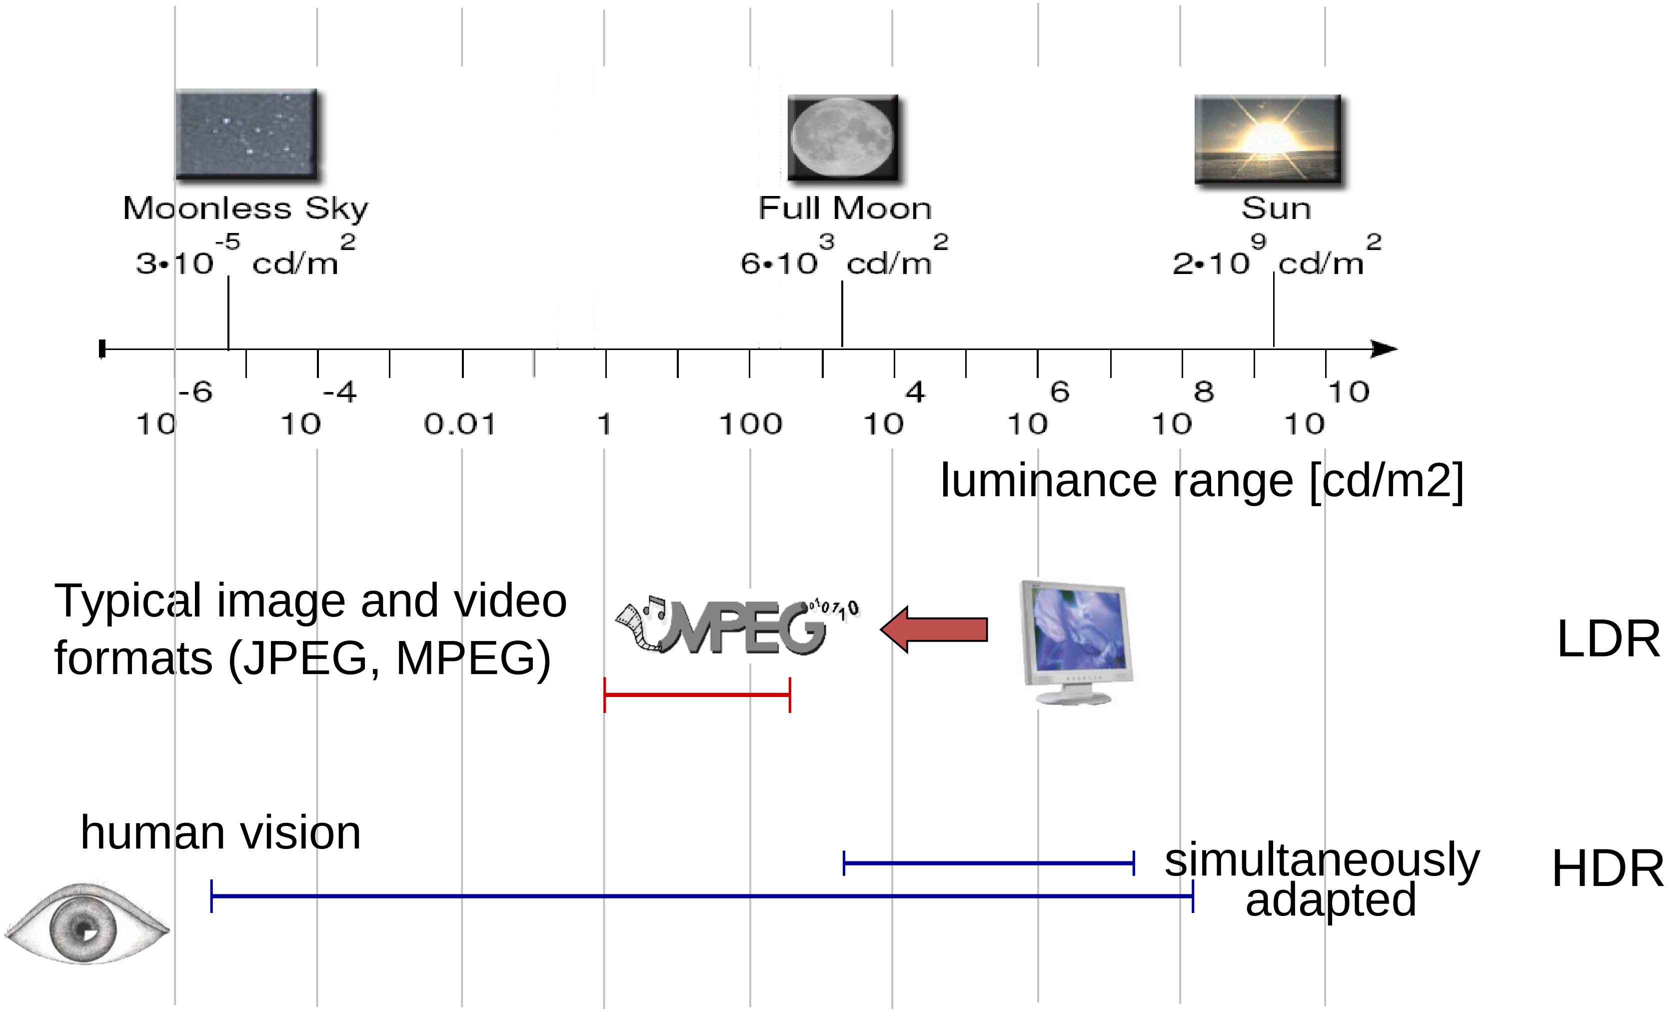

HDR Imaging Pipeline

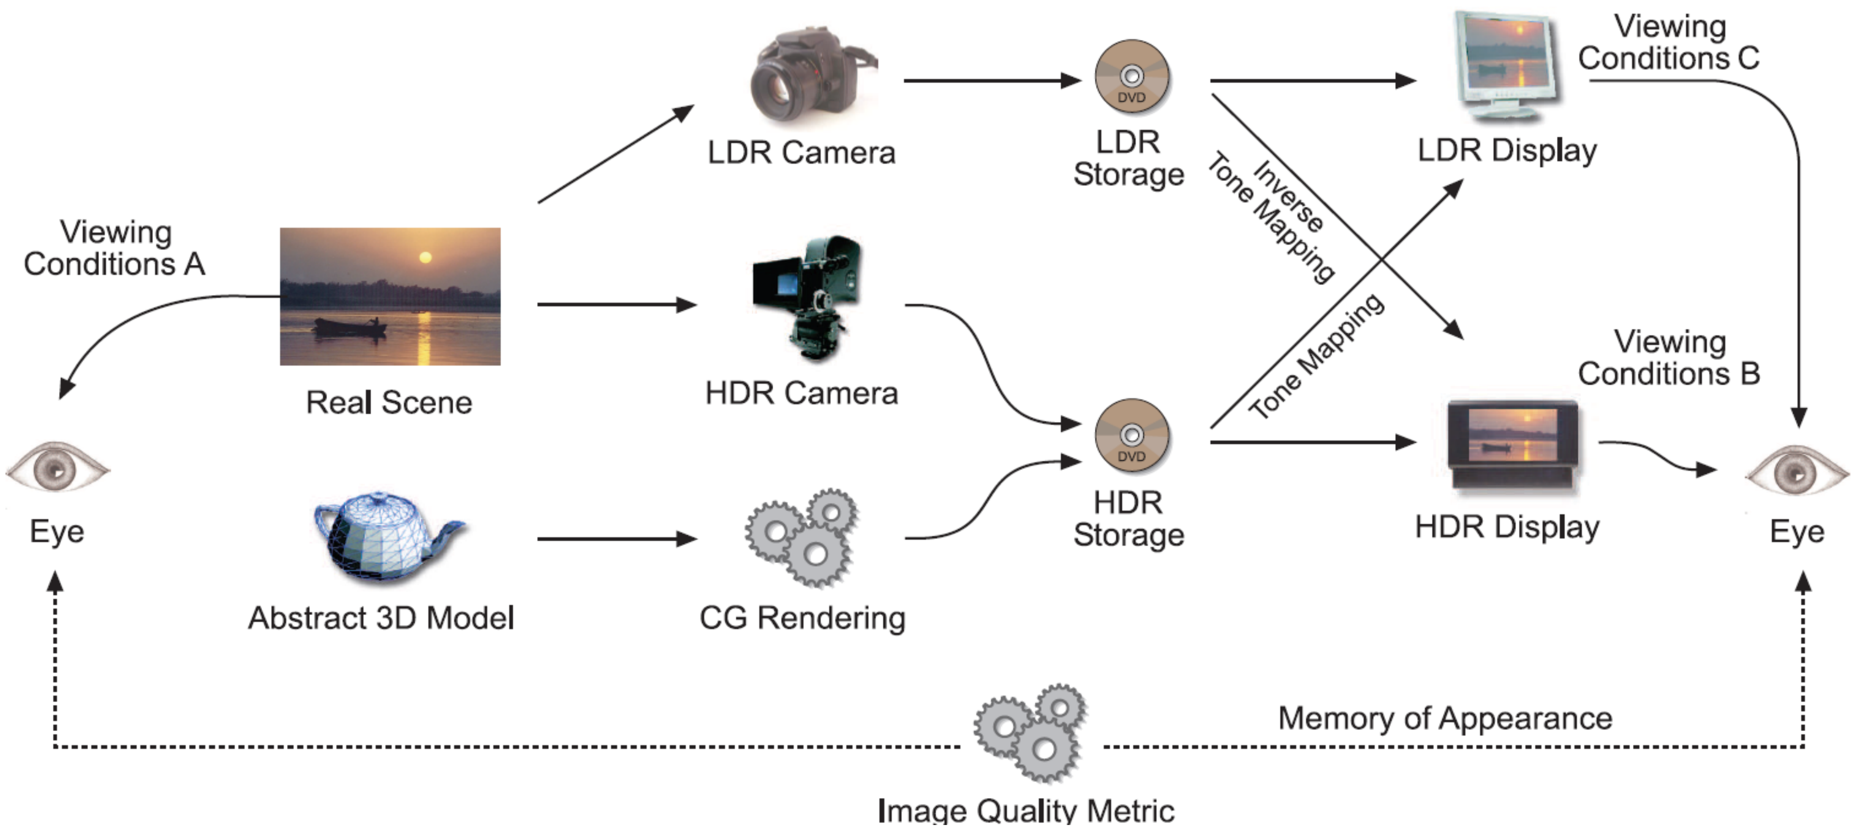

Tone Mapping

Naïve
non-linear compression

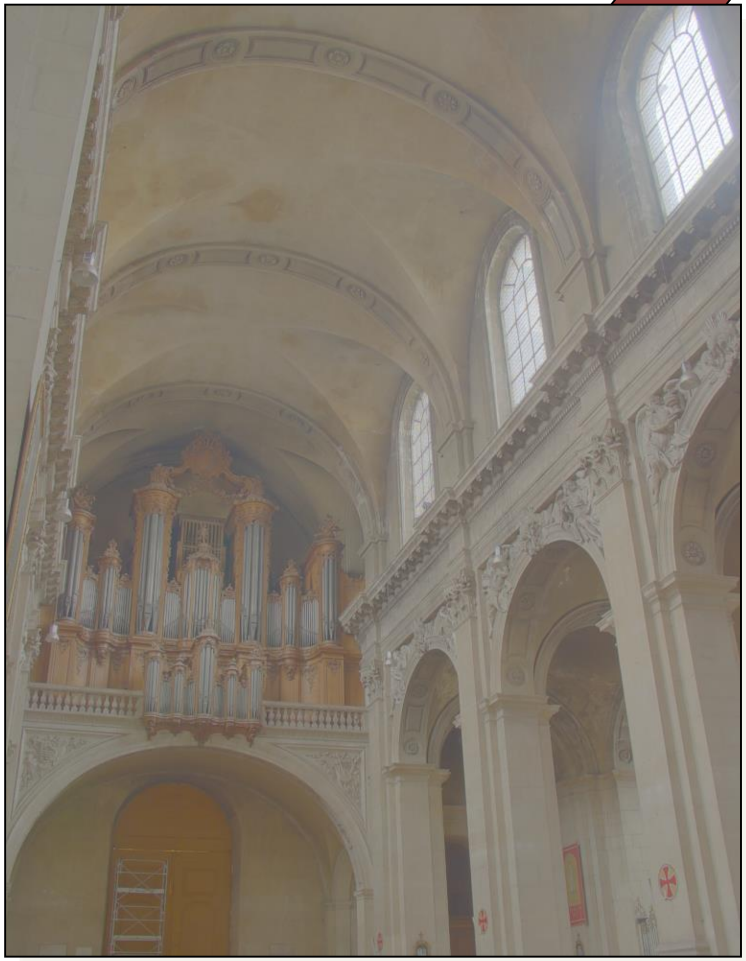

Advanced
contrast perception model

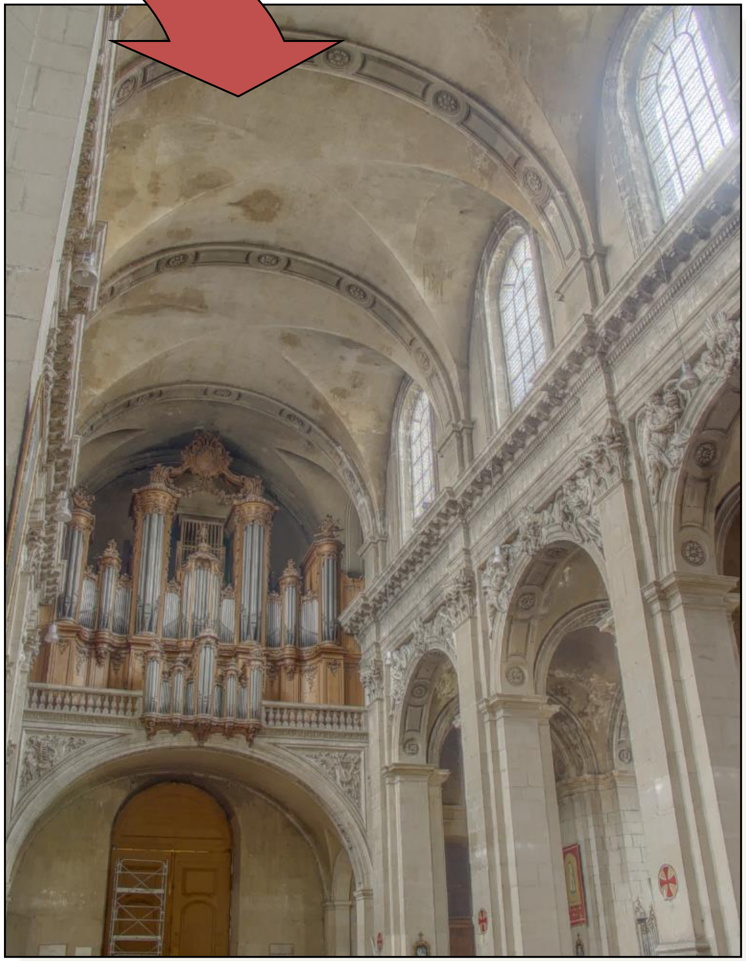

MPI HDR Software

PFSTools
For High Dynamic Range Images and Video

pfsin → pfstmo → pfsout

http://pfstools.sourceforge.net/

PFStmo

tone mapping operators

http://www.mpii.mpg.de/resources/tmo/

PFScalibration
Photometric Calibration of HDR and LDR Cameras

http://www.mpii.mpg.de/resources/hdr/calibration/pfs.html

HDR Visual Difference Predictor

http://www.mpi-sb.mpg.de/resources/hdr/vdp/index.html

GPL License

Overcoming Display Limitations

- Enhancing apparent (perceived) quality rather than improving technical aspects
- Take advantage of the visual system properties

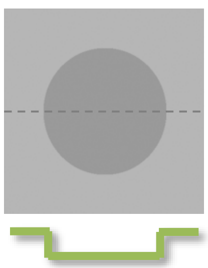

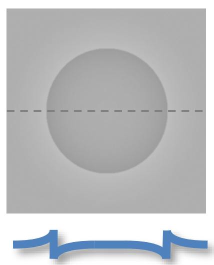

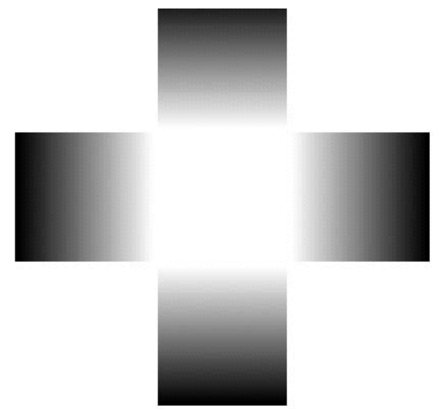

sharpen

blur

blur

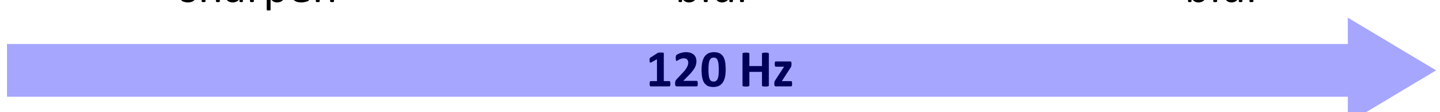

Cornsweet Illusion

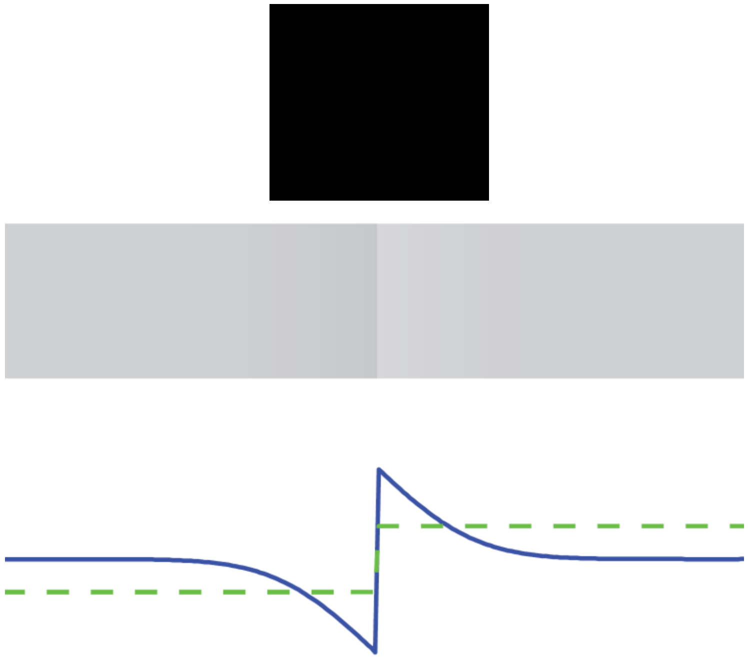

Glowing Effect

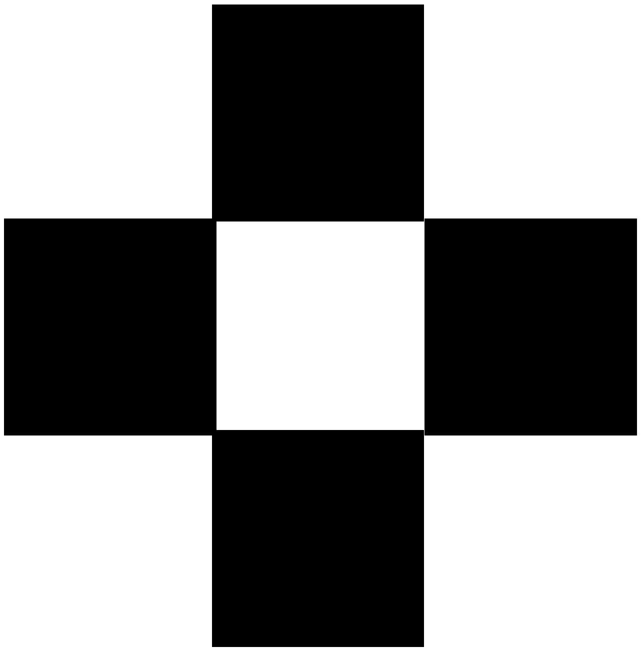

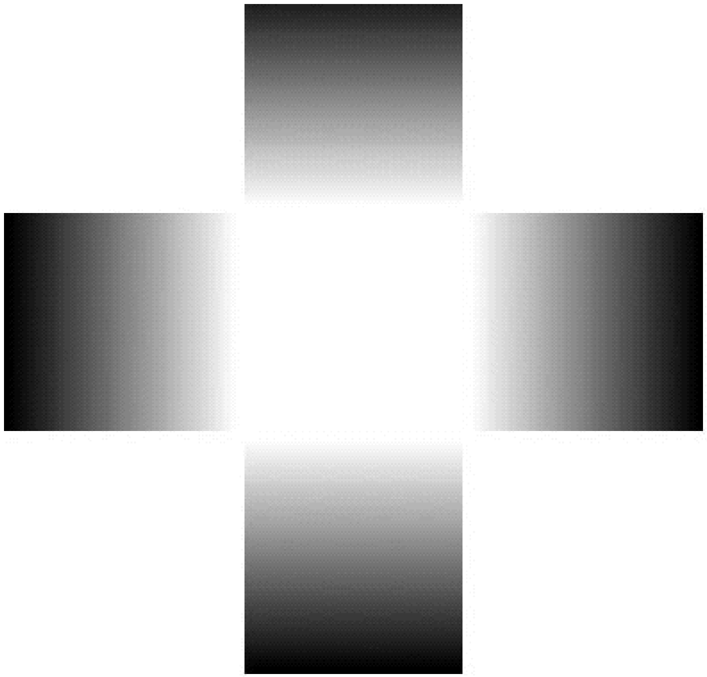

[Zavagno and Caputo 2001]

Apparent Resolution Enhancement

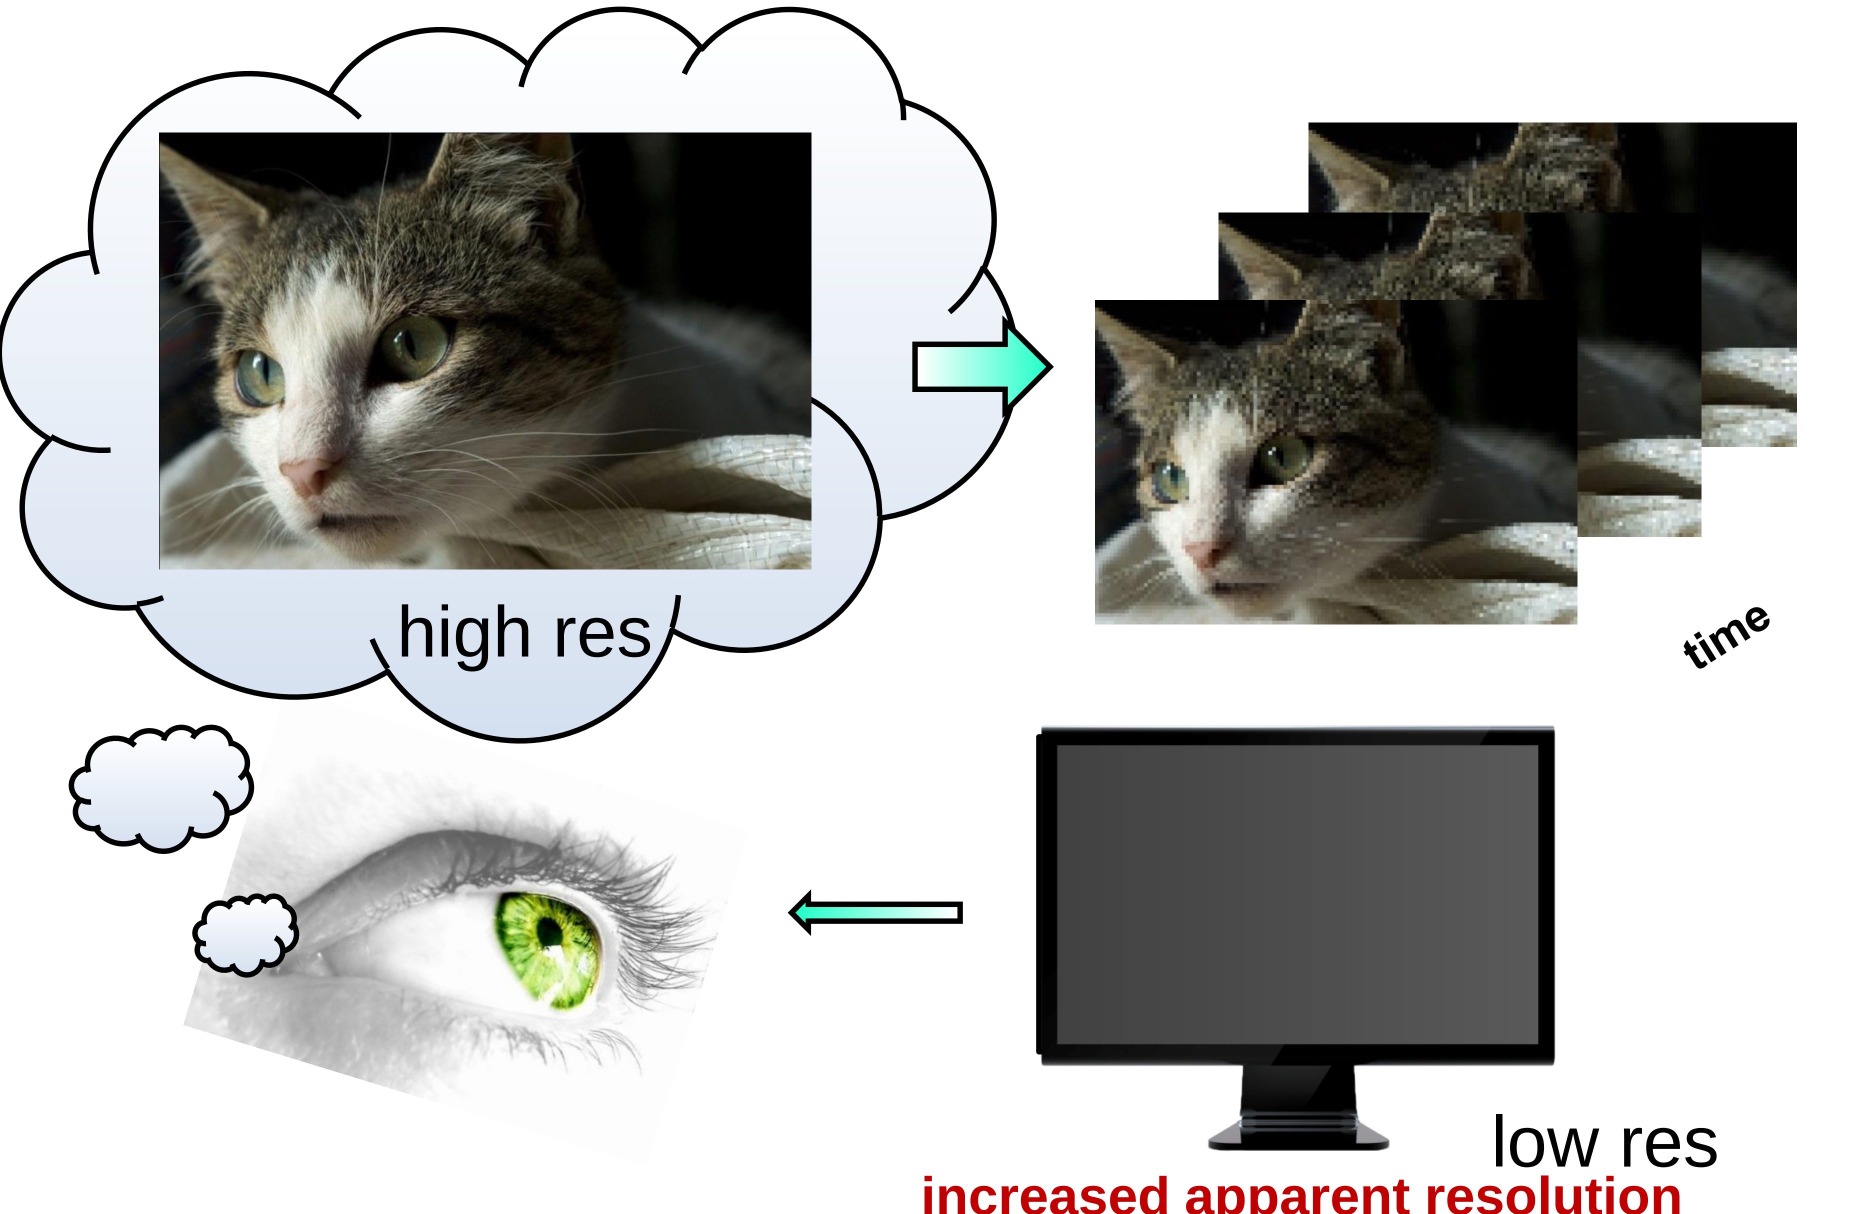

Optimization Result

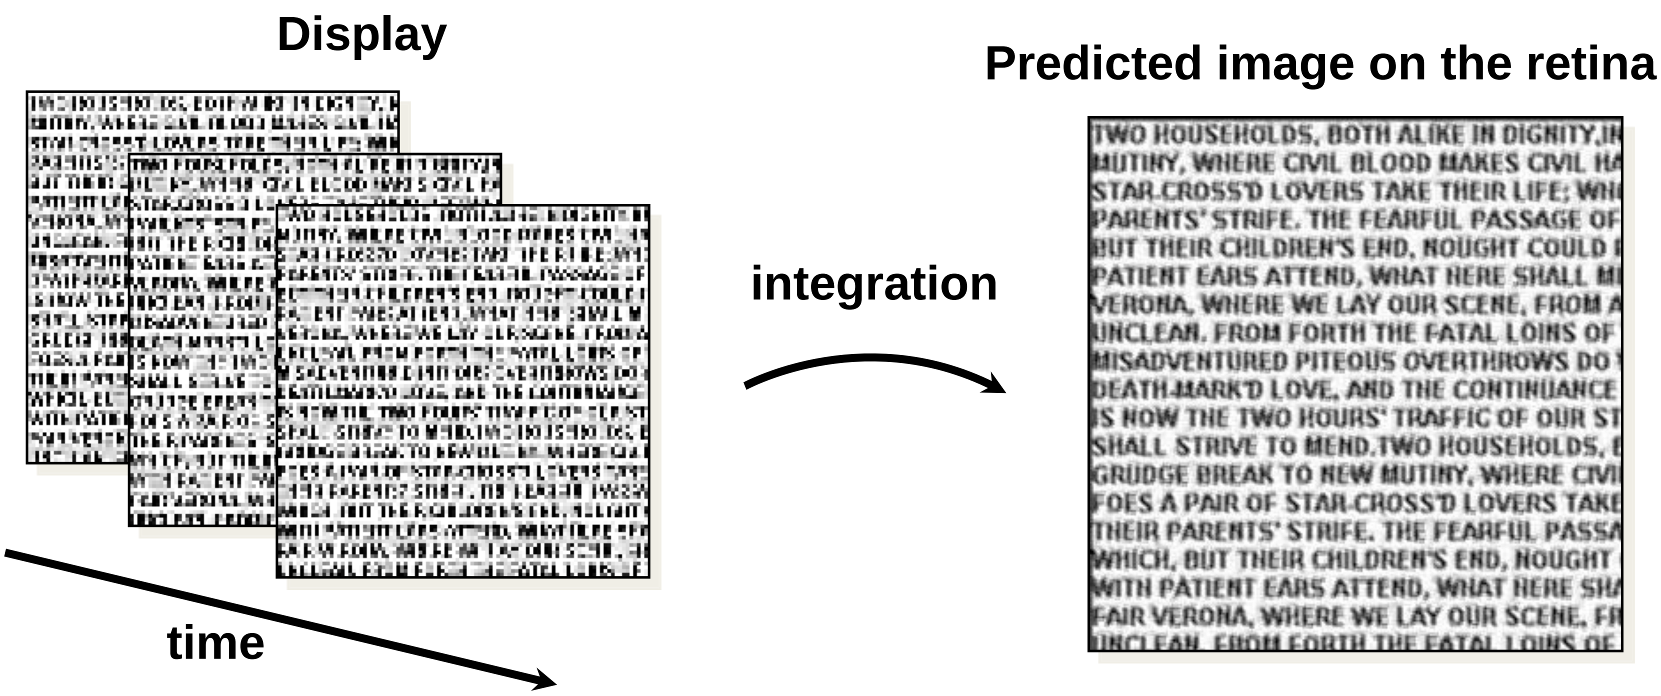

3D Image Retargeting

Input devices

Produce different depth ranges

Output devices

Reproduce different depth ranges

Visible Difference Metric (VDP)

- Can the human eye see the differences between two images?

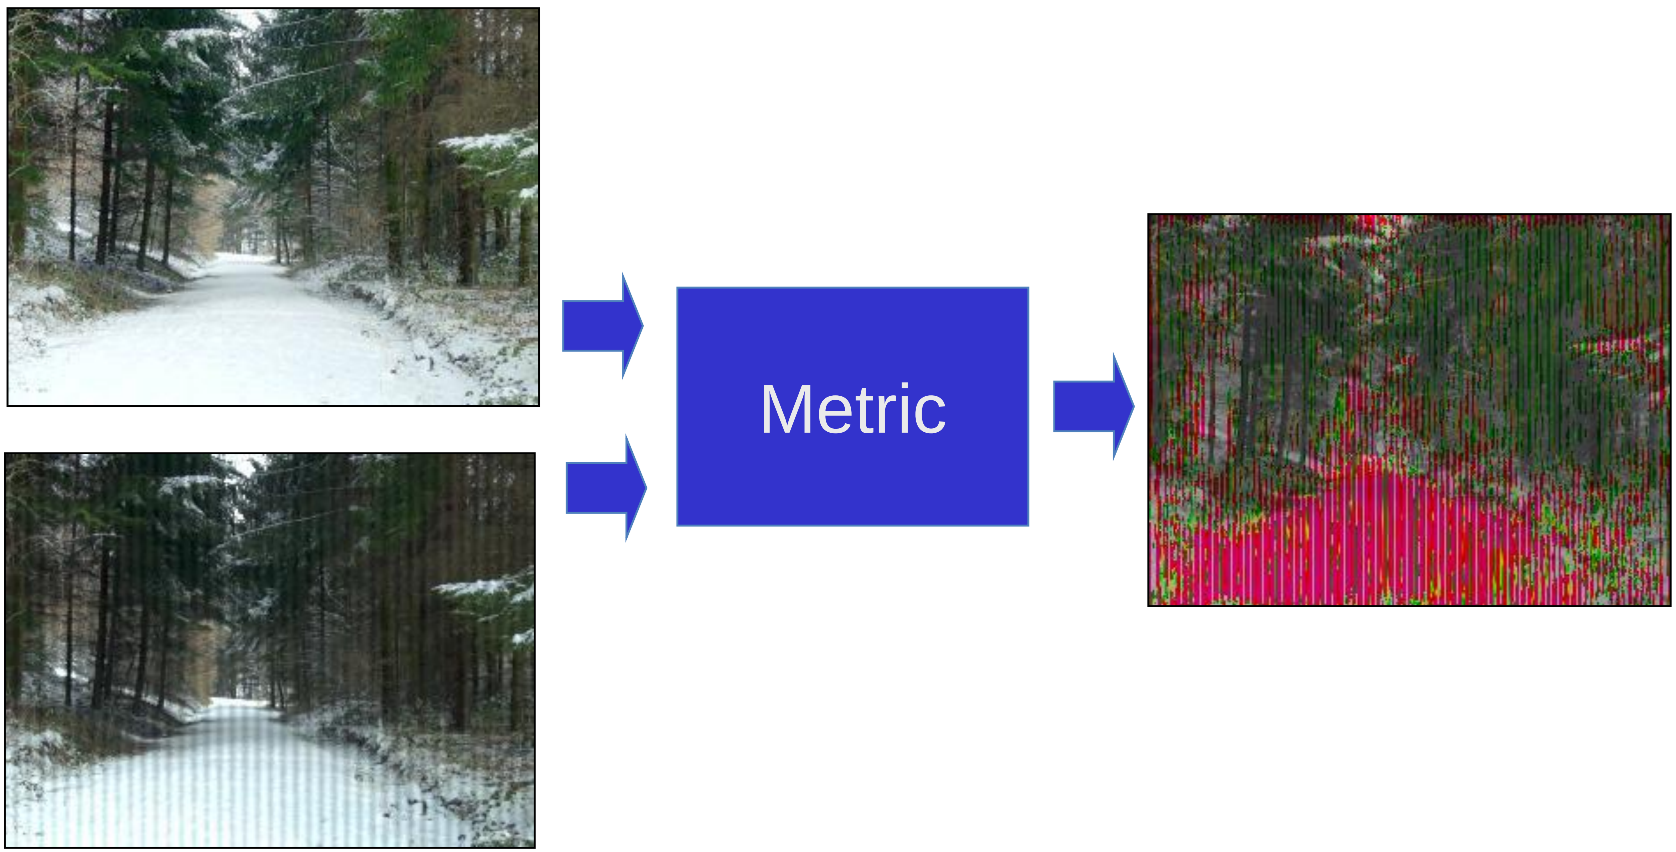

Dataset of Visible Distortions

Peter Panning

Z-fighting

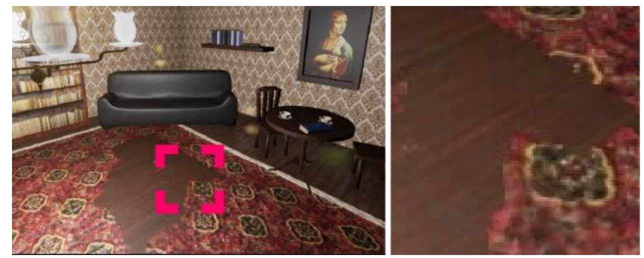

Shadow acne

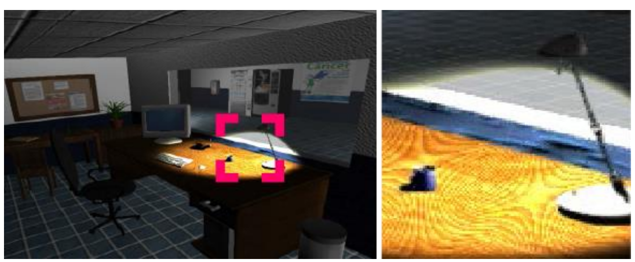

Shadowmap downsampling

[Piórkowski et al. 2017]

Dataset of Visible Distortions

Aliasing

[Piórkowski et al. 2017]

IBR

[Adhikarla et al. 2017]

Perception patterns

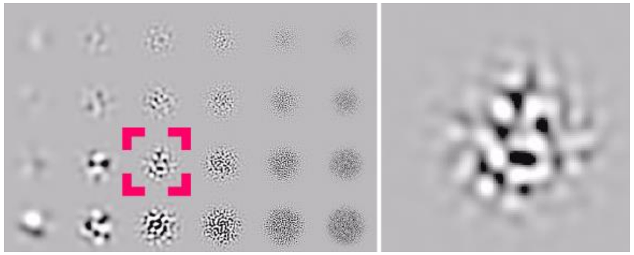

[Čadík et al. 2013]

Deghosting

[Karađuzović-Hadžiabdić et al. 2017]

Dataset of Visible Distortions

145 SCENES

1-3 LEVELS

296! IMAGES

Label Creation

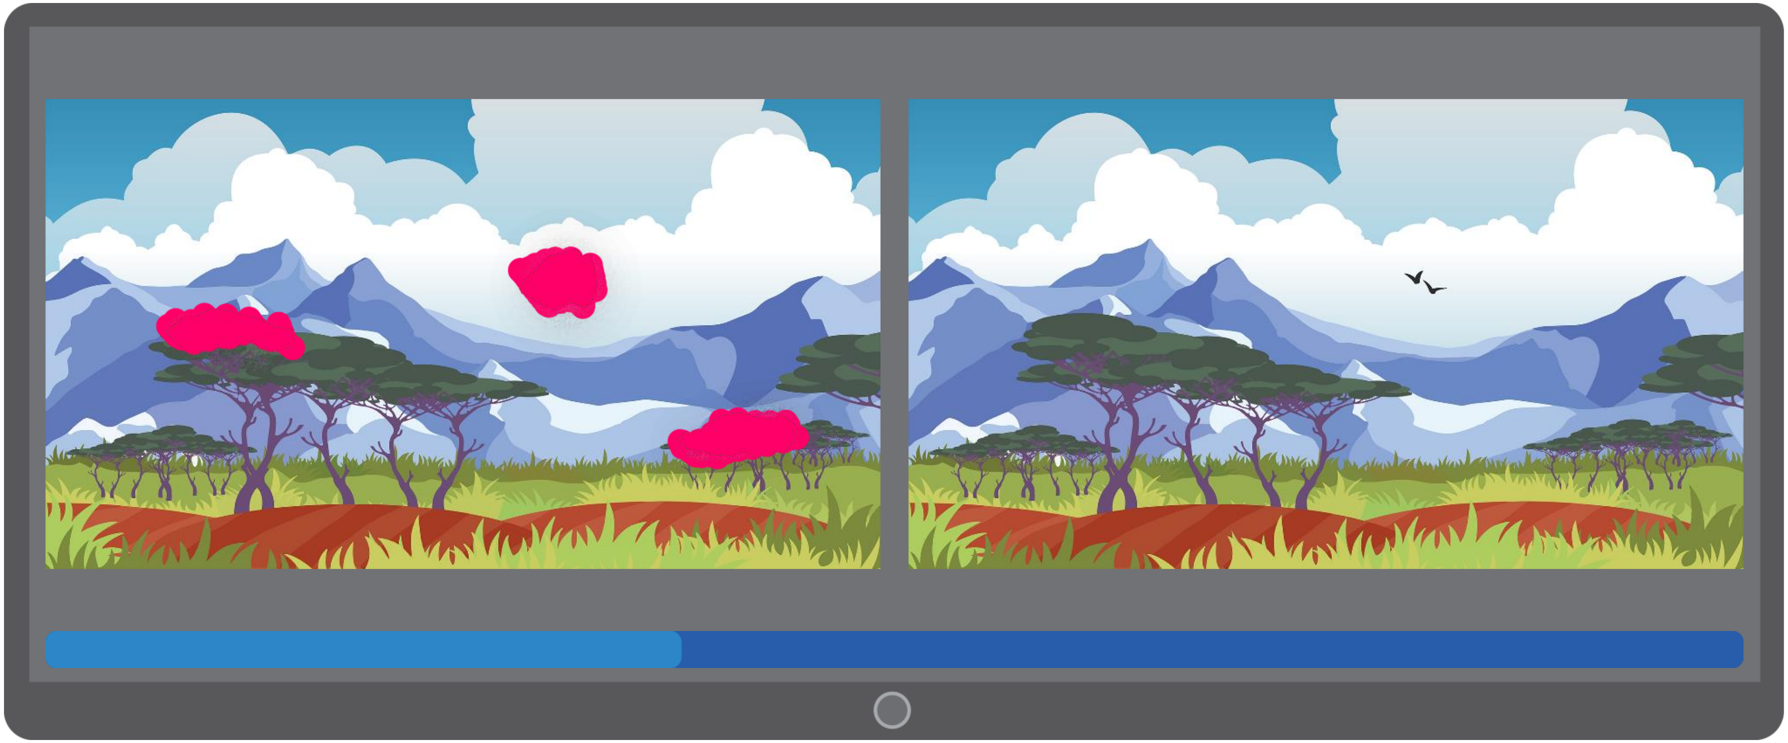

Label Creation

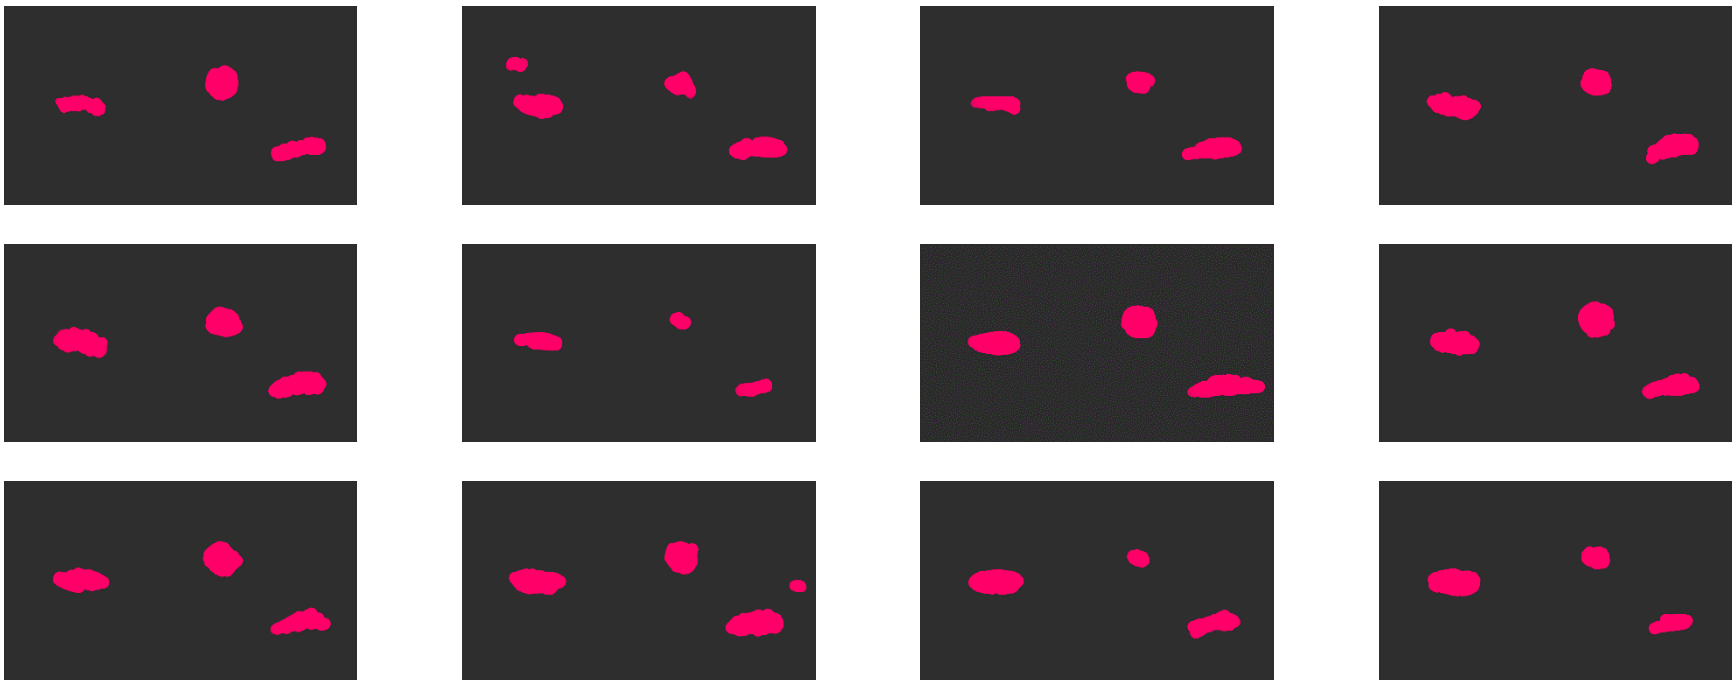

Neural Network Architecture

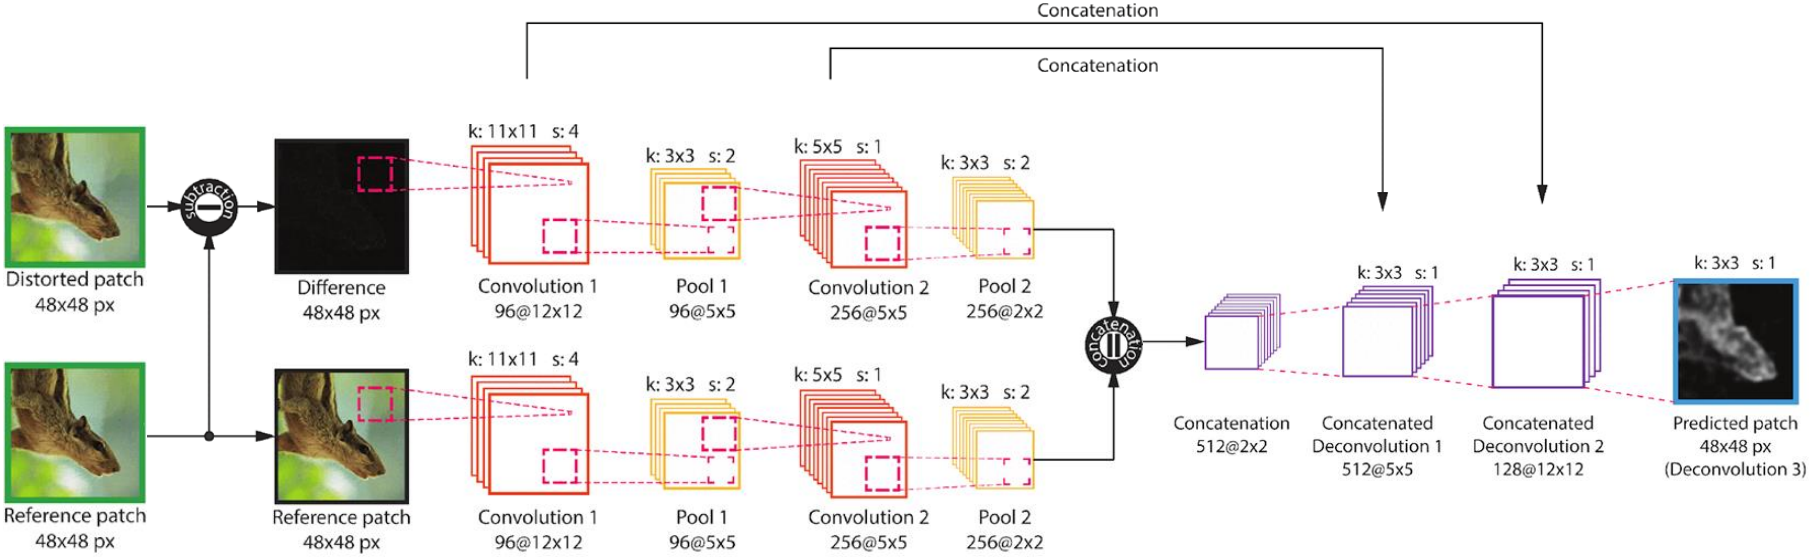

Multi-material Printing

Stratasys J750 (poly-jetting printer)

Cyan

Magenta

Yellow

blacK

White

Vero Opaque materials
(not actually opaque!)

3D Appearance Printing

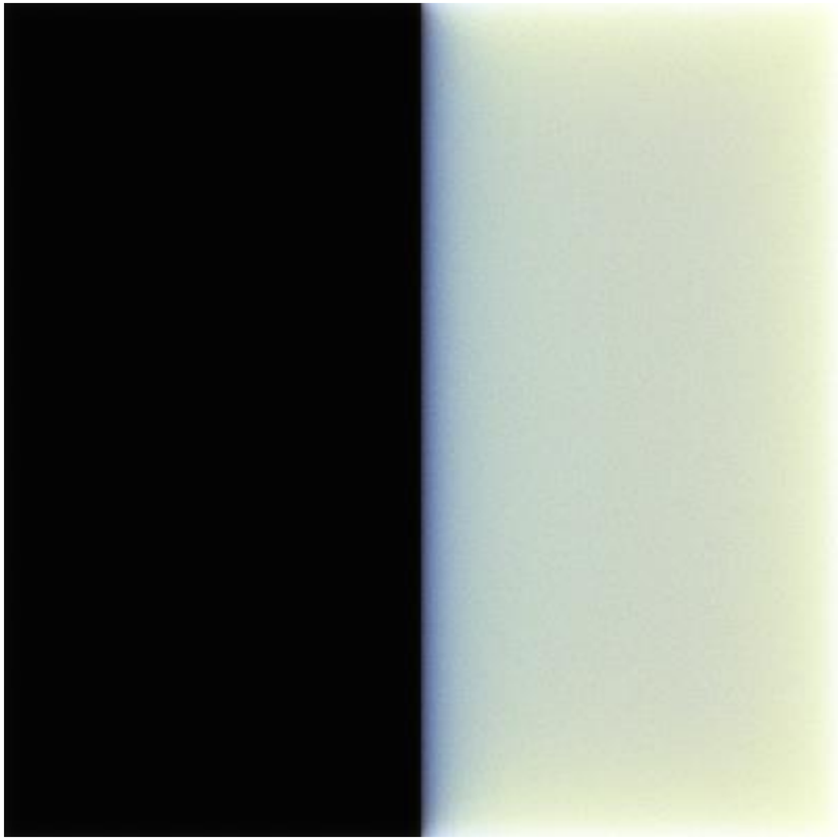

Goal: Visually Reducing Light Diffusion in the 3D Printed Material

3D Appearance Printing

Without correction (MC simulation)

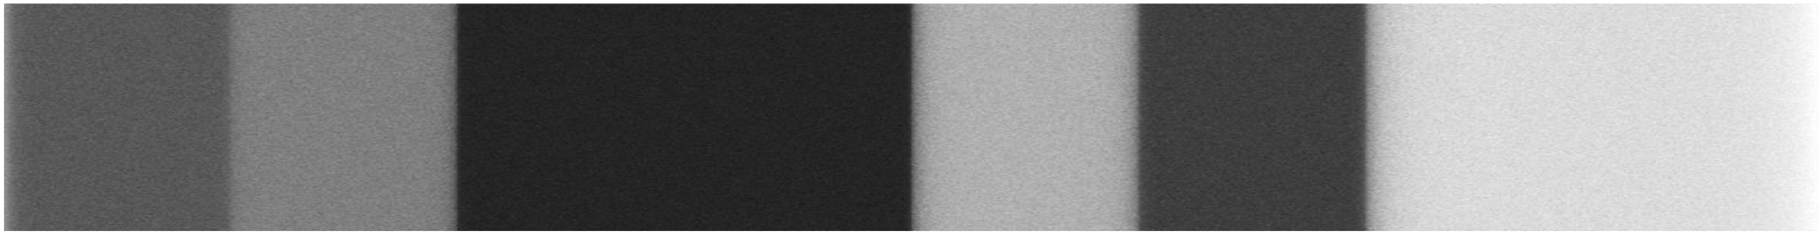

Target texture

With correction (MC simulation)

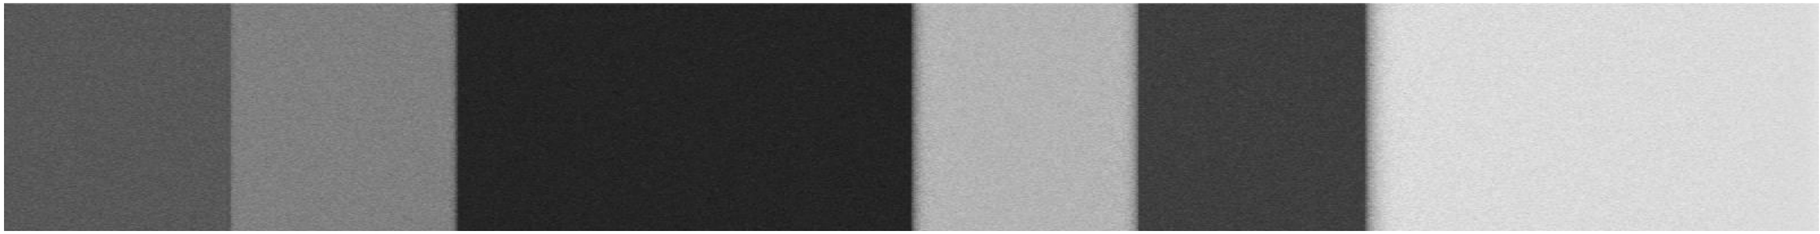

3D Appearance Printing

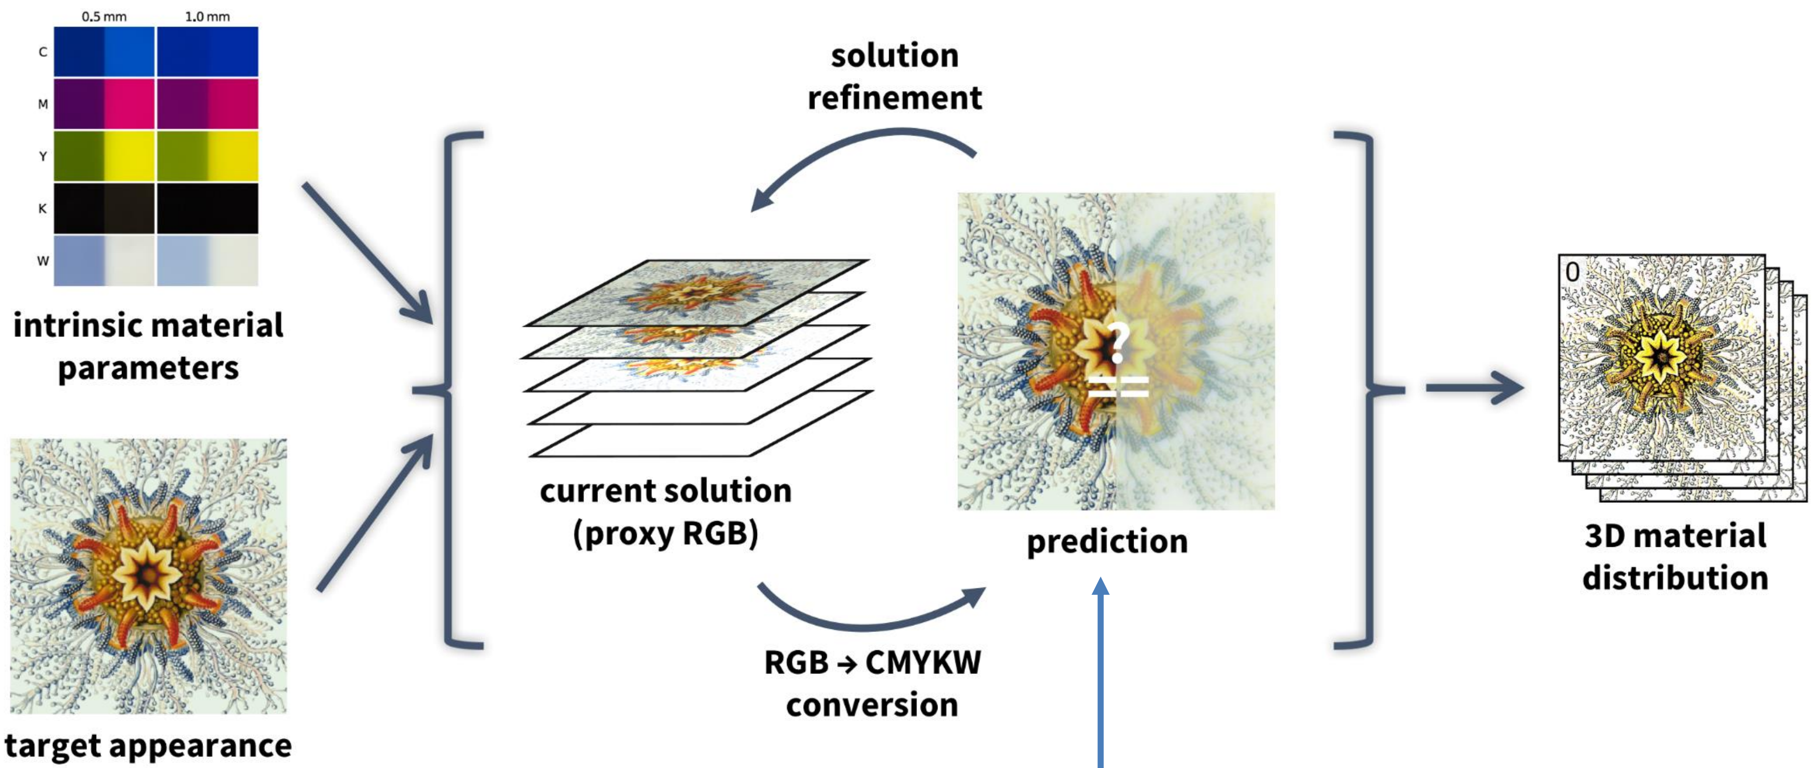

Volumetric MC global illumination simulation

3D Appearance Printing

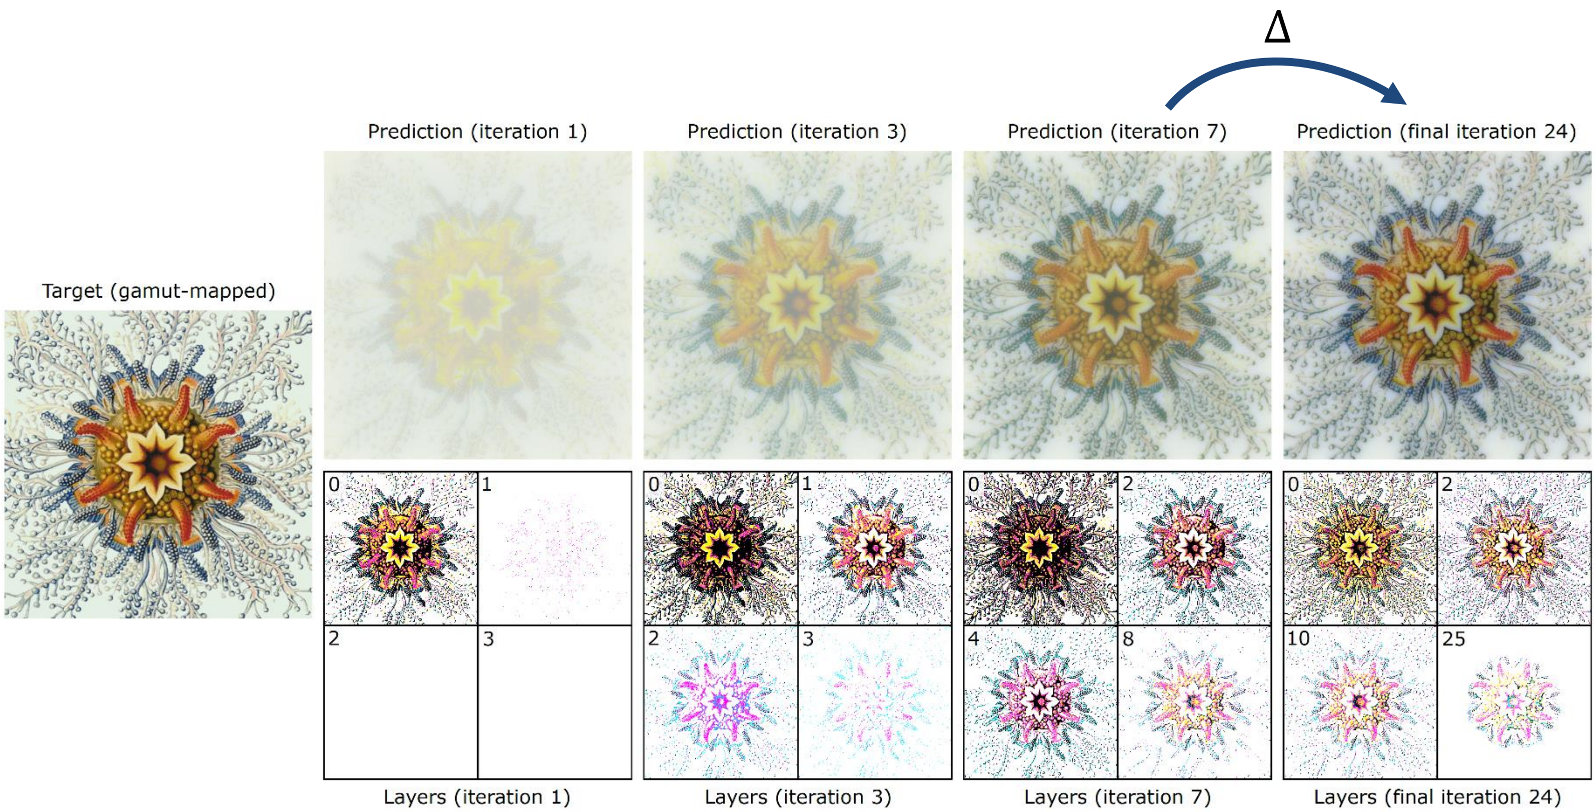

despite the non-linearity of the appearance, it changes monotonically
→ simple residual energy minimization

3D Appearance Printing

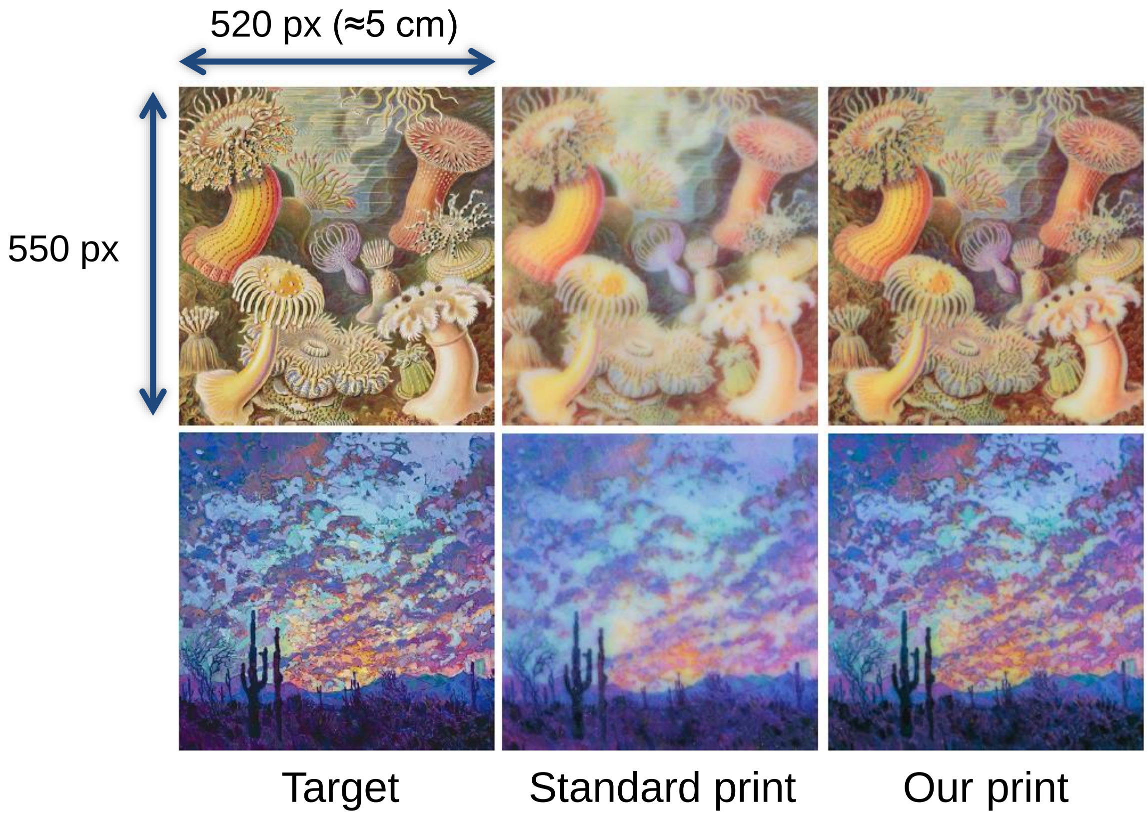

Varifocal Displays

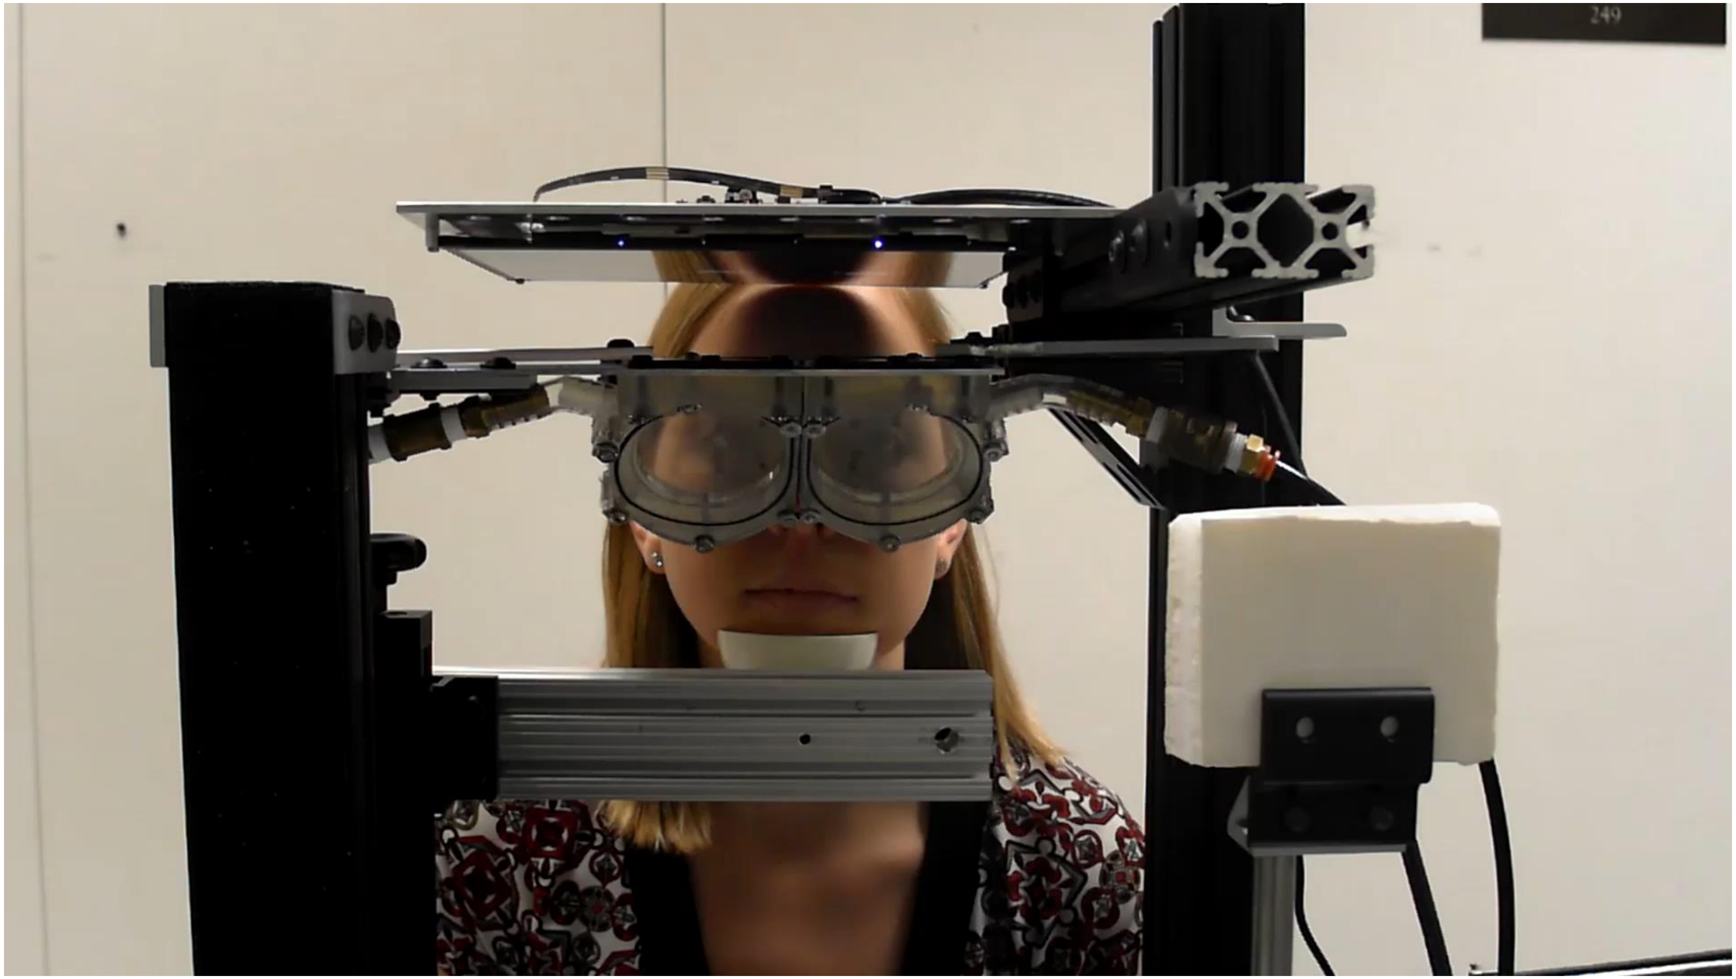

Membrane AR – Dunn et al.

Deformable Beamsplitter

Membrane AR – Dunn et al.

Dynamic focal depth: objects at any depth

Wide field of view

Optics are simple

Deformable Beamsplitter

Membrane AR – Dunn et al.

Multi-focal Plane Display

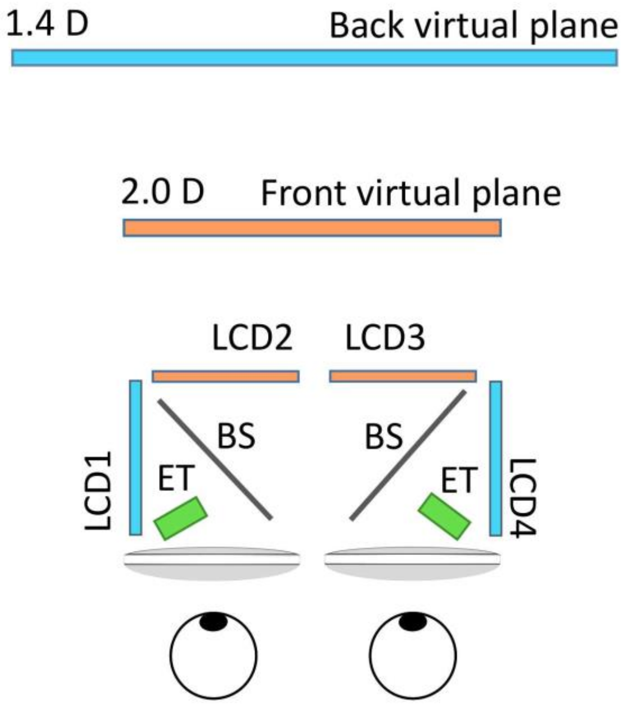

15cpd, 40 deg, 1200x1200 pixels

Multi-focal Plane Display

Temporal coherency

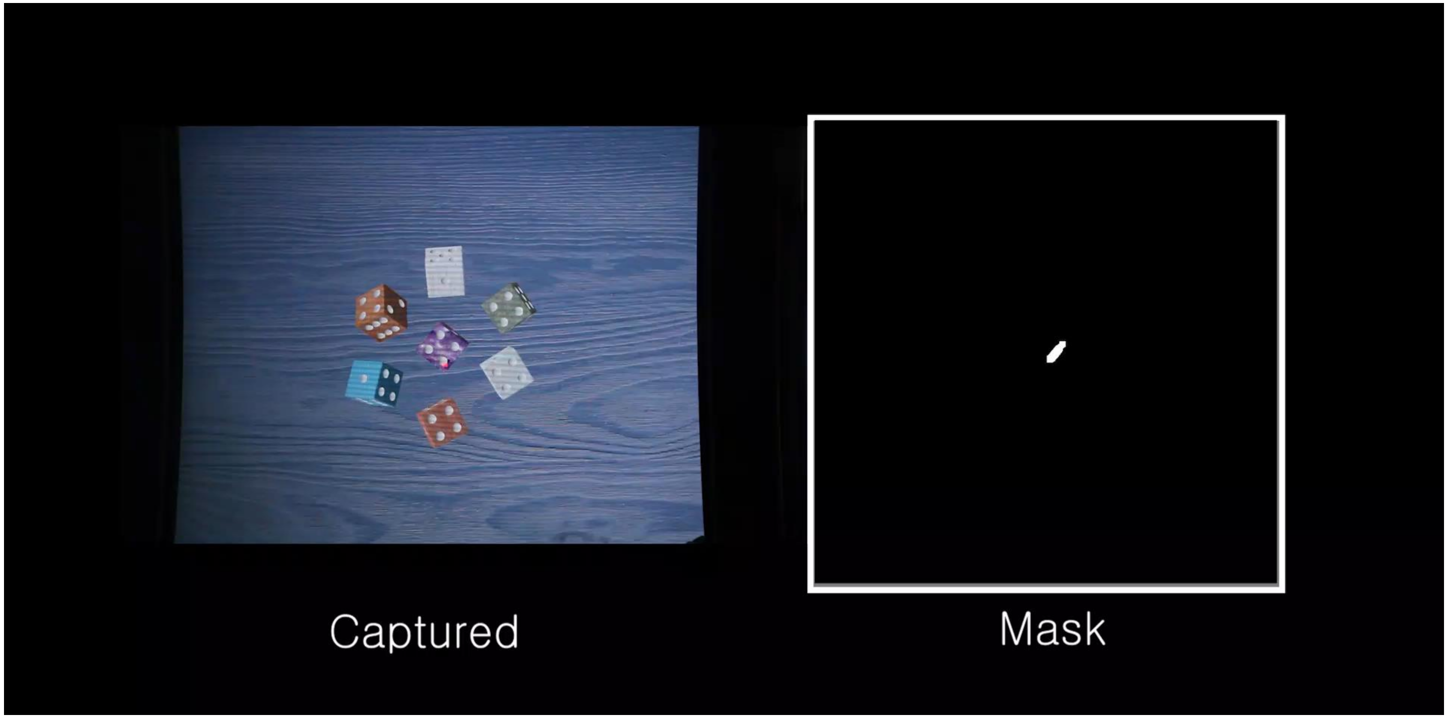

Saccade in Foveated Rendering

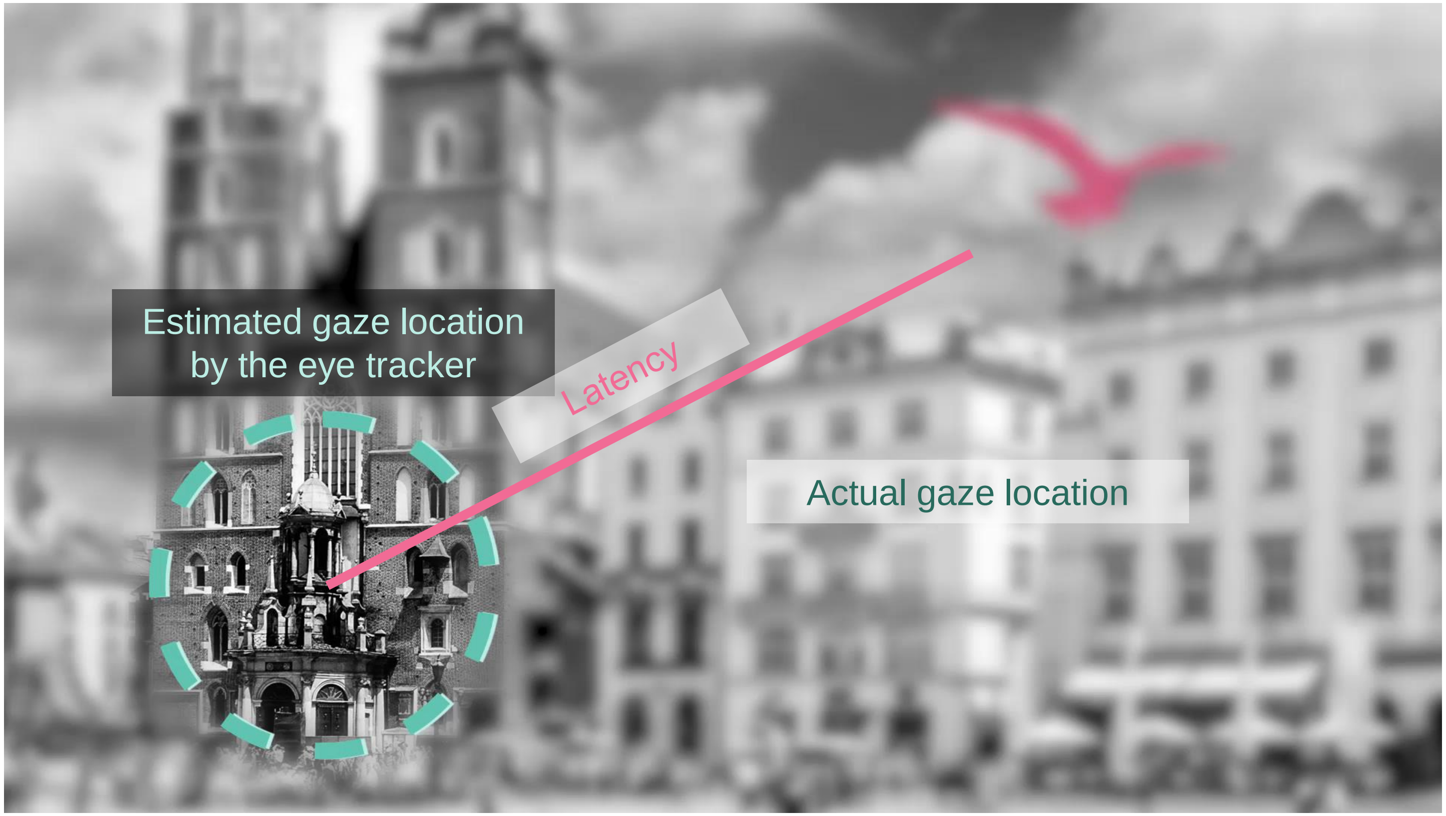

Saccade Landing Position Prediction for Gaze-Contingent Rendering

Saccade in Foveated Rendering

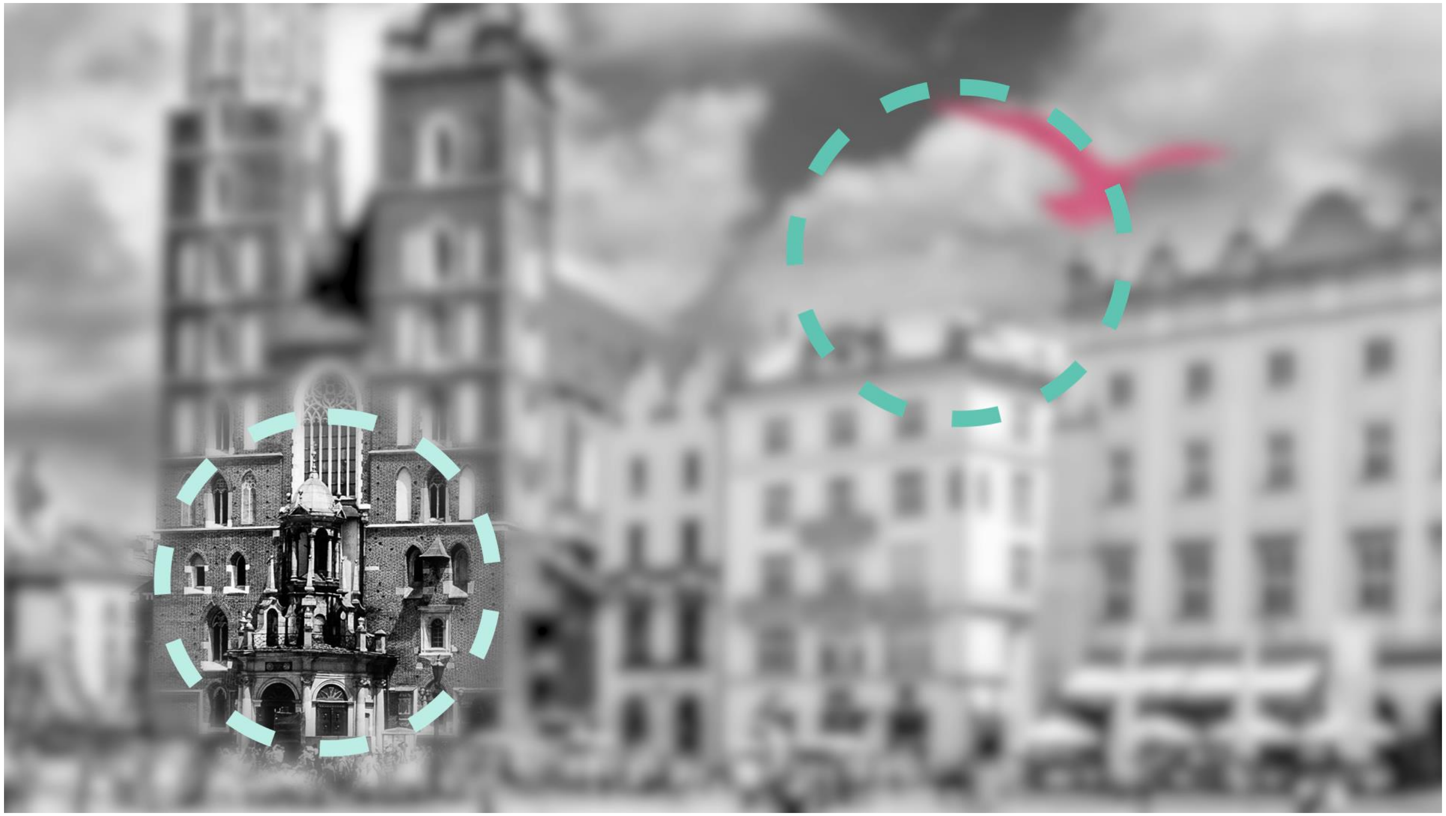

Saccade Landing Position Prediction for Gaze-Contingent Rendering

Eye Tracking: Saccade Landing Prediction

Eye Tracking: Saccade Landing Prediction

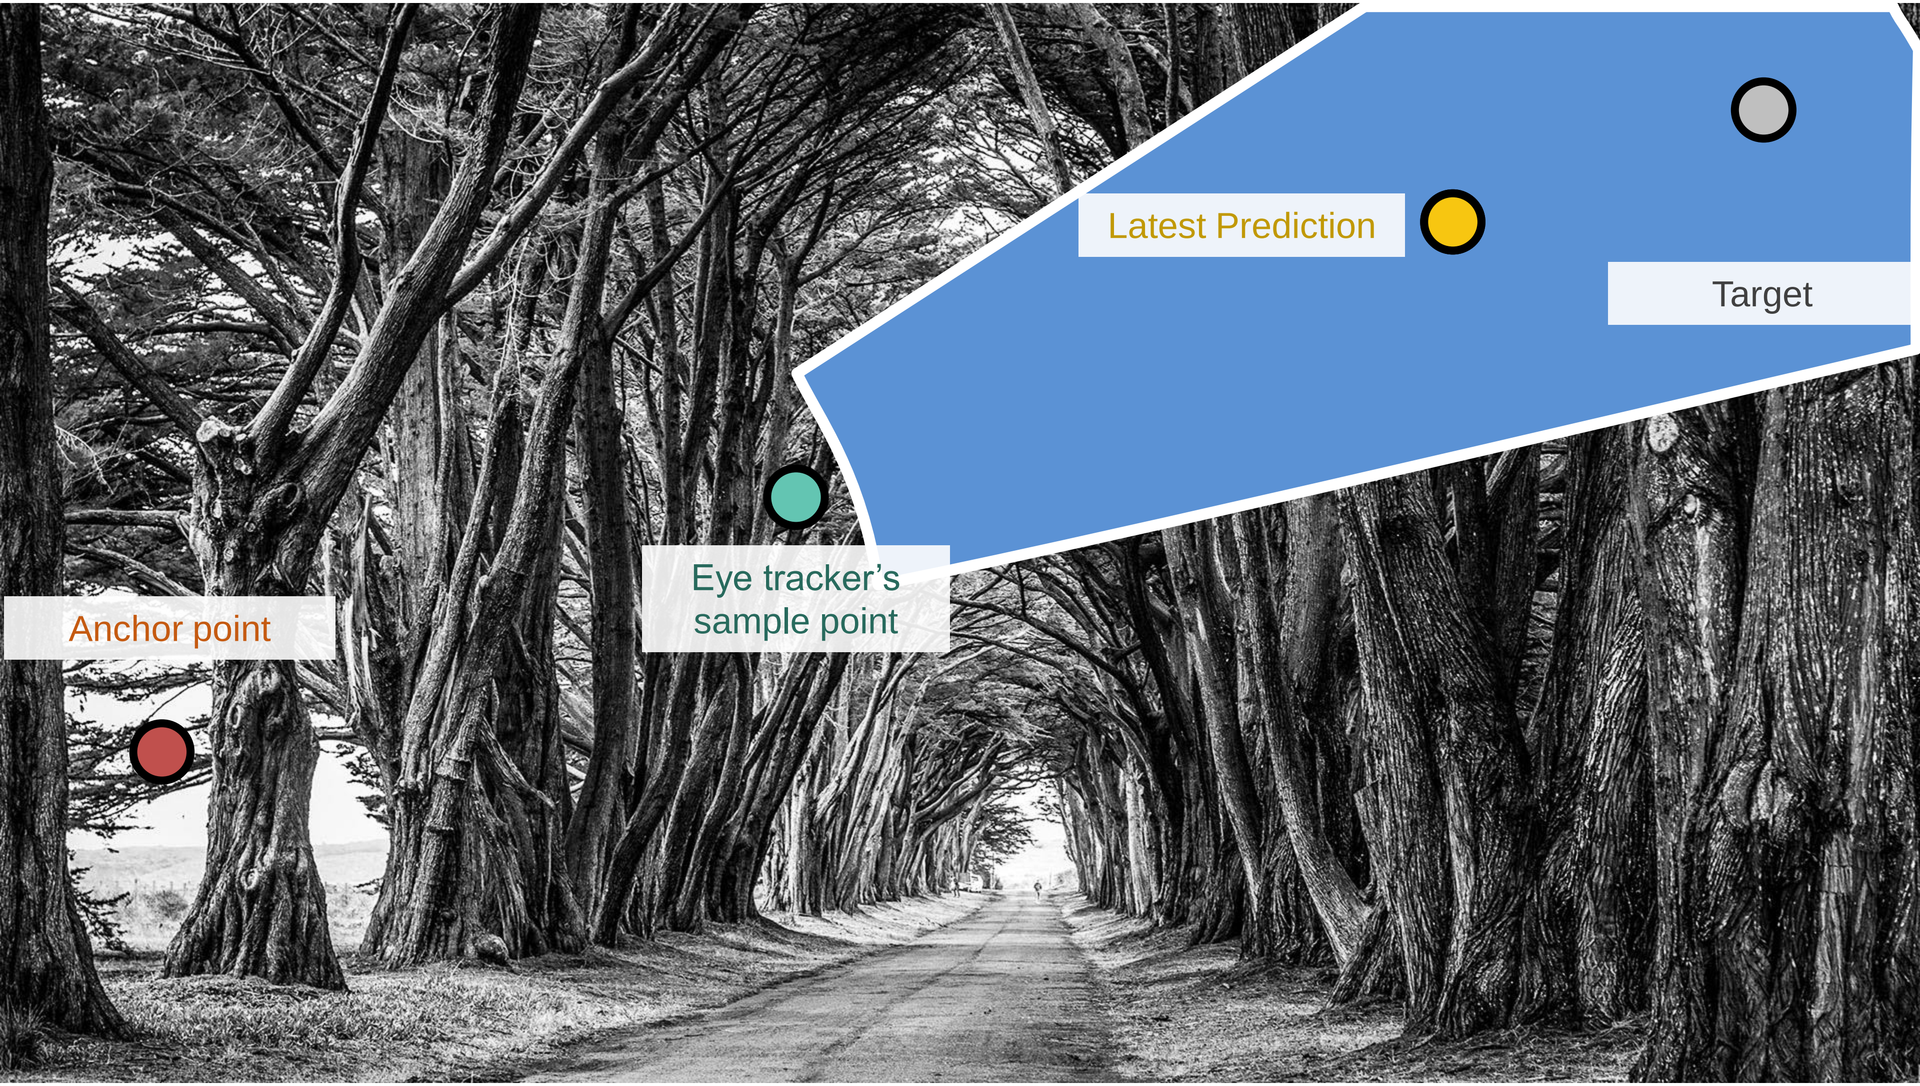

Eye Tracking: Saccade Landing Prediction

Eye Tracking: Saccade Landing Prediction

Eye Tracking: Saccade Landing Prediction

Eye Tracking: Saccade Landing Prediction

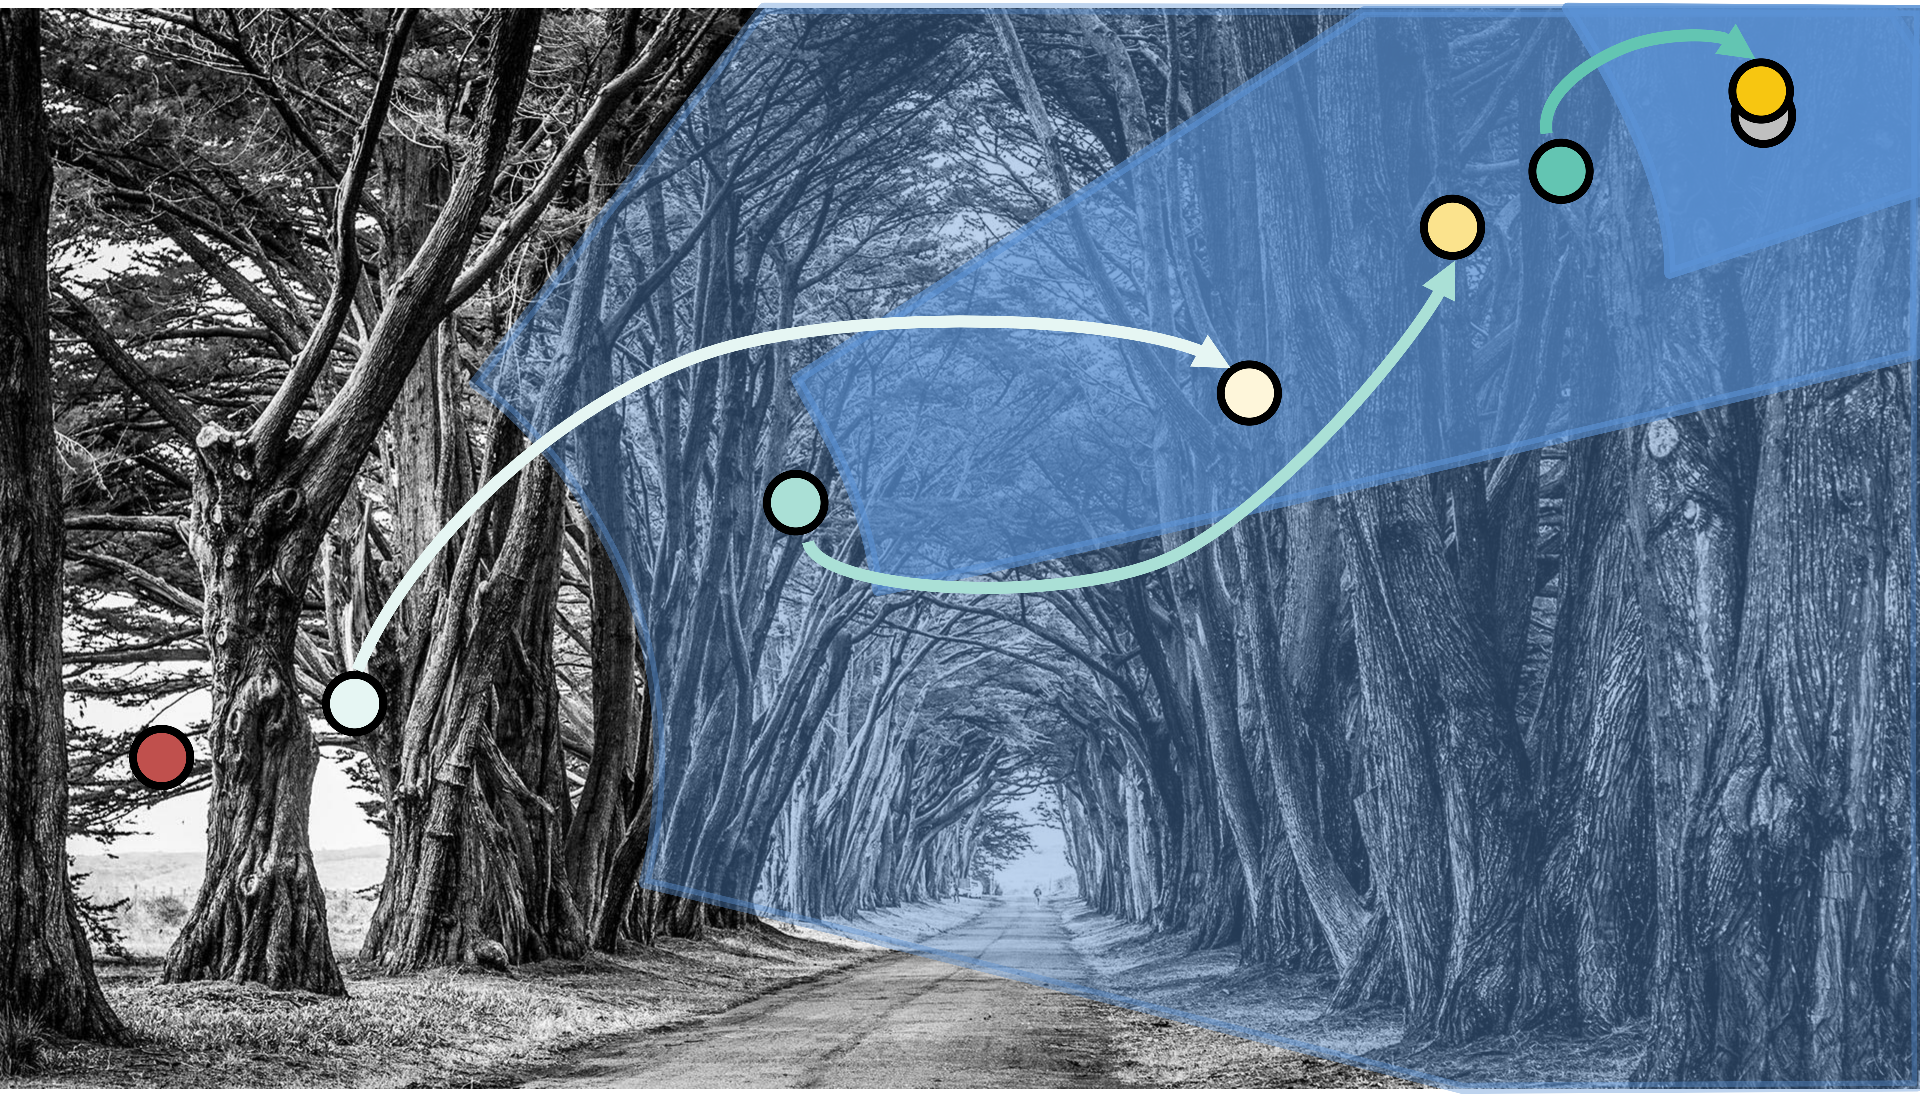

Eye Tracking: Saccade Landing Prediction

Luminance-Contrast-Aware Foveated Rendering

Luminance-Contrast-Aware Foveated Rendering

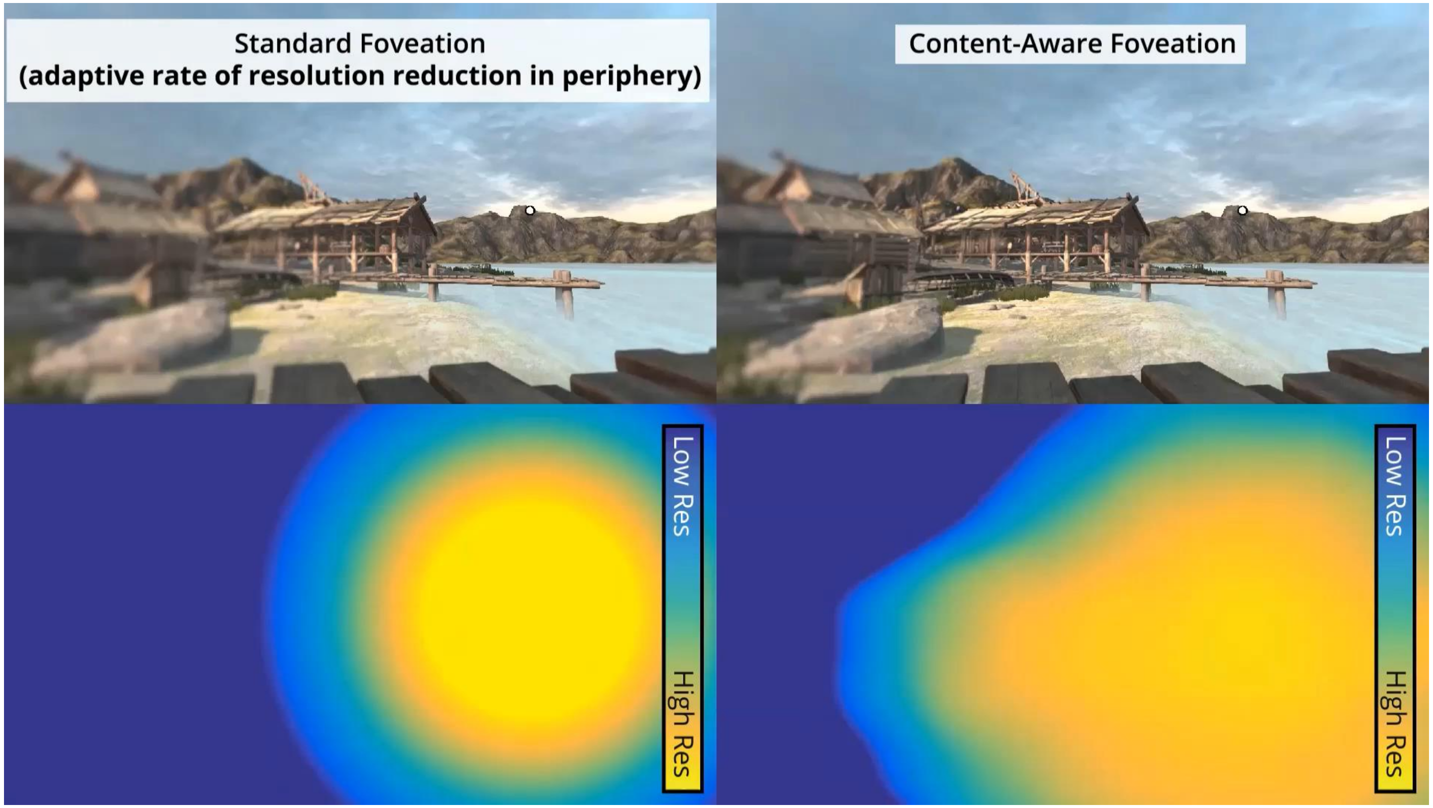

Sampling Patterns

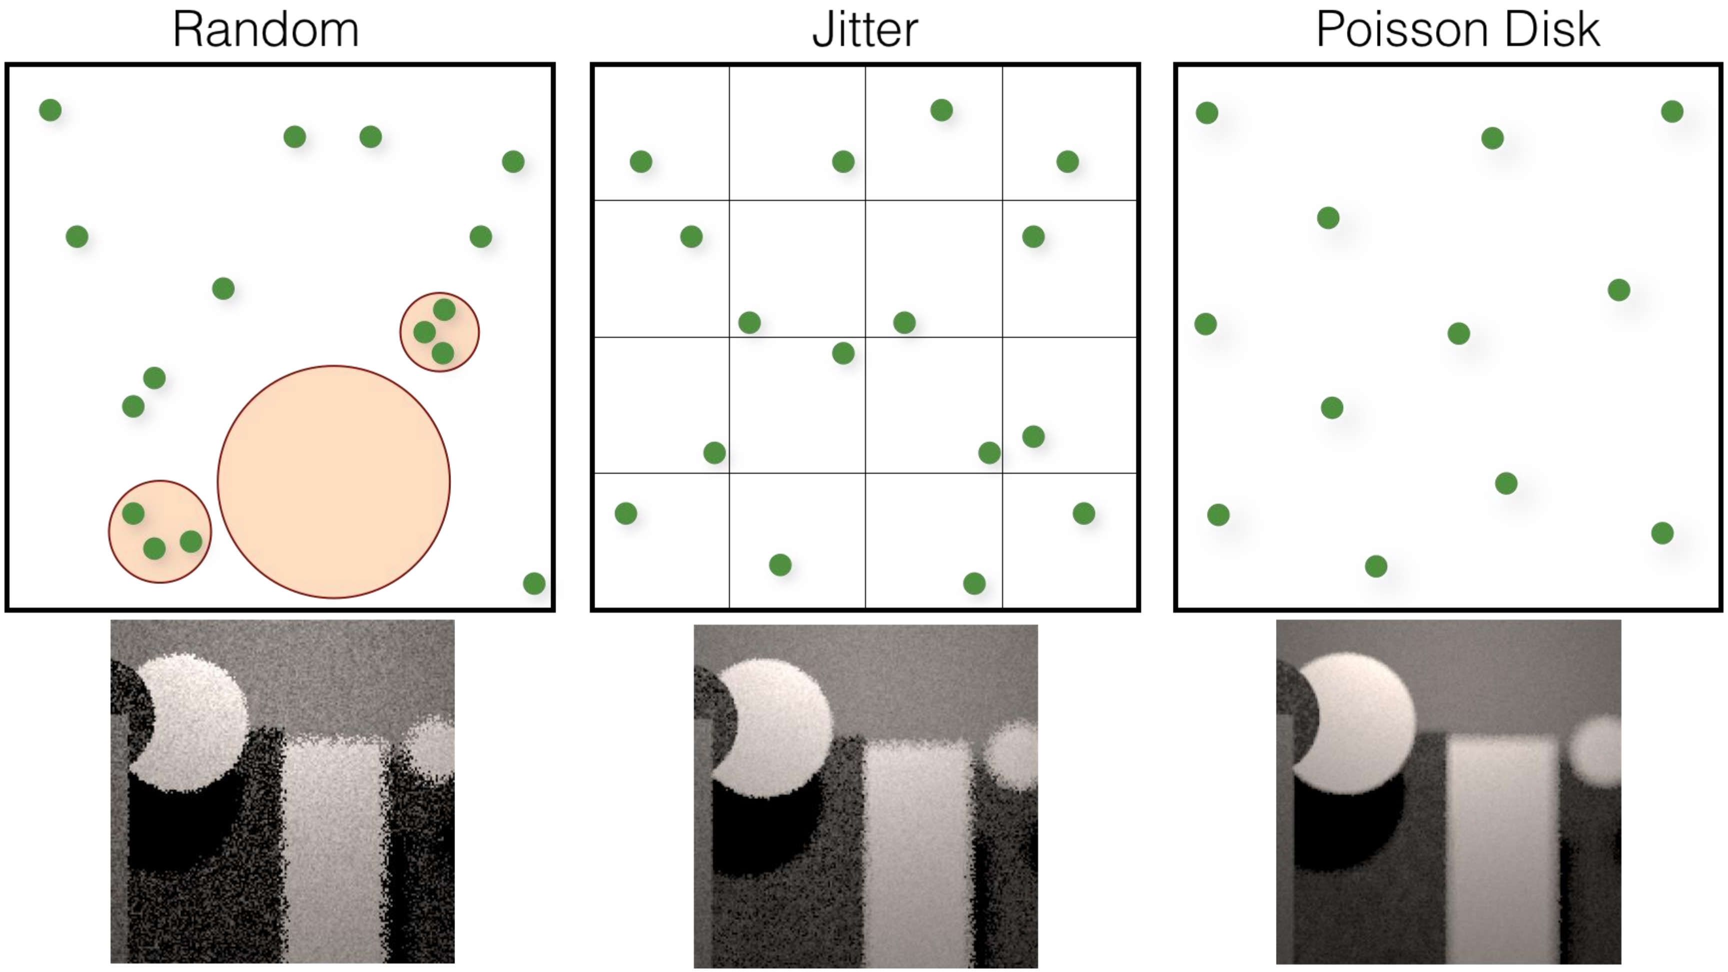

Advanced Sampling

How error in MC integration is affected by different sampling patterns?

Spatial domain statistics: Pair Correlation Function / Discrepancy

Fourier domain statistics

Define Error in terms of Spatial and Fourier domain statistics

Learn to Render: Path to Neural Networks

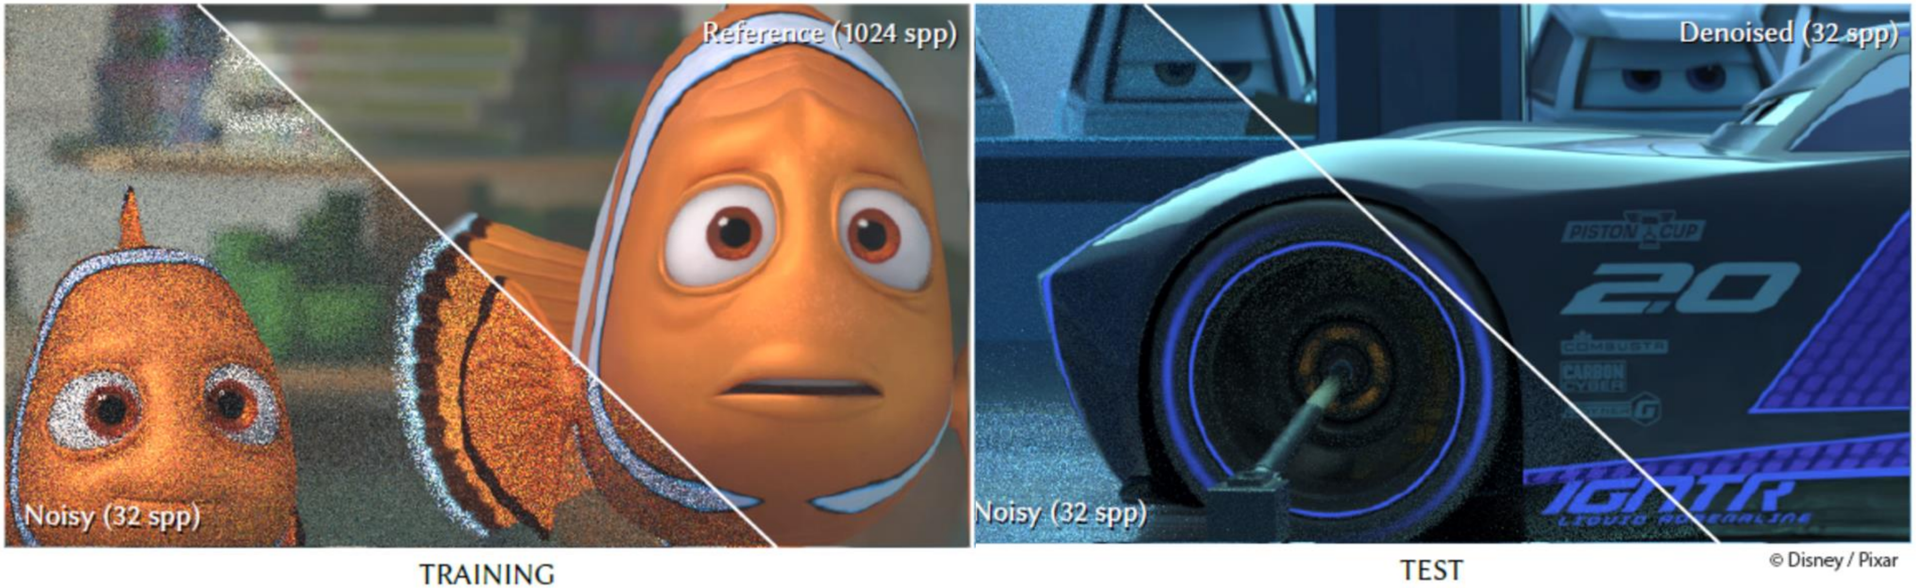

Bako et al.[2017]

Our Focus: Learn to Render

- ML/NN algorithms for denoising
- CNNs/GANs (unstructured)
- Learning Light Transport the Reinforced Way
- Learning to Importance Sample