Realistic Image Synthesis

Introduction


Which One is Real?
A
B
© 3DAS

What?

- Goal: Create photorealistic images
- Applications
    - Movies and games
    - Design and architecture
    - Visualization and simulation
    - Optimization, inverse rendering
    - AI and machine learning

Who?

- 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
- Tutor
  - NN

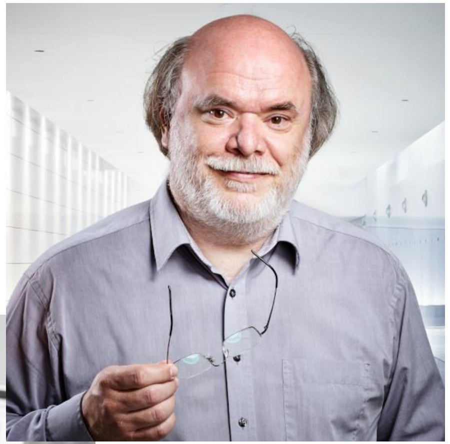

Administrative information

- Type
  - Advanced lecture
  - 9 credit points
- Prerequisites
  - Interest in math, physics
  - Basic programming experience in C++
  - Core lecture “Computer Graphics” recommended but not required
- Web-page: https://graphics.cg.uni-saarland.de/courses/ris-2024/
- MS Teams (Join via link on the webpage)
  - Announcements, Q&A, …
  - Assignments posted and submitted

Grading

- Exam admission requires
  - 50% of the total points across all assignments
  - 30% of the maximum points in **every** assignment
- Final grade
  - Assignments: 50%
  - Final exam: 50%

Assignments

- Irregular rhythm
  - Sometimes 1 week, sometimes 2
- Type
  - A few theoretical assignments
  - Mostly practical ones
- Teamwork
  - Can be done in groups of two
  - Make sure you understand everything your partner worked on!
- Published, handed-in, and graded via MS Teams

Reading materials

- Pharr, Jakob, and Humphreys. *Physically based rendering: From theory to implementation*. Morgan Kaufmann, 2016.
    - Free e-book: http://www.pbr-book.org/
- More listed on the website

Applications

Where are the things you will learn here used?

Movies: Visual Effects (VFX)

Game of Thrones

Avatar: The Way of Water

Movies: Animated Films

The Lion King (2019)

The Sea Beast

Video Games

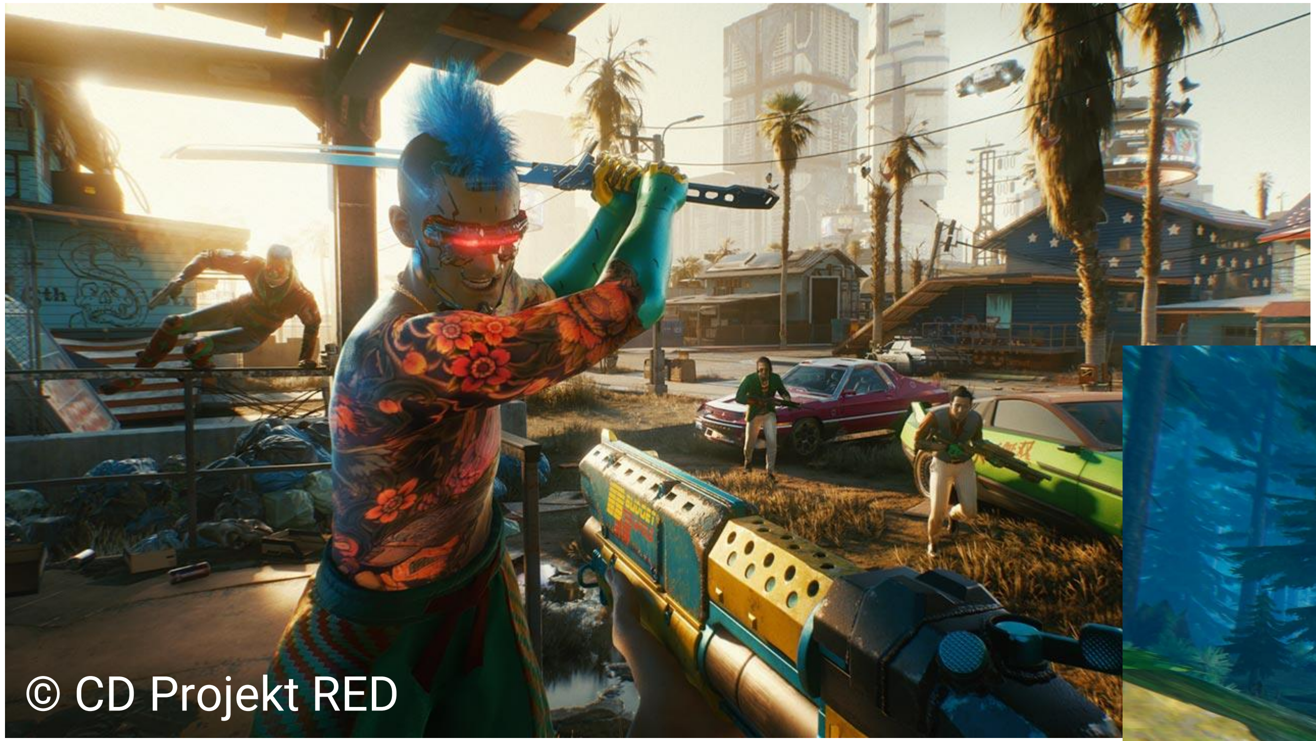

Cyberpunk 2077

Valheim

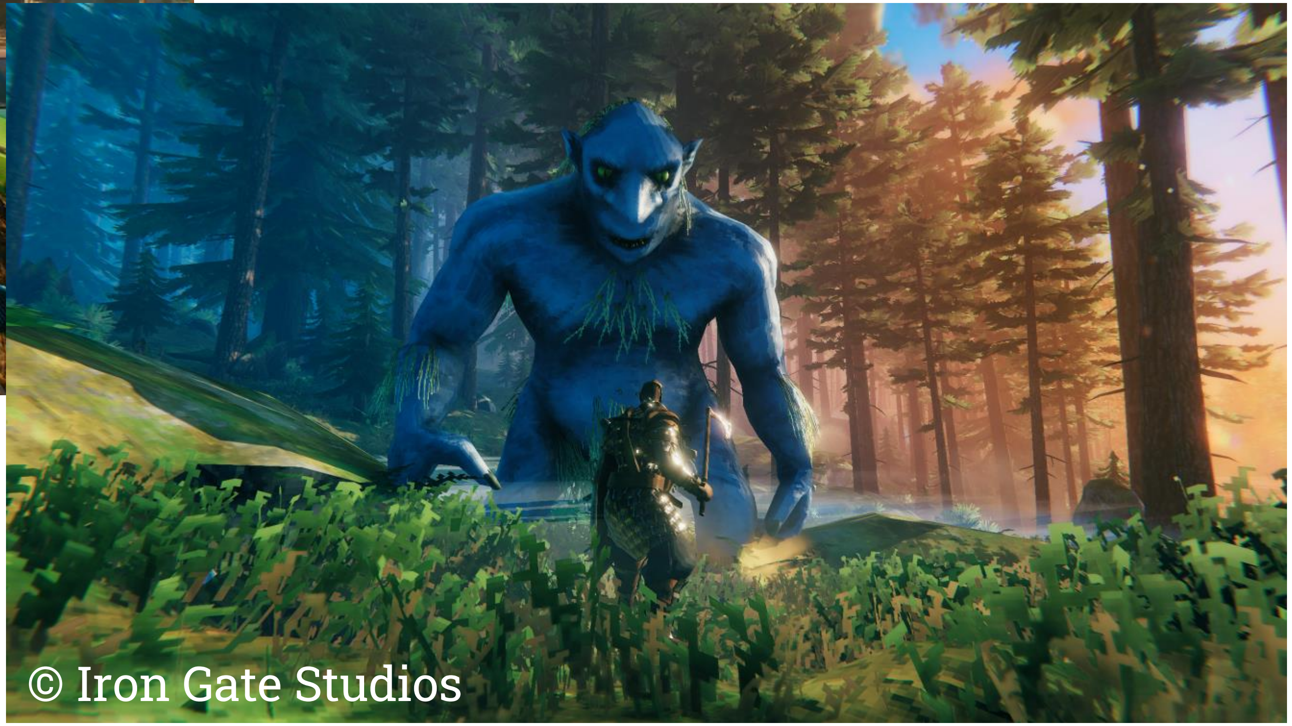

Simulation

© Thomas Angus / ICL

Design and Engineering

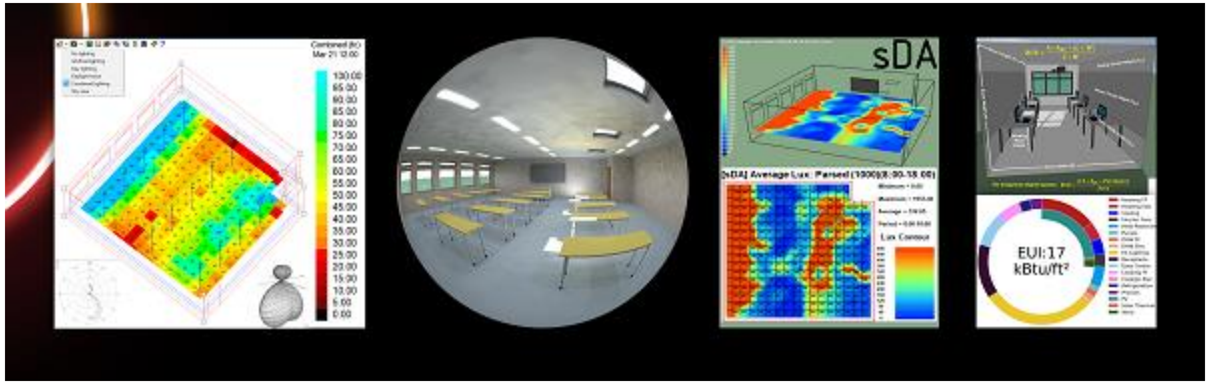

© IES

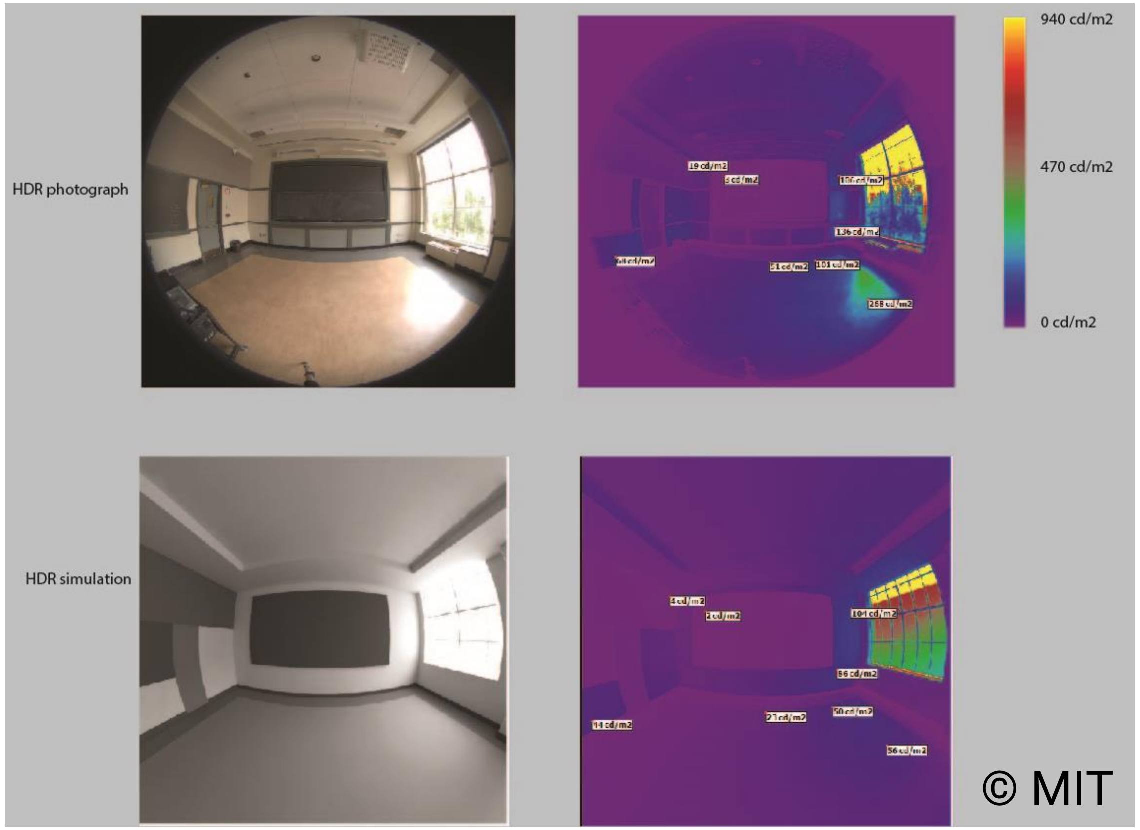

Product Visualization and Advertisement

© IKEA

© Nissan

Architecture

© Pixelcraft Work

Optimization and Inverse Rendering

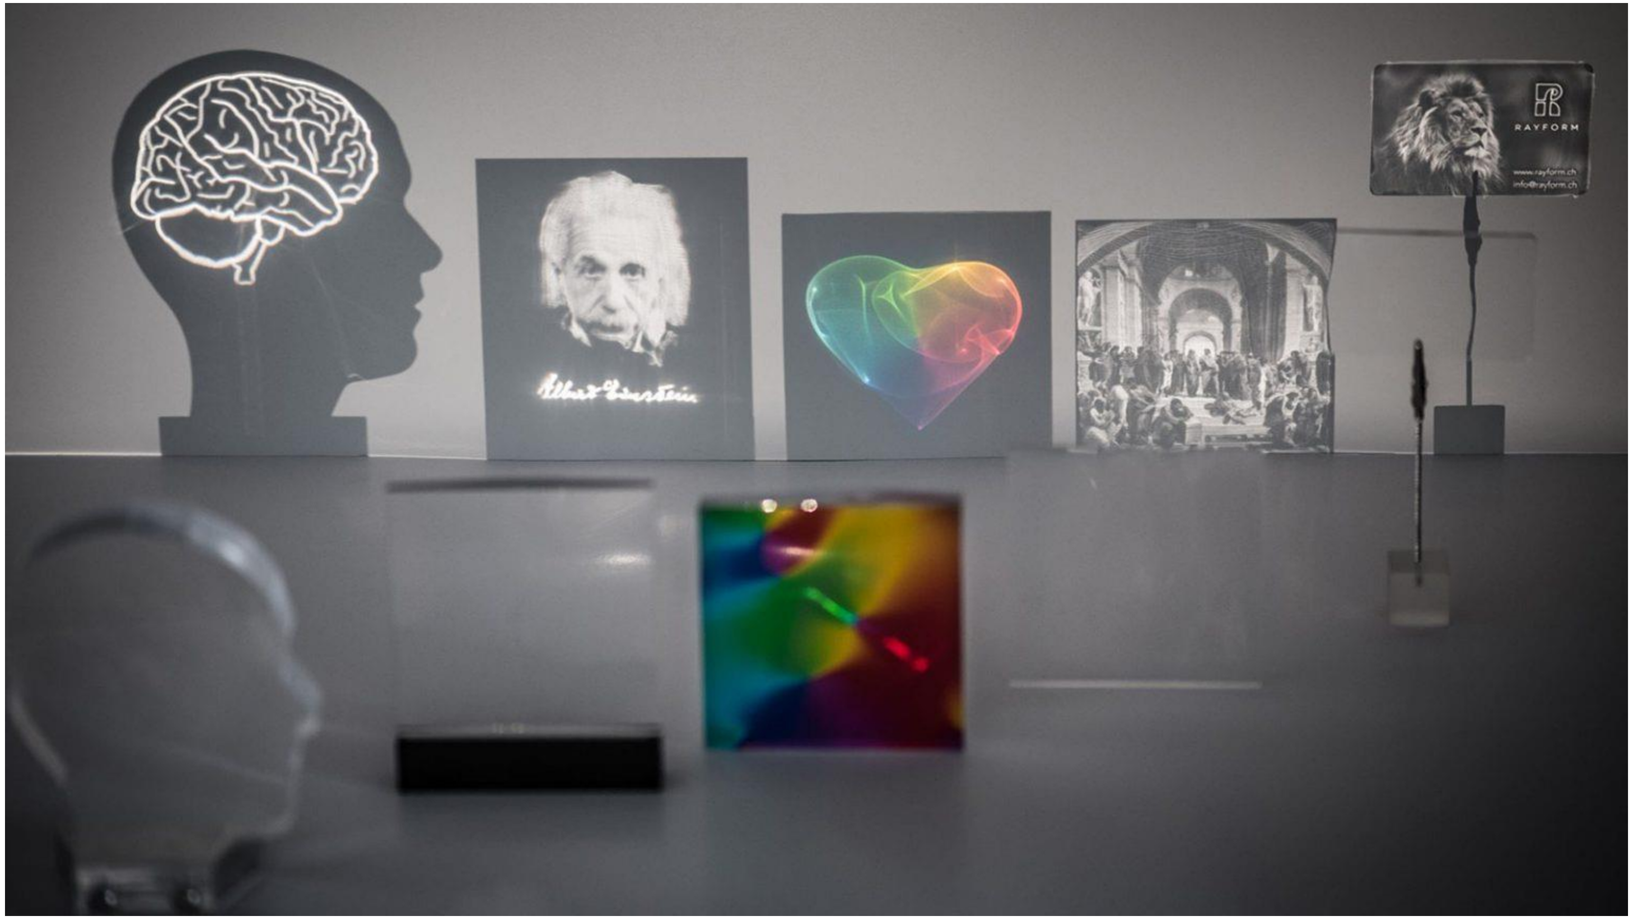

© Schwartzburg et al. 2014

Artificial Intelligence

Course overview

What will you learn?

Course Overview

- Core concepts
  - Rendering equation
  - Radiosity
  - Probability theory and Monte Carlo integration
  - BRDFs and path tracing
  - Advanced sampling
- Bidirectional and adaptive algorithms
  - Bidirectional methods
  - Markov chain Monte Carlo
  - Path guiding
- Advanced effects
  - Volume rendering
  - Radar / Spectral
- Perception and imaging
  - HDR and tone mapping
  - Perception and modern display technology
- Machine learning
  - Denoising
  - Differentiable rendering

Rendering Equation

Outgoing light

Emitted light

Incident light
(recursively given by the same equation)

Projection

$$L_o(x,\omega_o) = L_e(x,\omega_o) + \int_\Omega L_i(x,\omega_i) f_r(\omega_o, x, \omega_i) \cos\theta_i \, d\omega_i$$

Integral over all directions
(computes reflected light)

BSDF: Material properties

Monte Carlo Integration and Path Tracing

$$\int_X f(x)\,dx \approx \frac{1}{n}\sum_{i=1}^{n}\frac{f(x_i)}{p(x_i)}$$

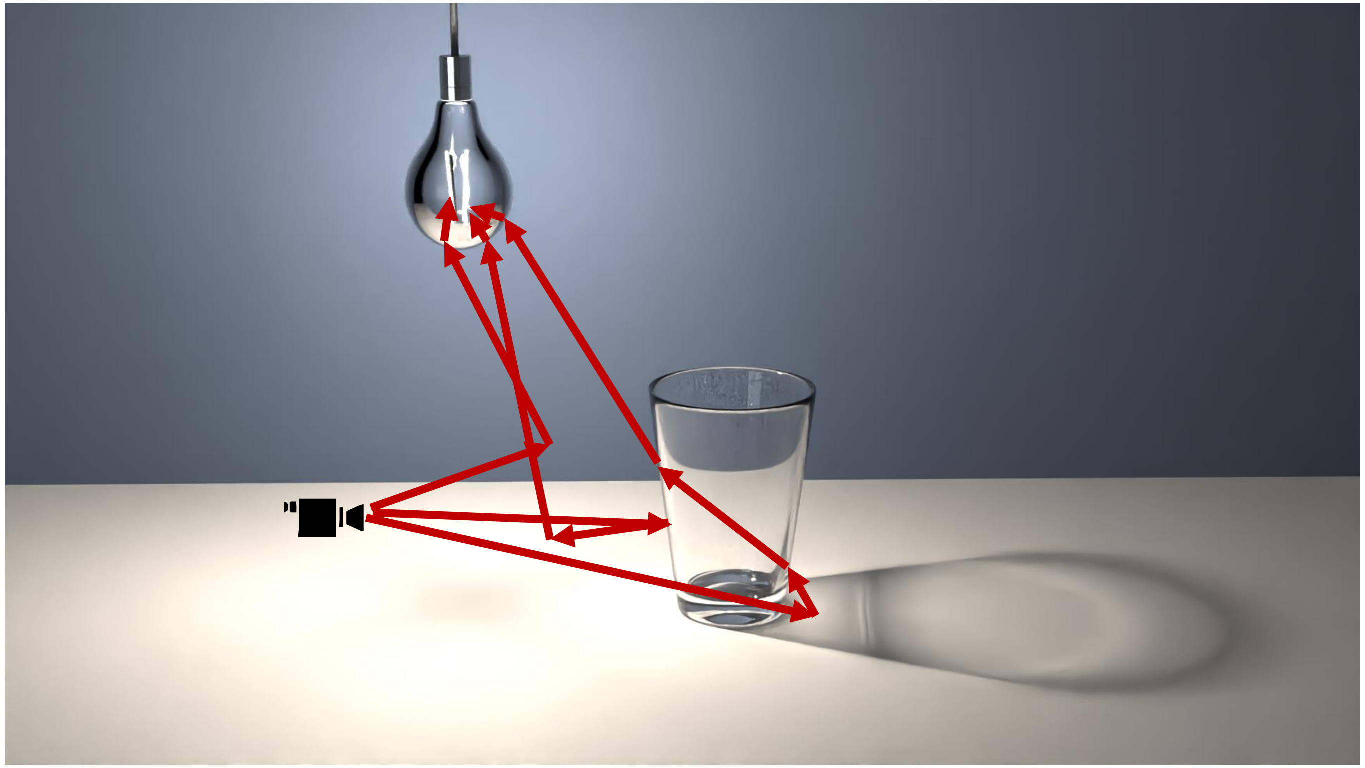

Advanced Sampling

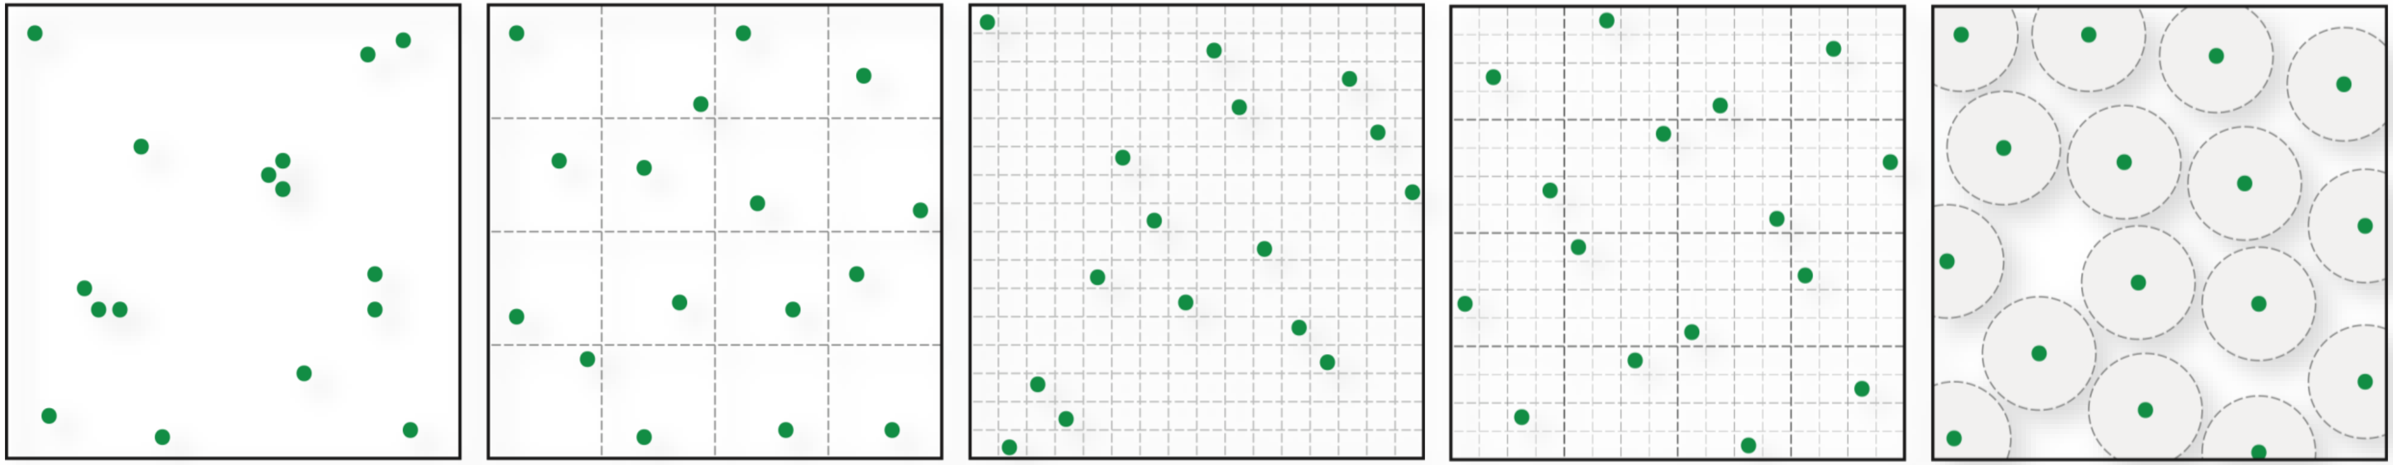

Bidirectional Methods

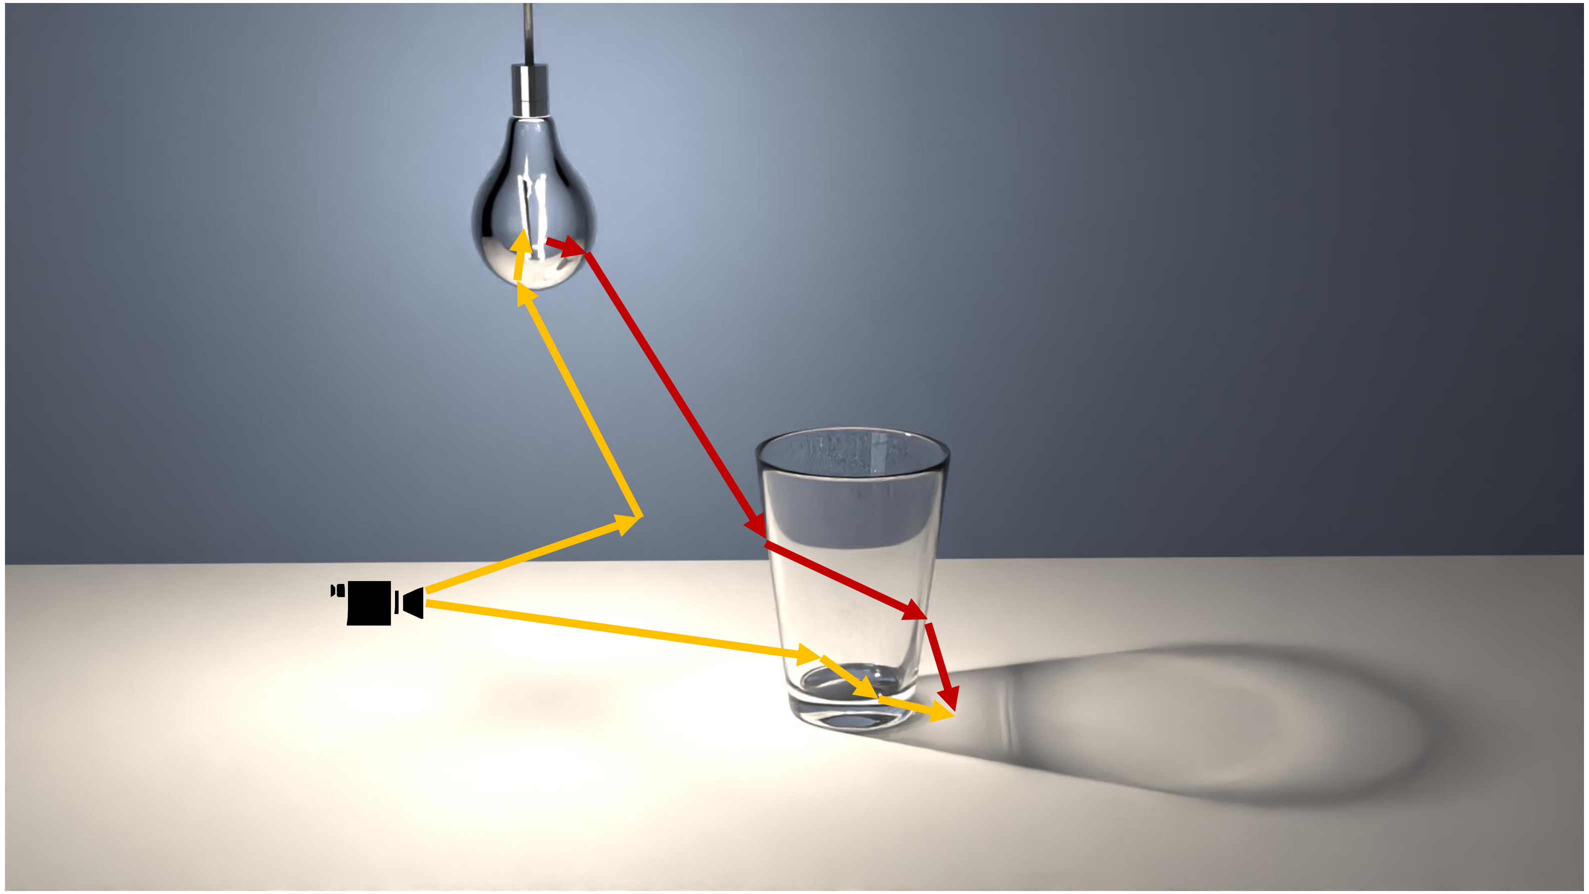

Adaptive / Learned Sampling

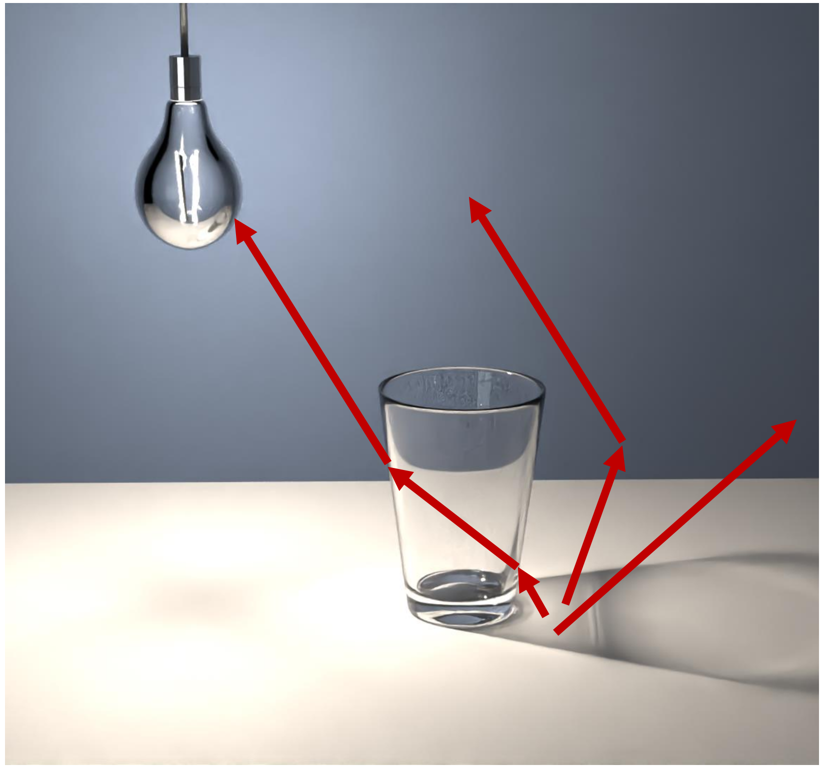

Initial training samples

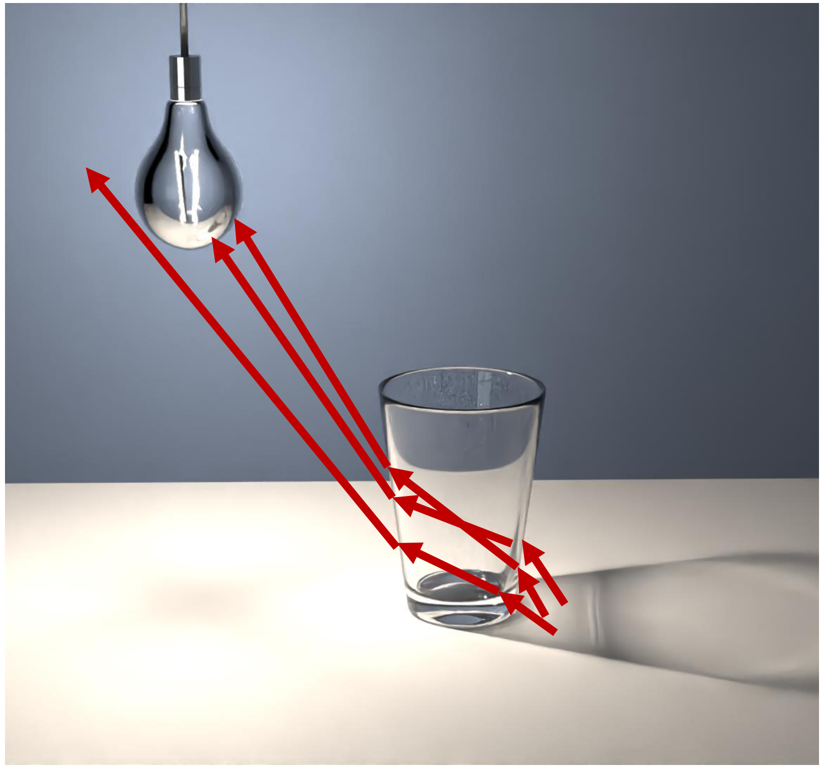

Guided samples

Volume Rendering

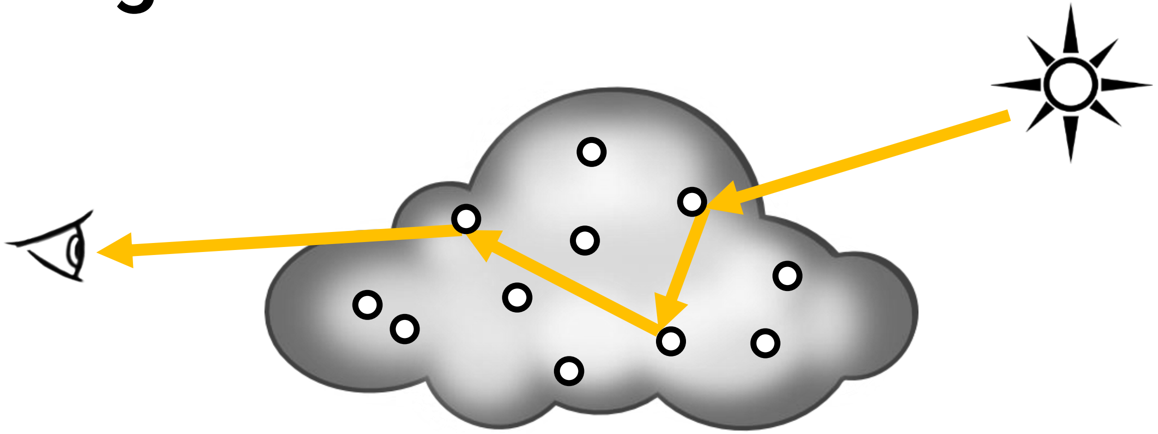

http://coclouds.com

http://wikipedia.org

http://commons.wikimedia.org

HDR and Tone Mapping

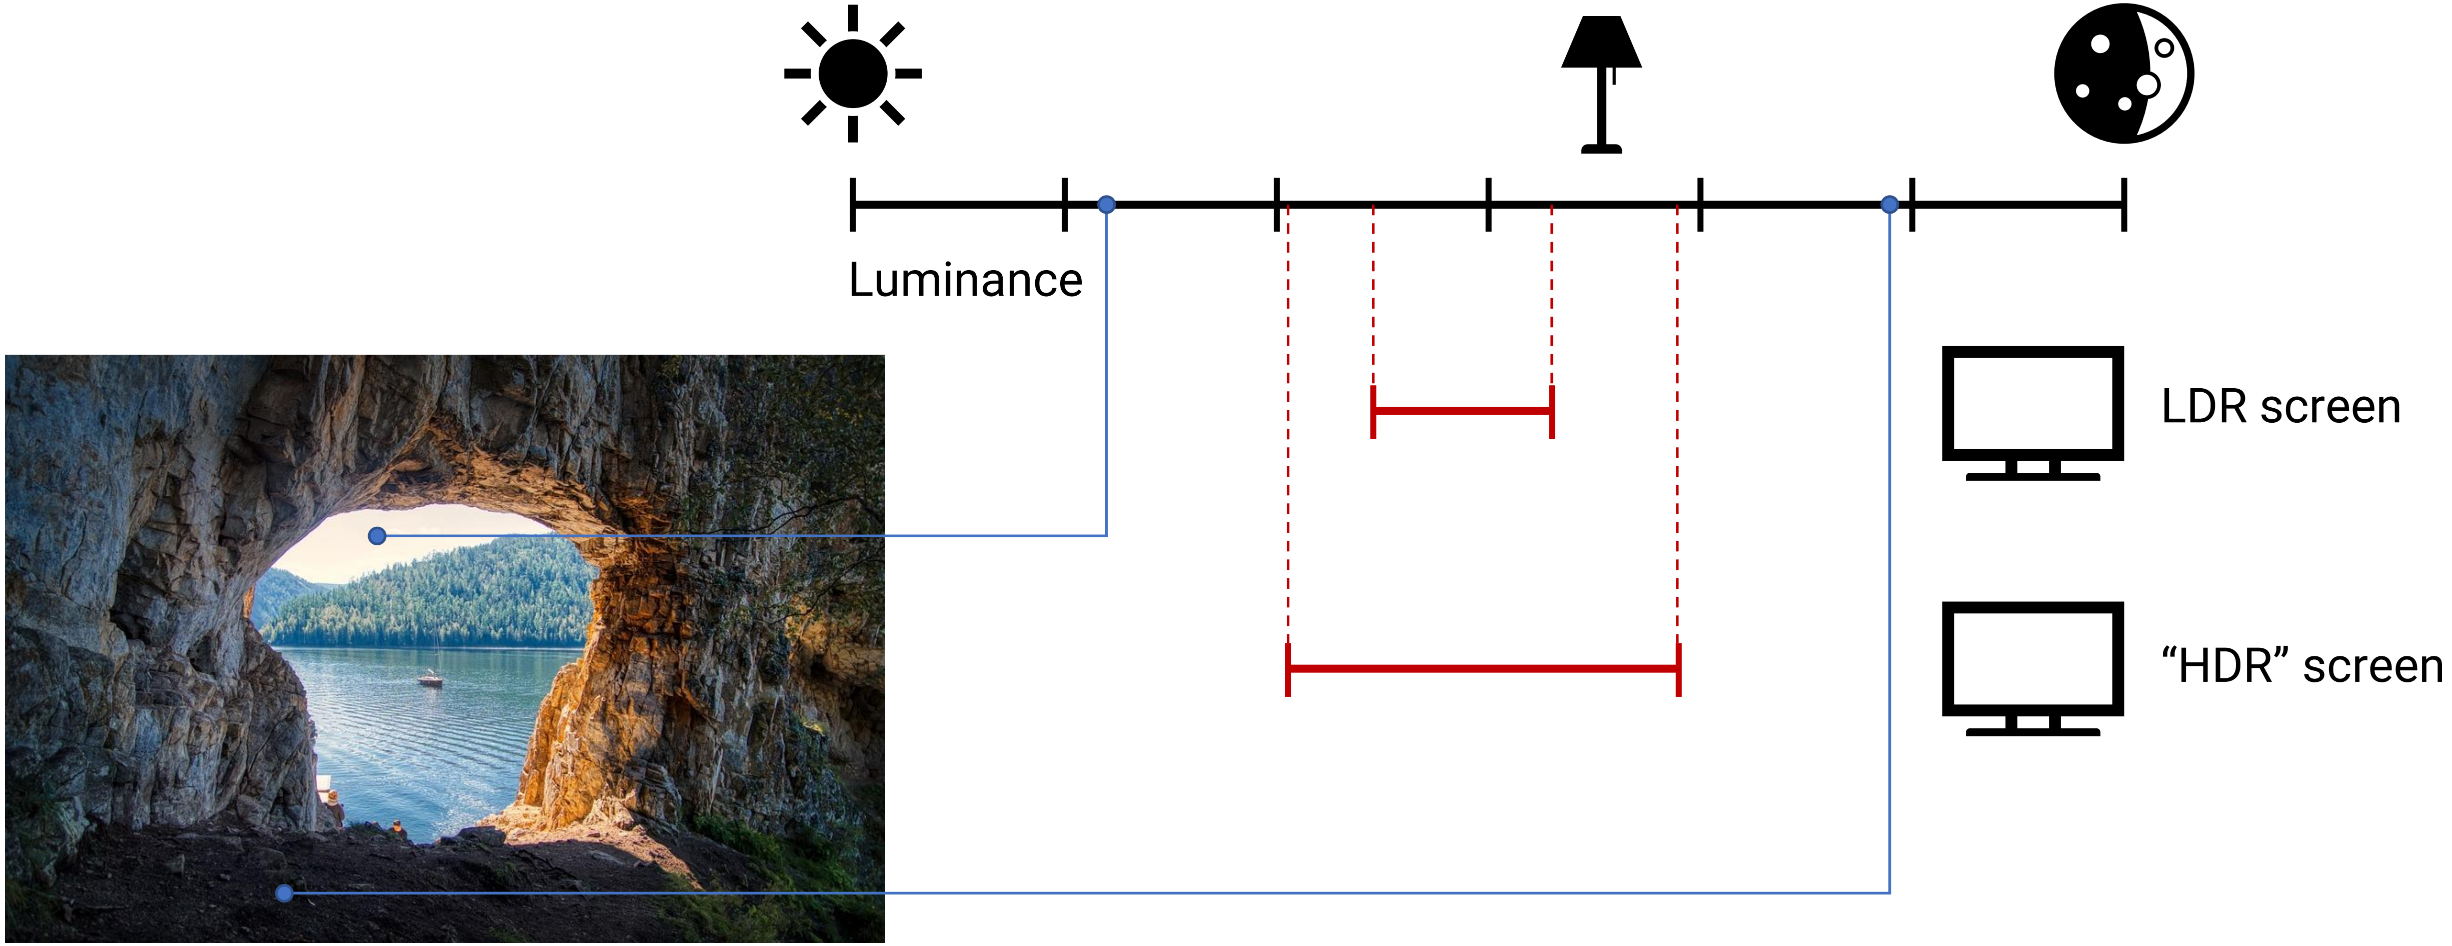

Denoising

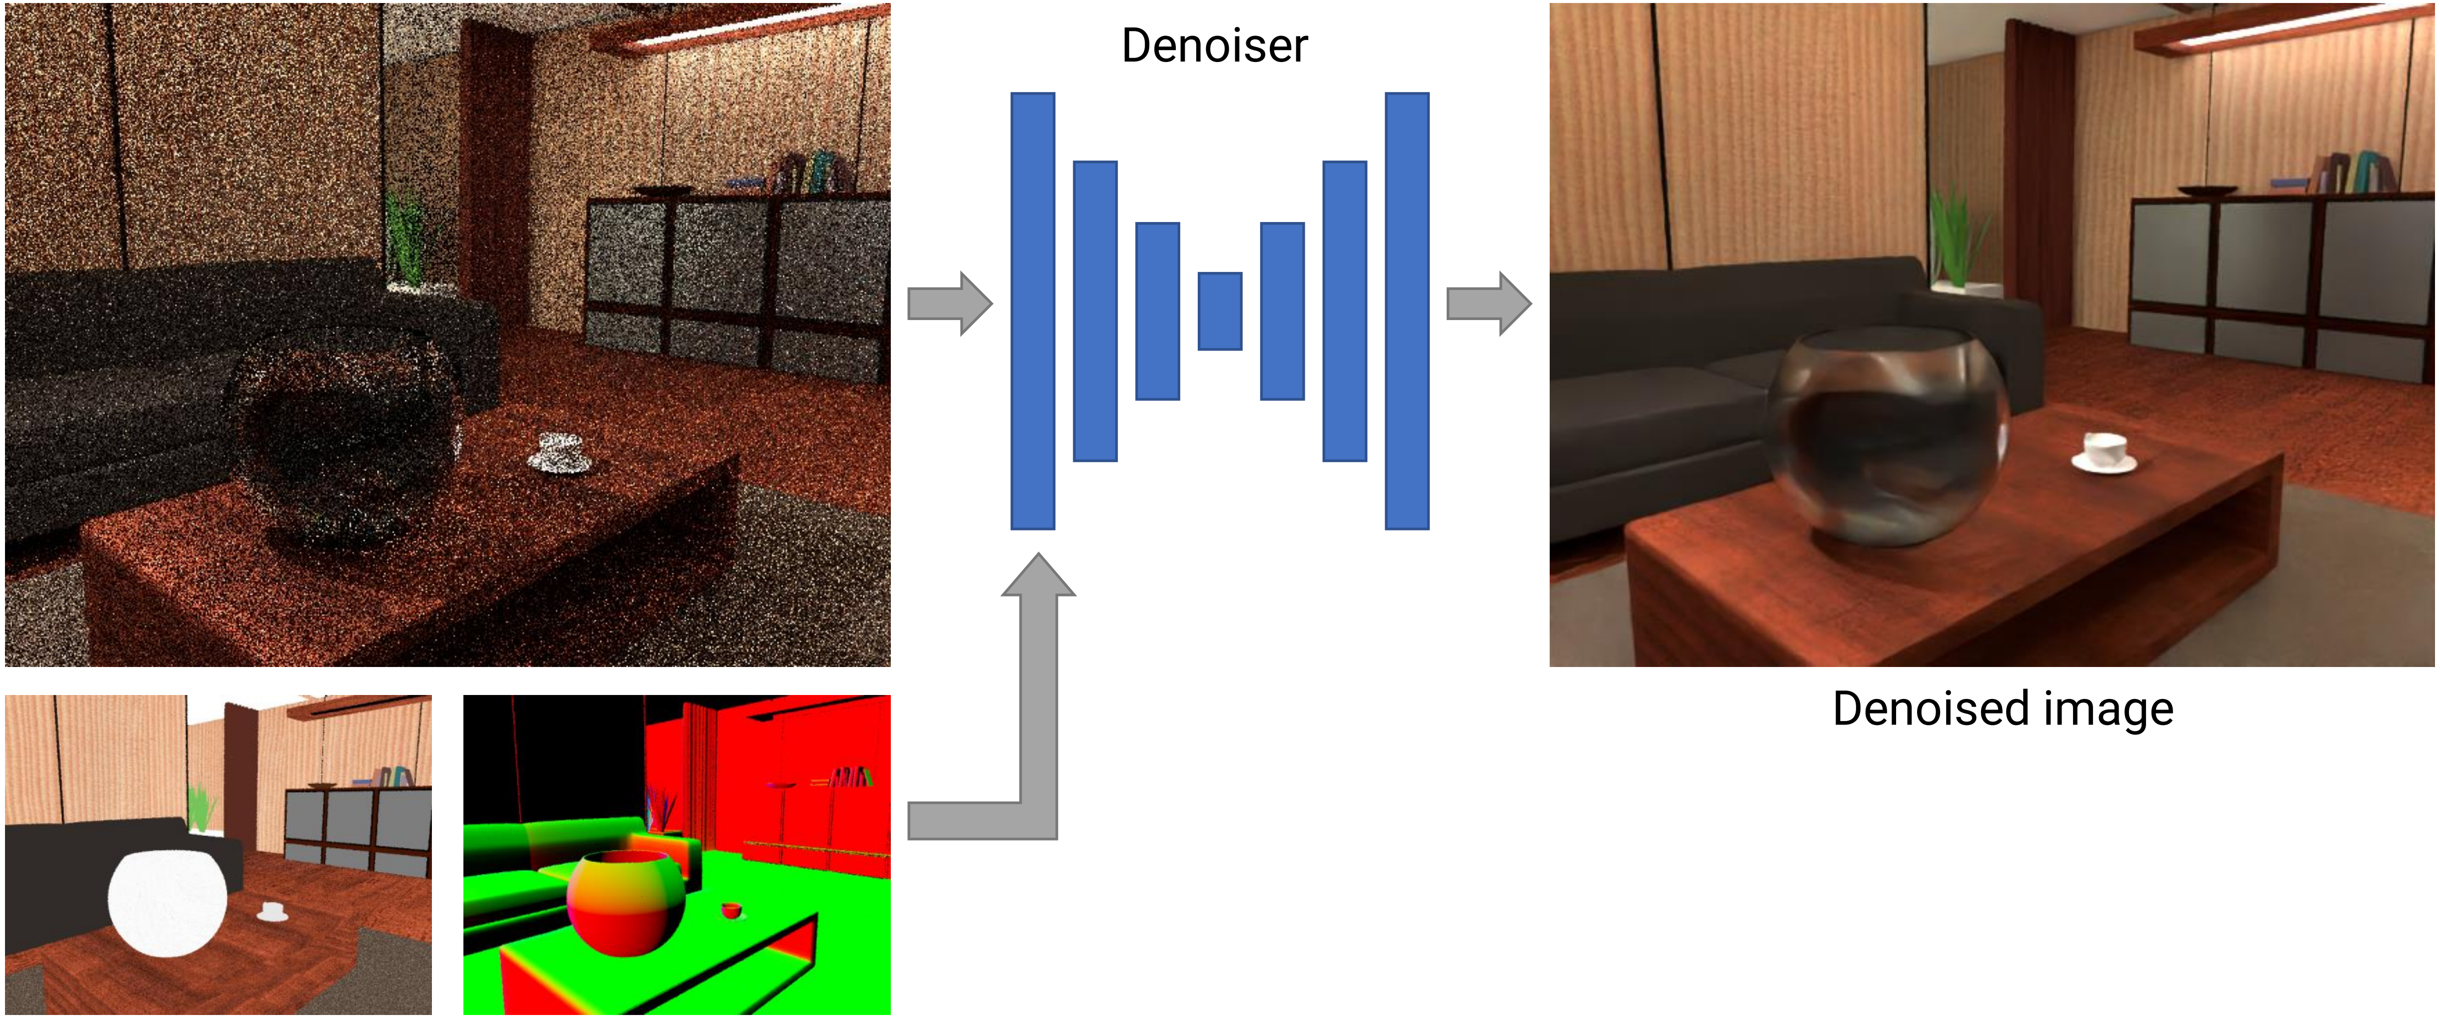

Differentiable Rendering

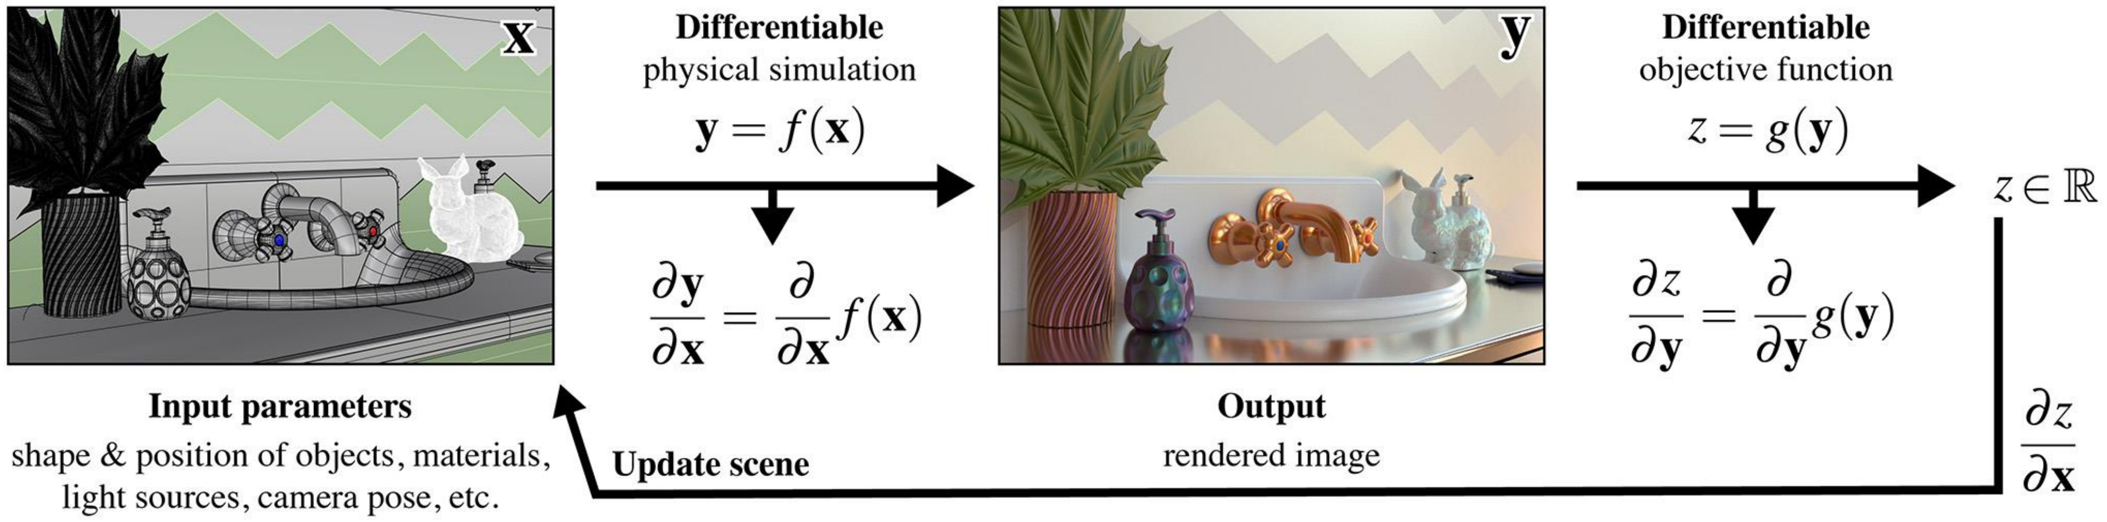

Beyond this course

How and where can you apply what you will learn?

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).

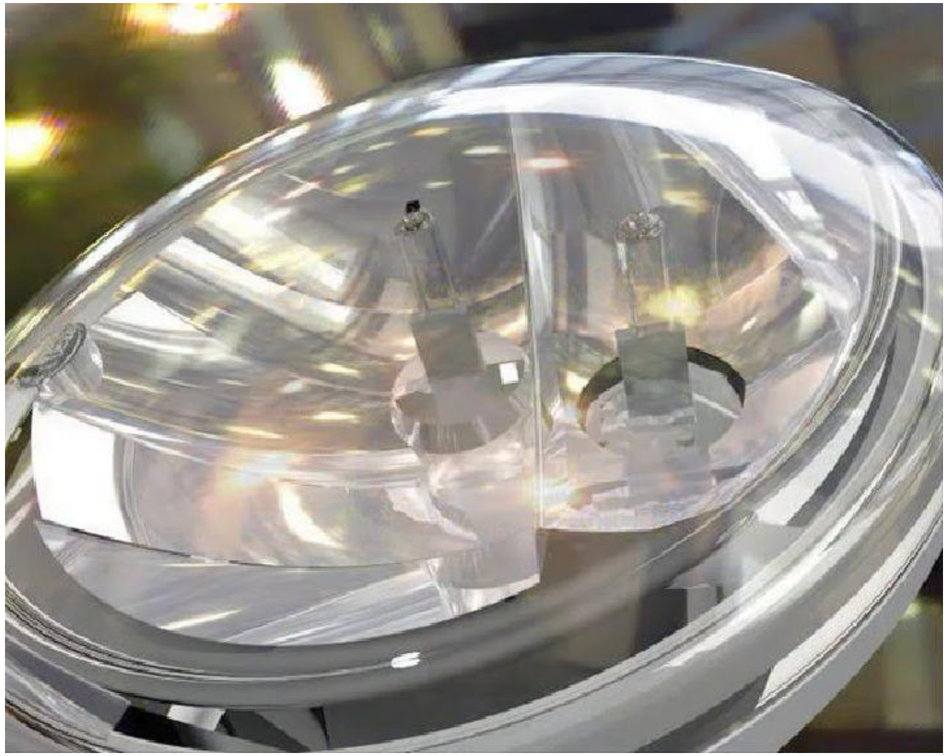

Instant Global Illumination

- 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.

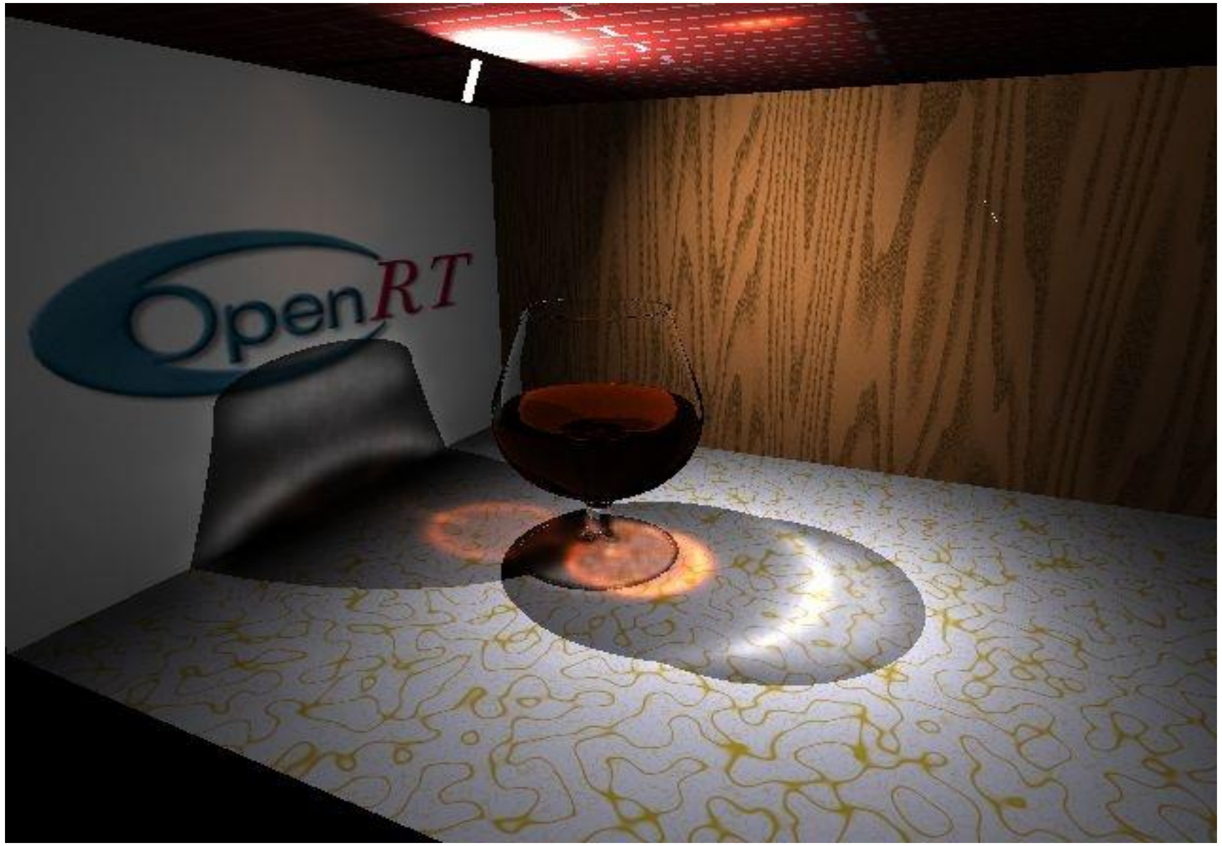

Photon Mapping

- 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).

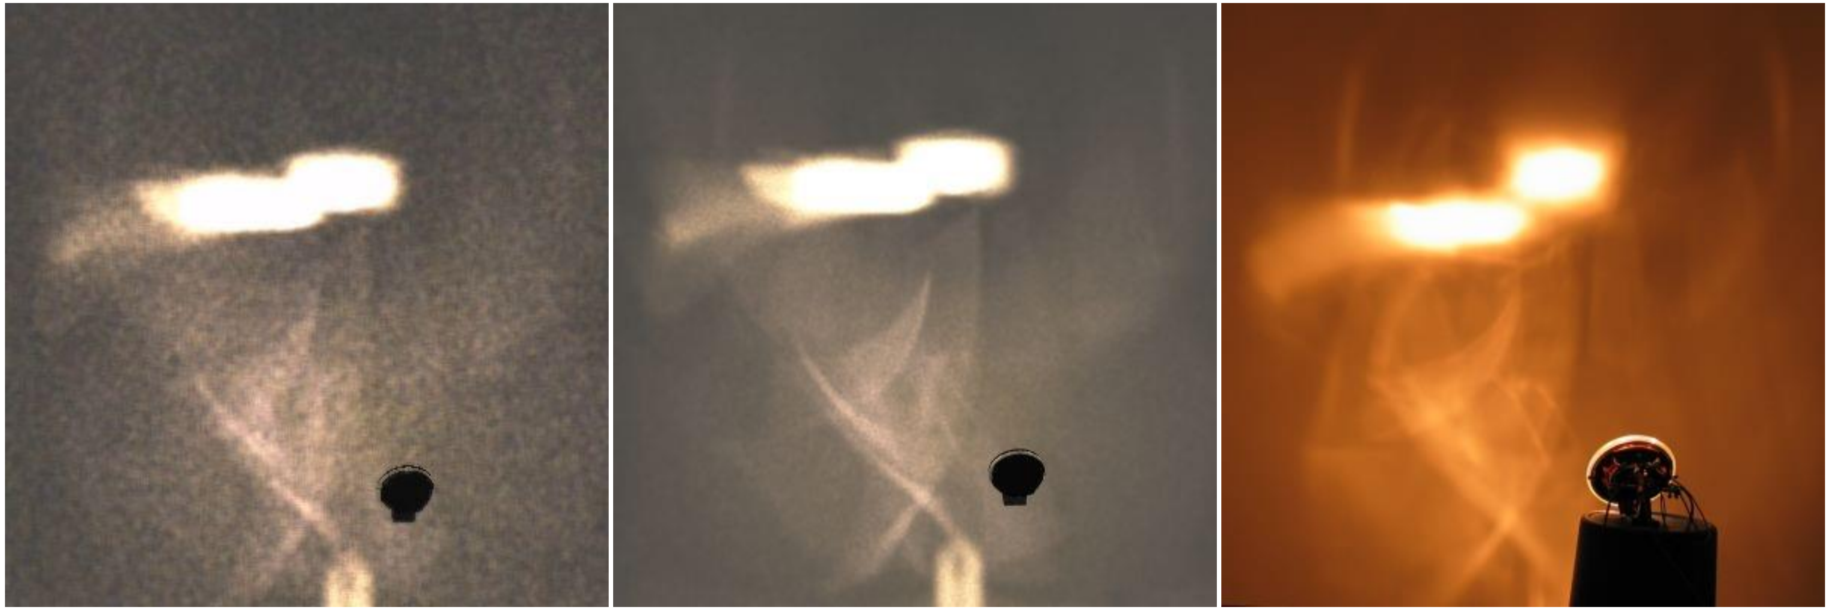

Light Transport Simulation

- 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.

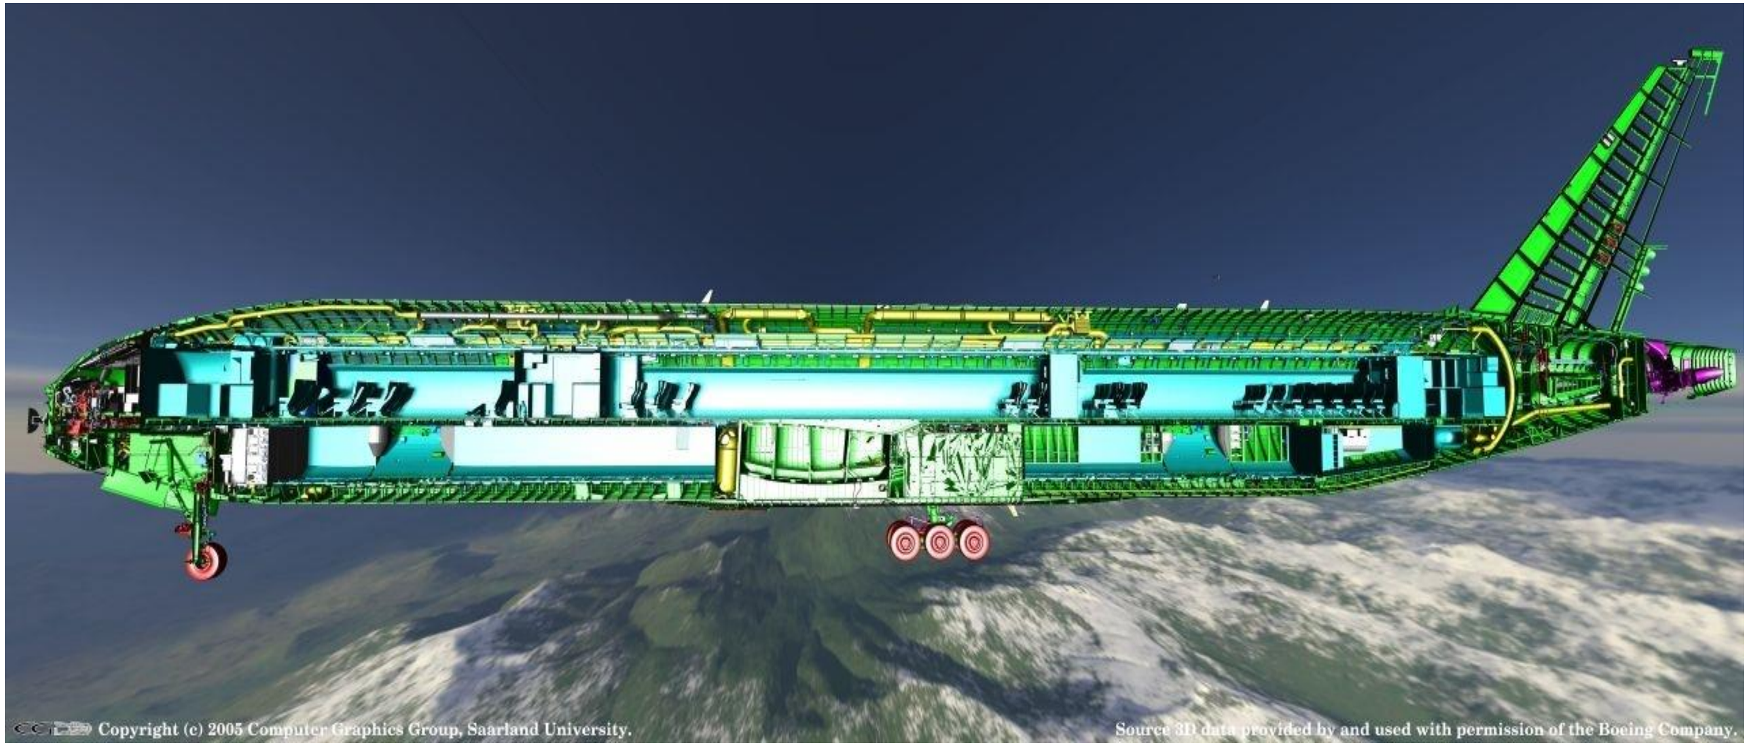

Massive Models

- Visualization of large outdoor scenes (300x300m2) with 365k plants and several billion triangles.

Massive Models

- Much larger outdoor scene (80x80 km²) with realistic lighting and full vegetation ($90*10^{12}$ triangles)

High-Performance Simulation

- Advanced rendering techniques in games

Physically-Based Image Synthesis with Real-Time Ray Tracing

Technical Oscar
Feb 2021

Custom Ray Tracing Processor [Siggraph'05]

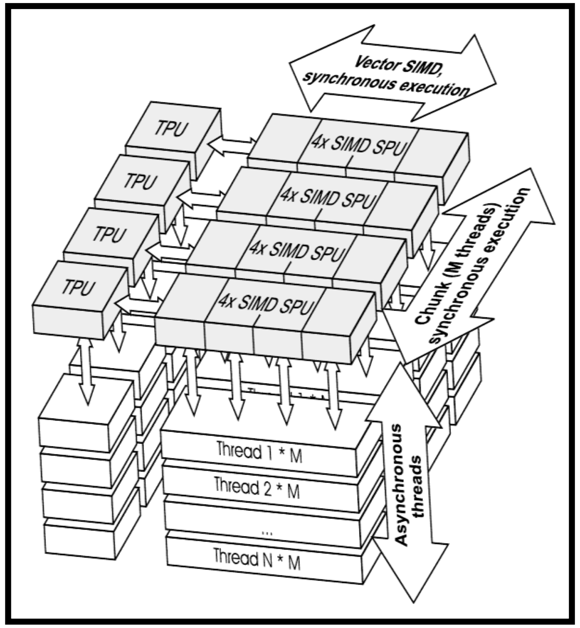

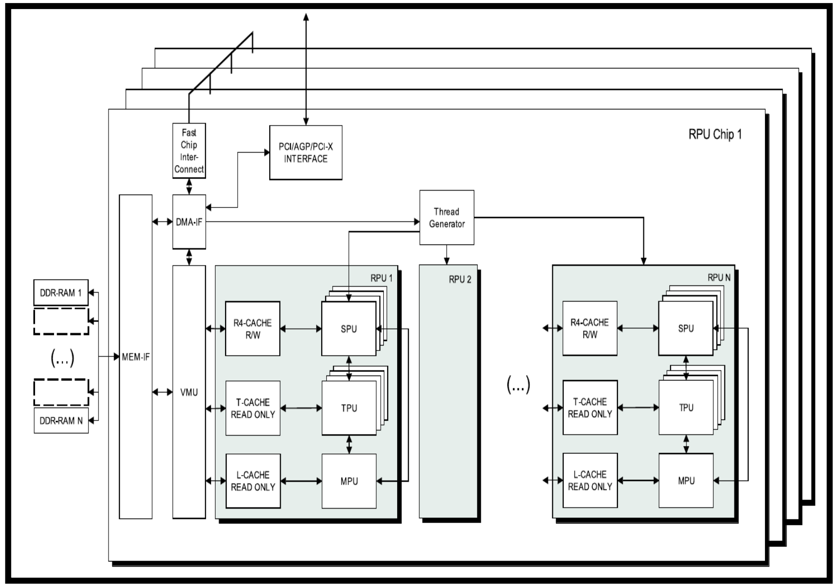

AnyDSL Compiler Framework

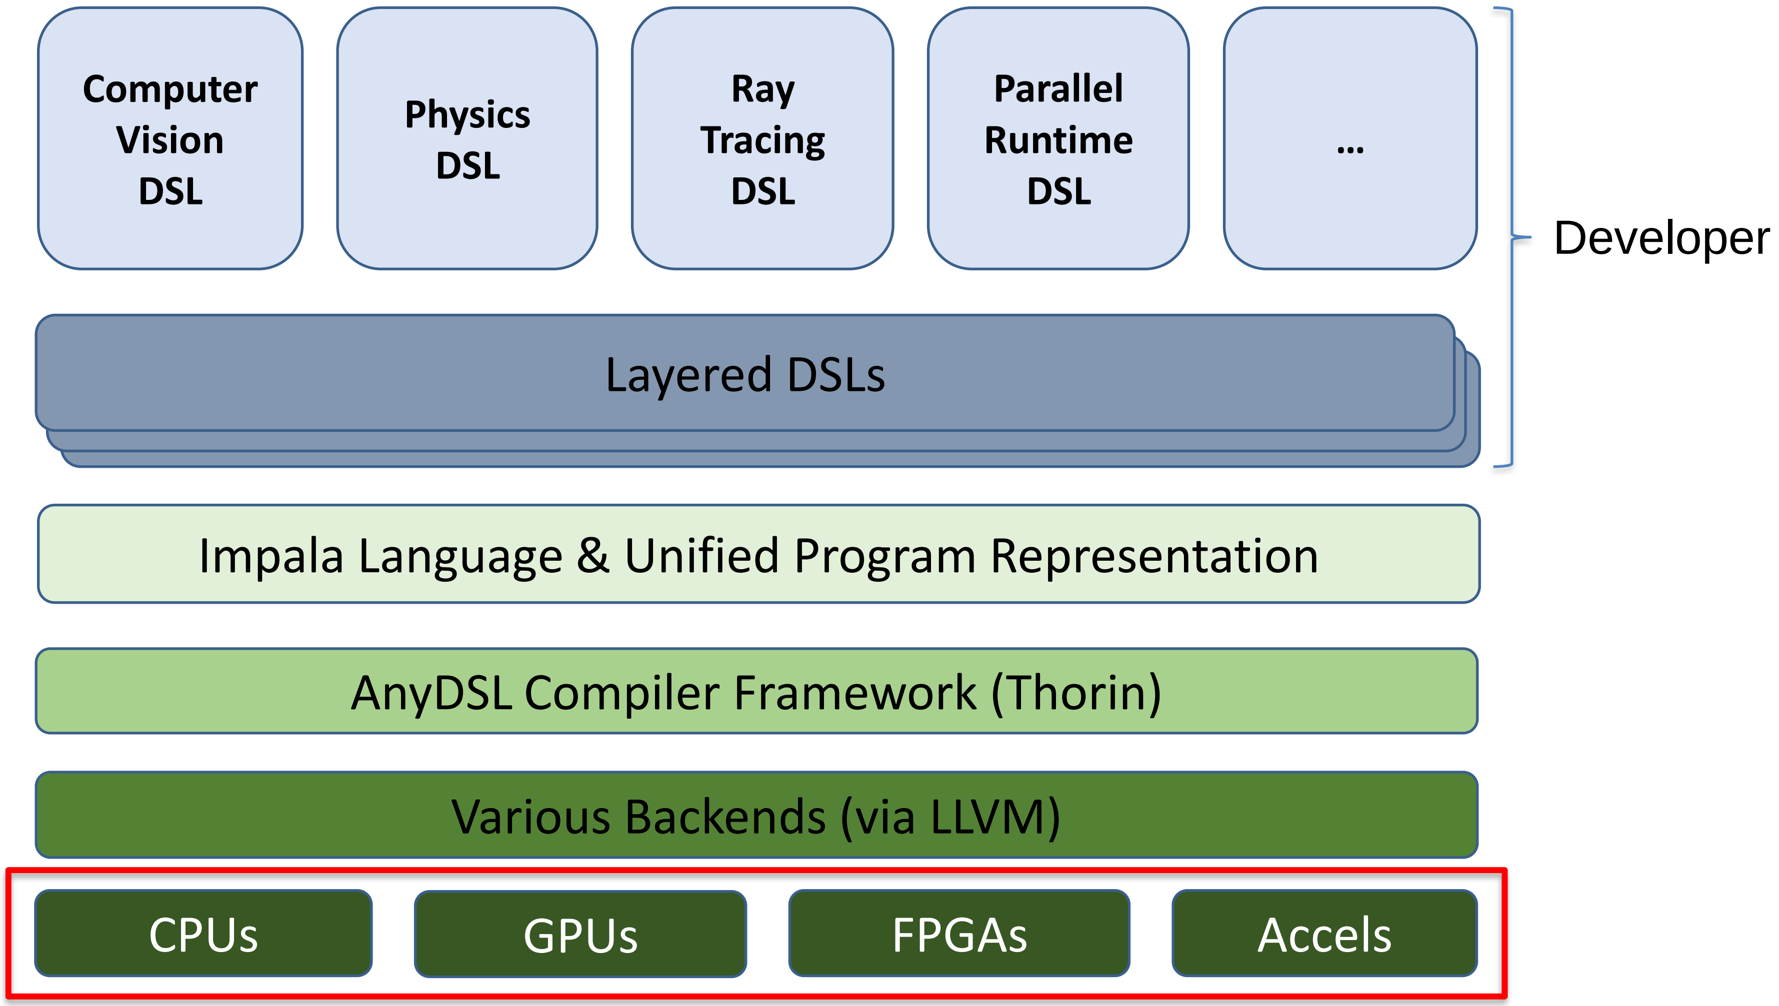

Importance Caching

- Iliyan Georgiev, et al. [Eurographics 2012]

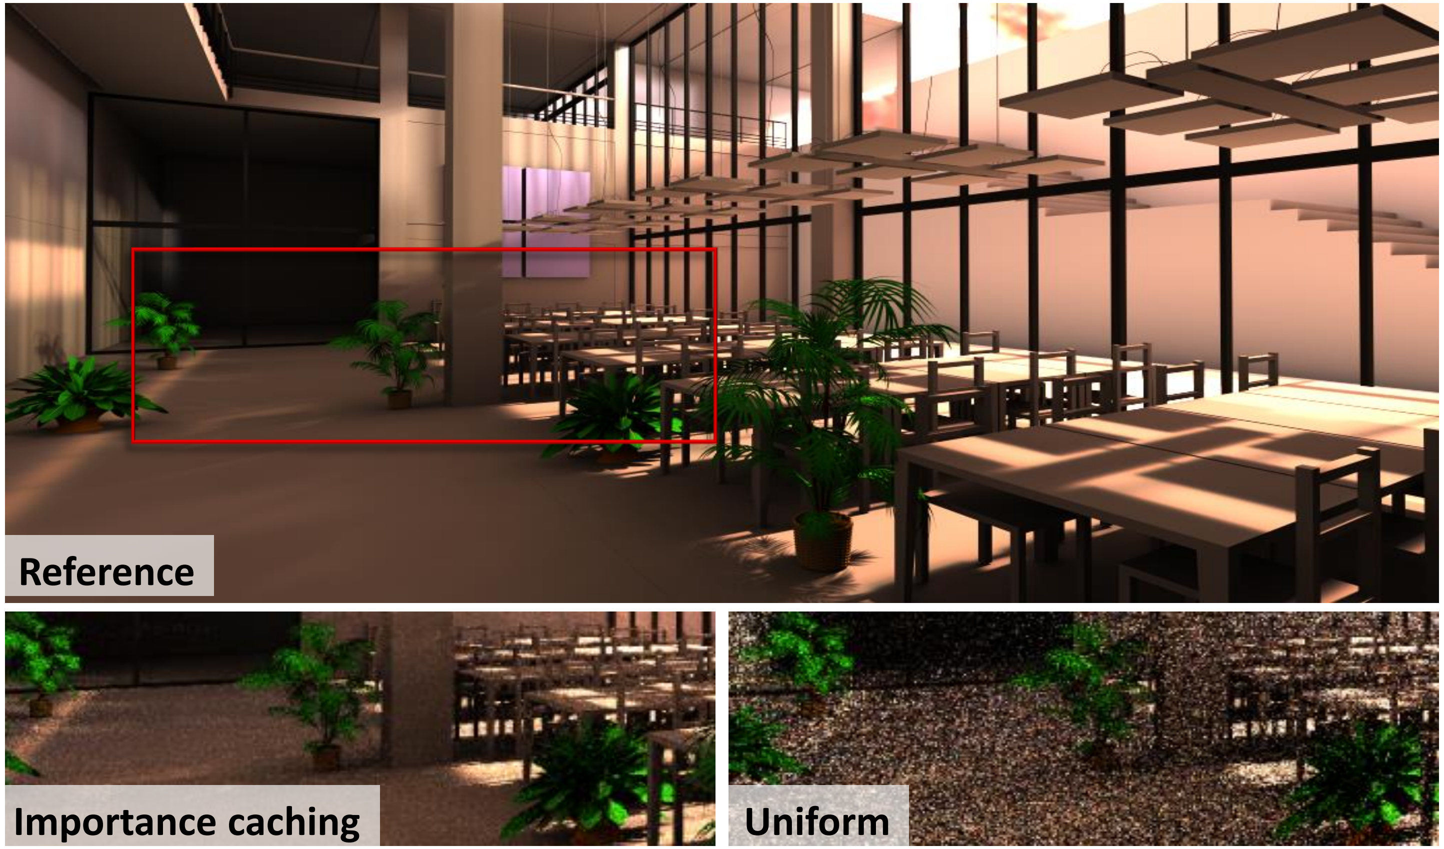

Monte-Carlo vs Density Estimation

- Vertex Connection & Merging, Ilijan Georgiev [SiggraphAsia´12]
  - Formulating Density Estimation algorithms as a Monte-Carlo (MC) techniques

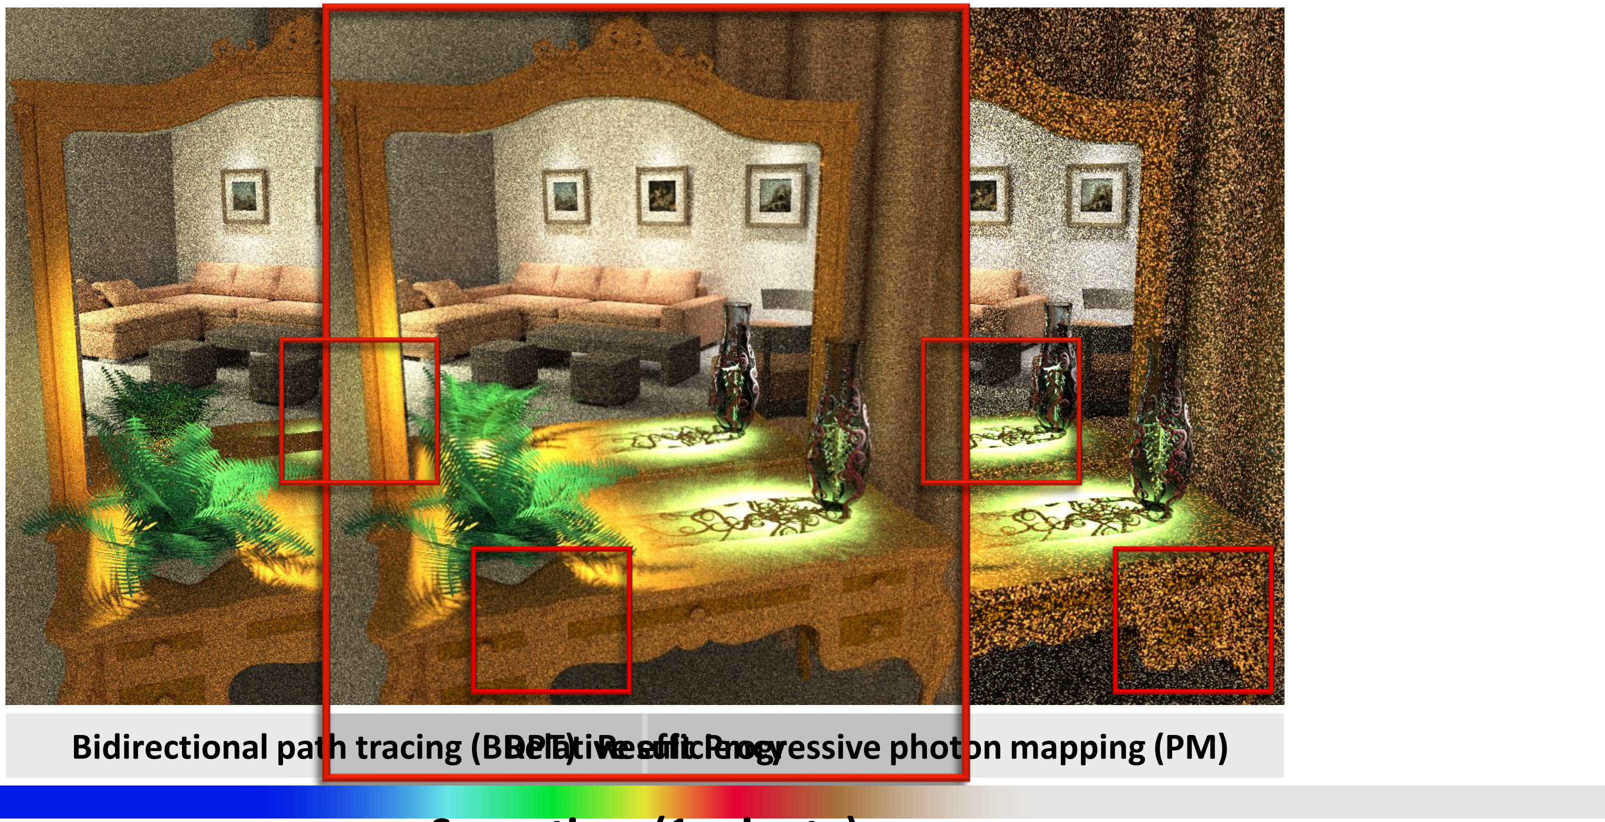

Same time (1 minute)

A Quick Glance at (Some of) Our Current Research

- Goal: General, robust, and efficient rendering algorithms
- “One algorithm to render them all”

- Methodology: Adapt the algorithm to the scene based on statistics from initial samples
  - Learn better sample distributions
  - Optimize parameter values and sample counts
  - Adapt weighting functions and combinations

Motivation

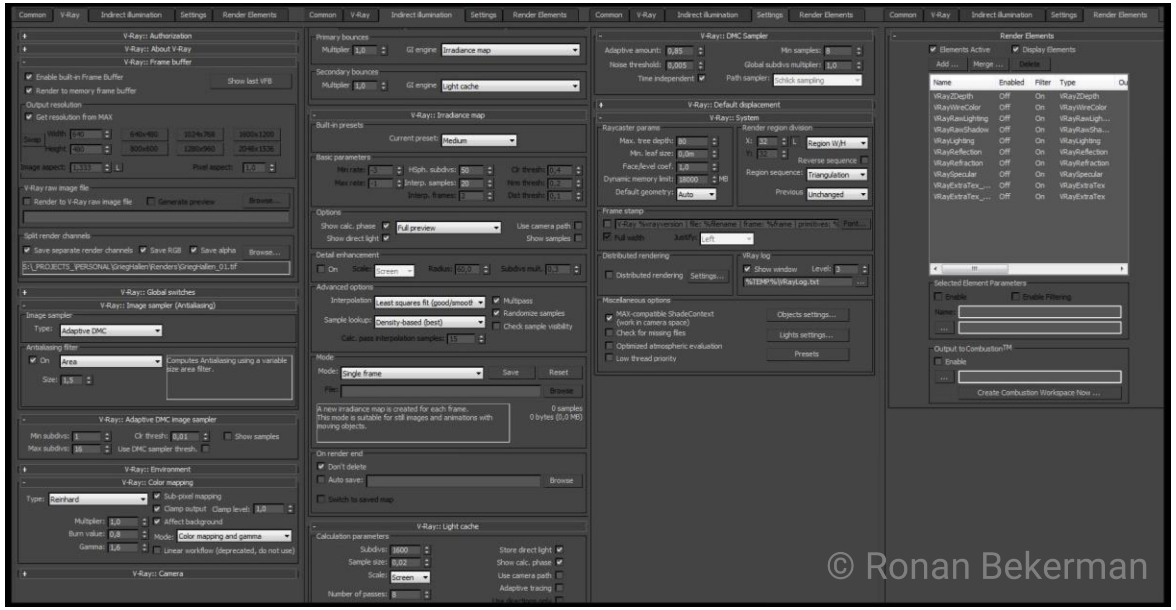

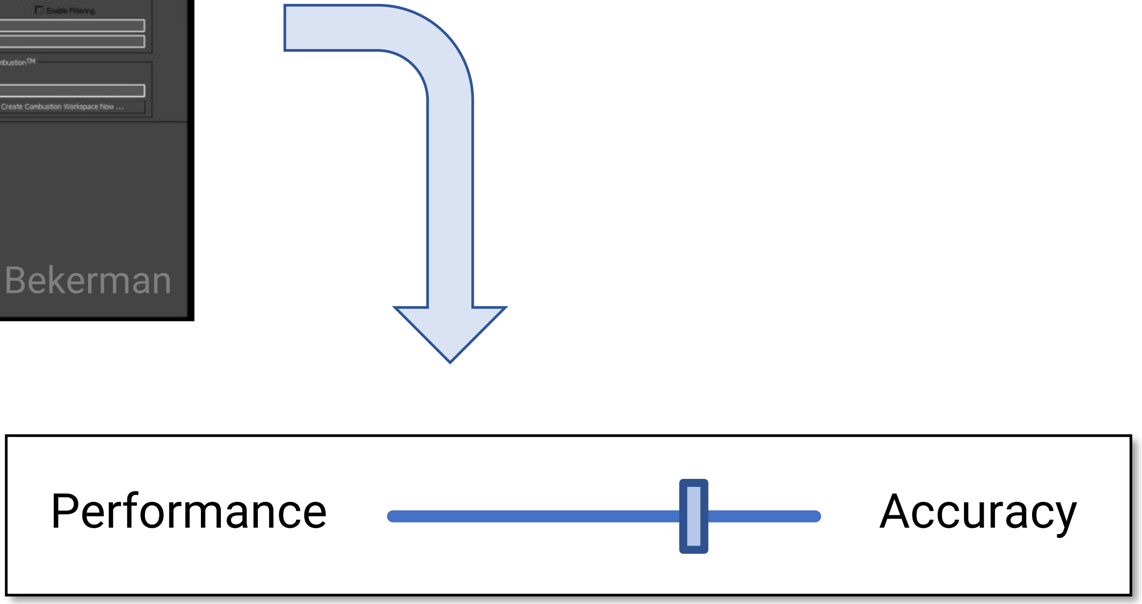

Adapting Parameters and Sample Counts

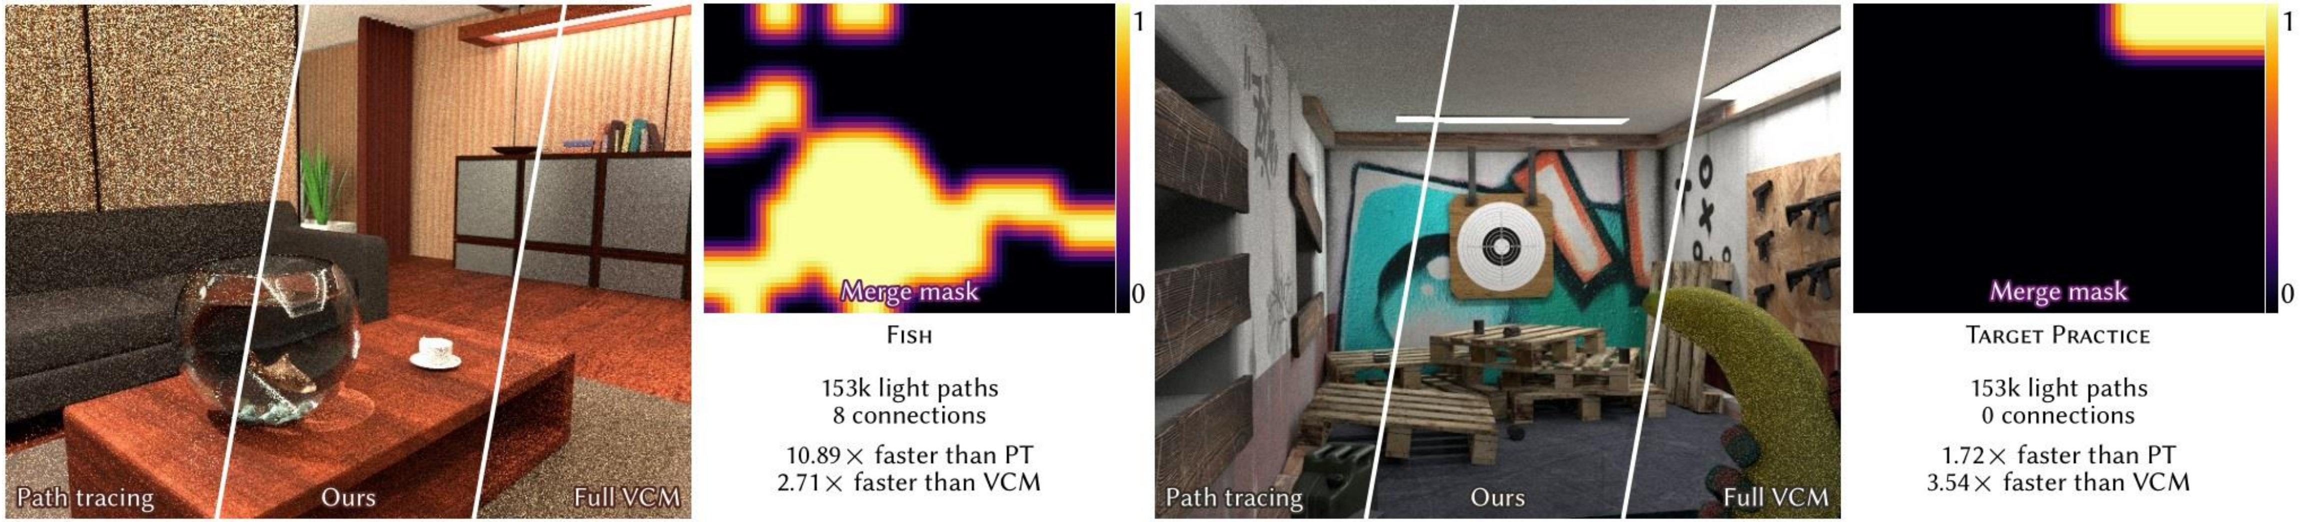

Grittmann et al. – Efficiency-aware multiple importance sampling
SIGGRAPH 2022

Lightweight Bidirectional Methods

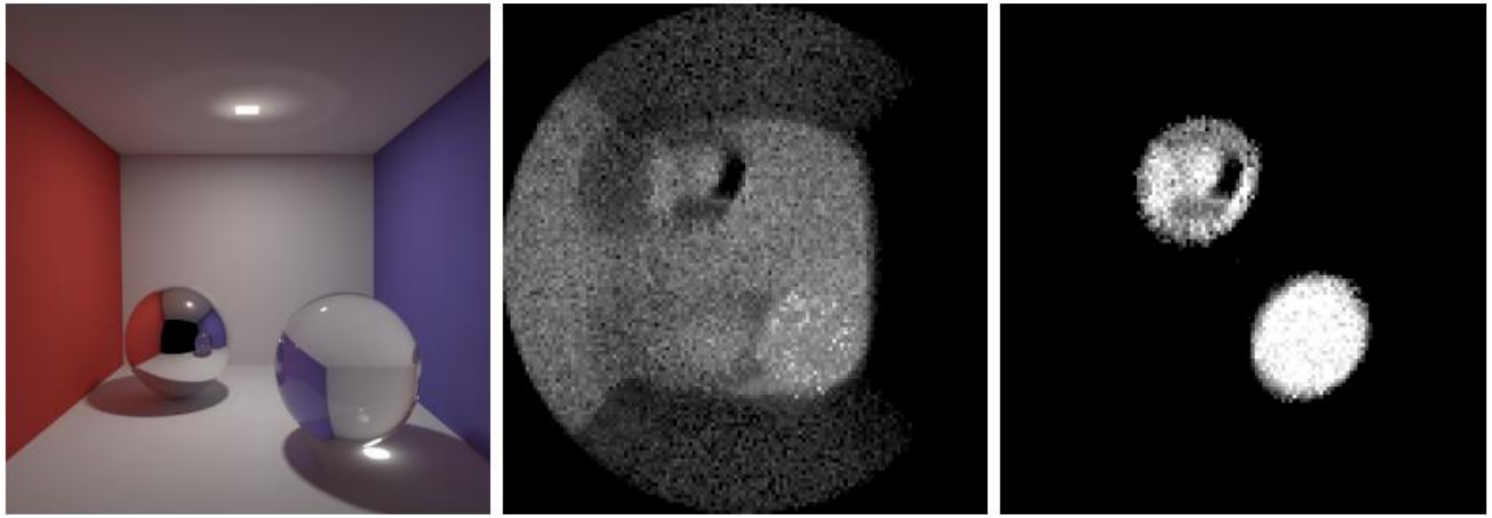

Grittmann et al. – Efficient caustic rendering with lightweight photon mapping
EGSR 2018

What Should Path Guiding Learn?

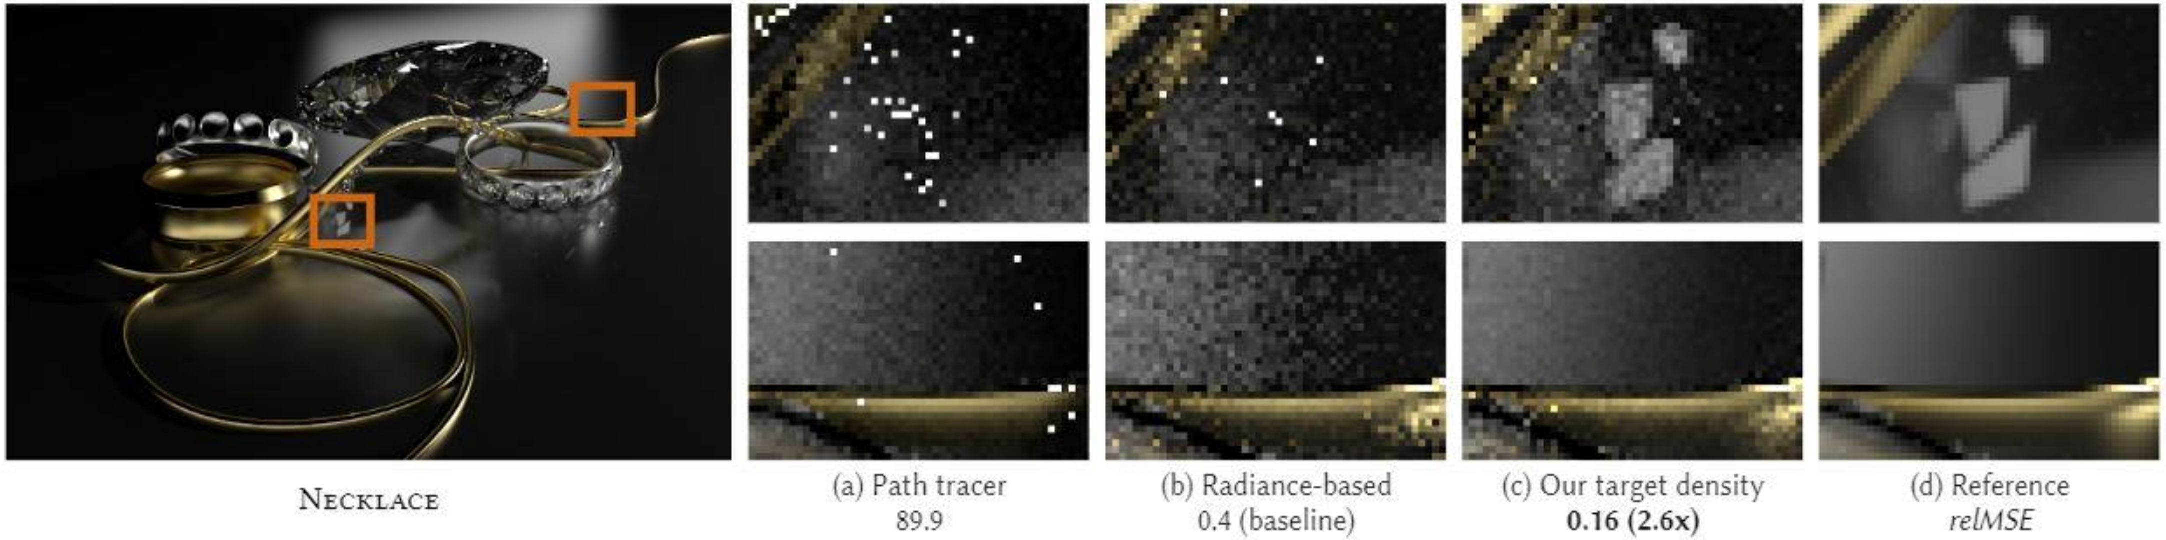

Rath et al. – Variance-aware path guiding. SIGGRAPH 2020

Path Termination and Splitting

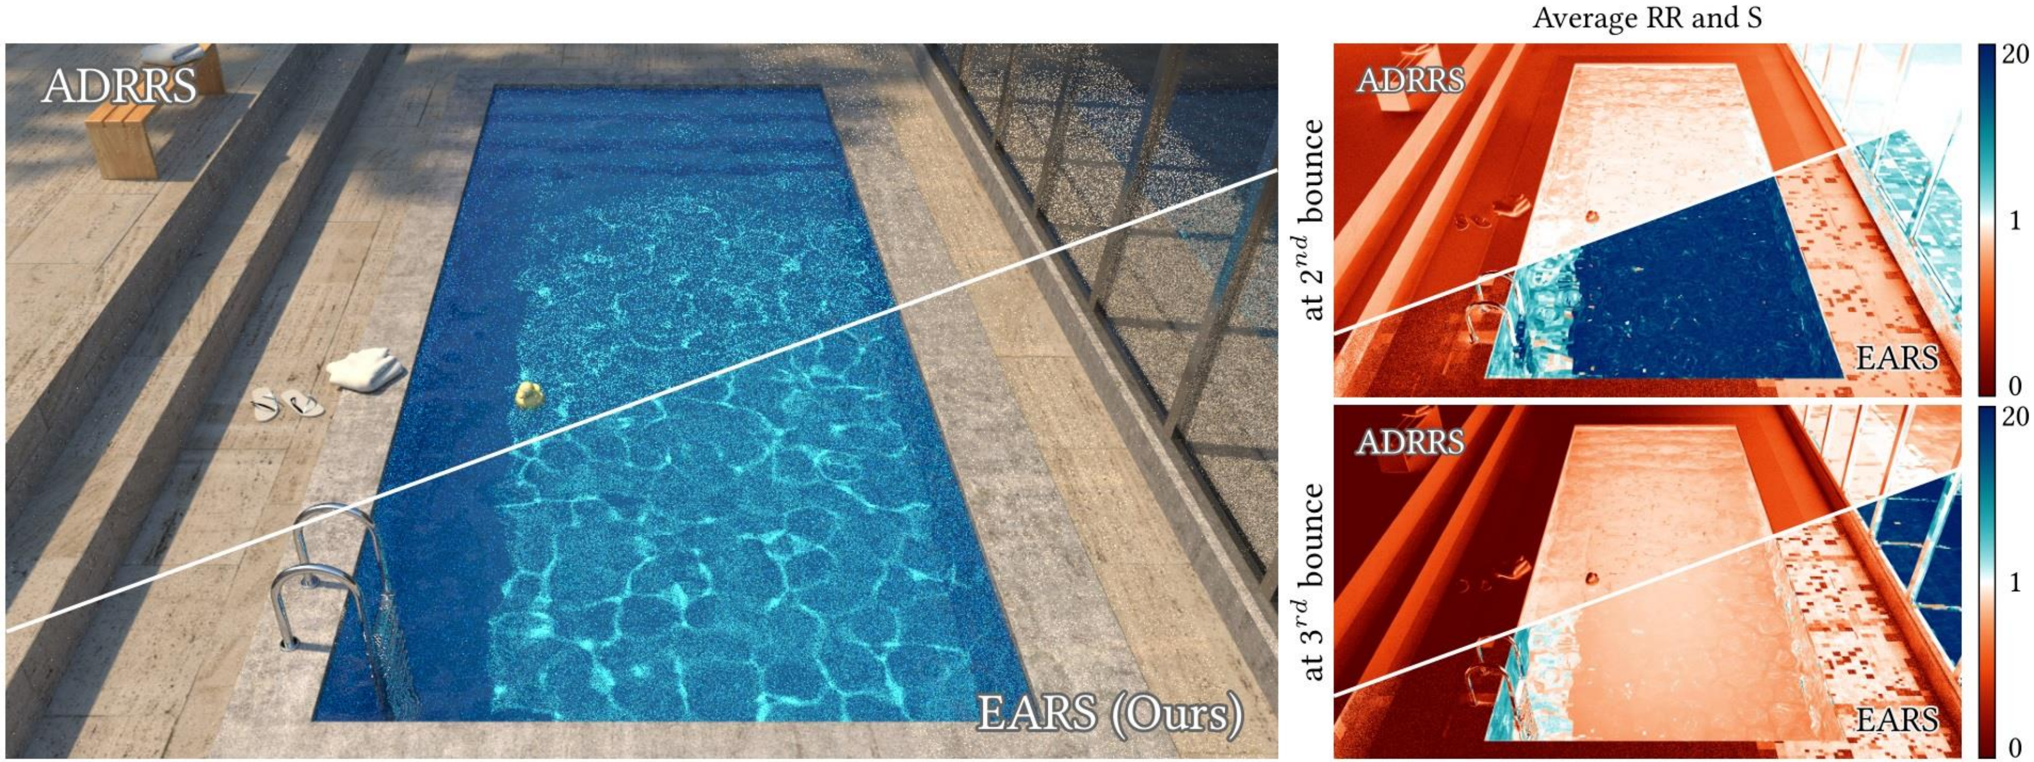

Rath et al. – EARS: Efficiency-aware Russian roulette and splitting
SIGGRAPH 2022

Optimal MIS

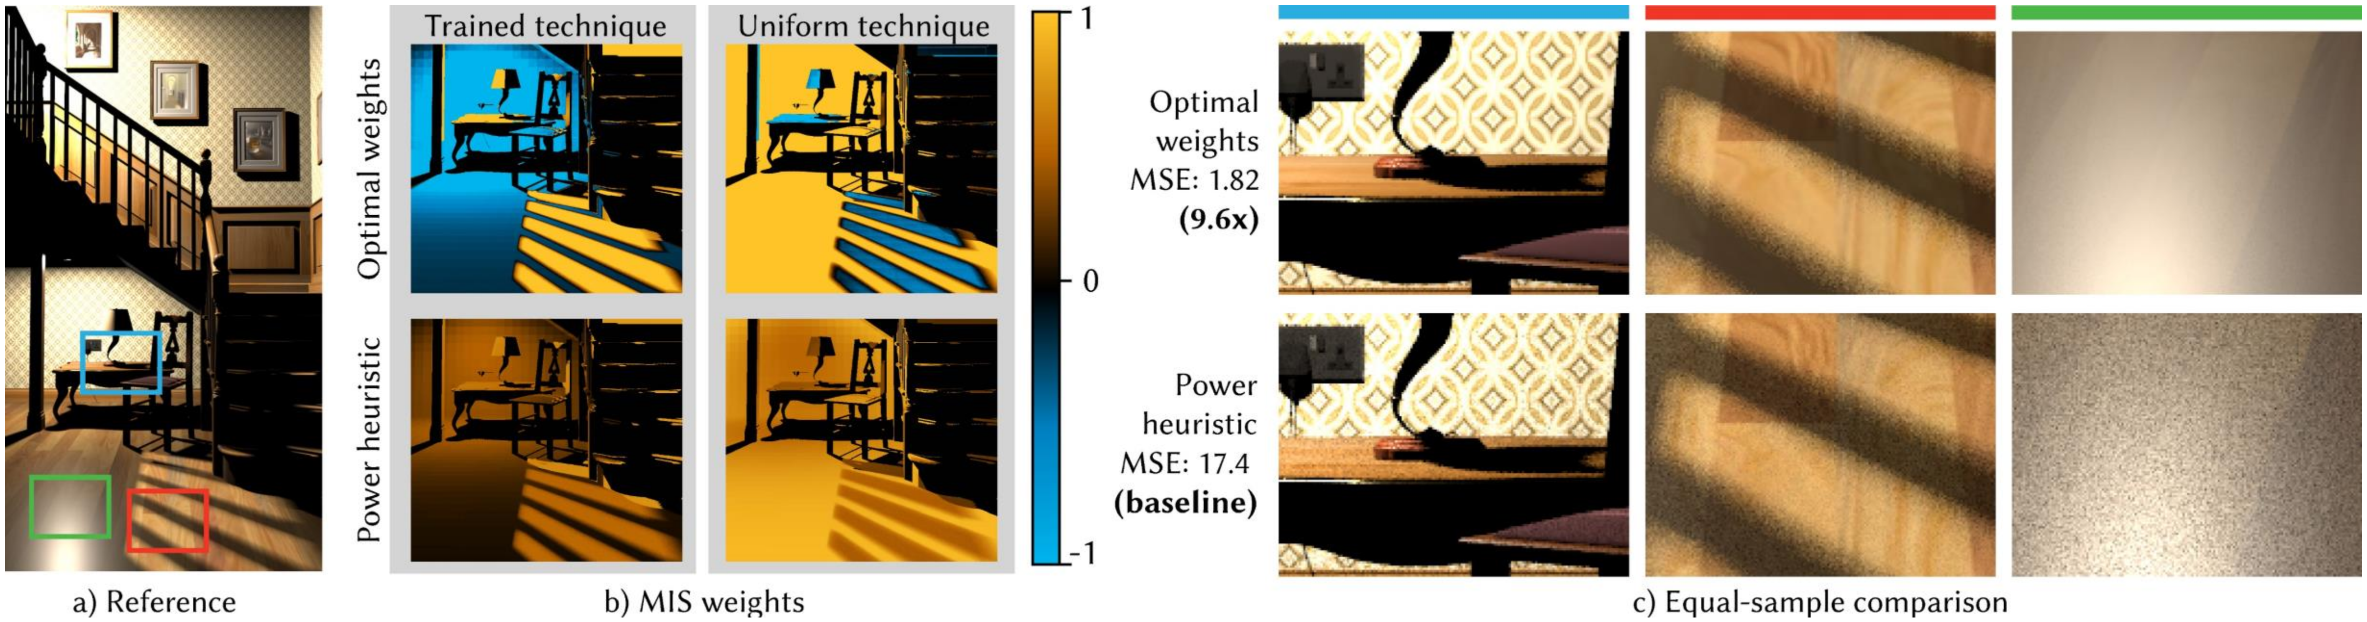

Kondapaneni et al. – Optimal multiple importance sampling
SIGGRAPH 2019

Fixing MIS for Bidirectional Methods

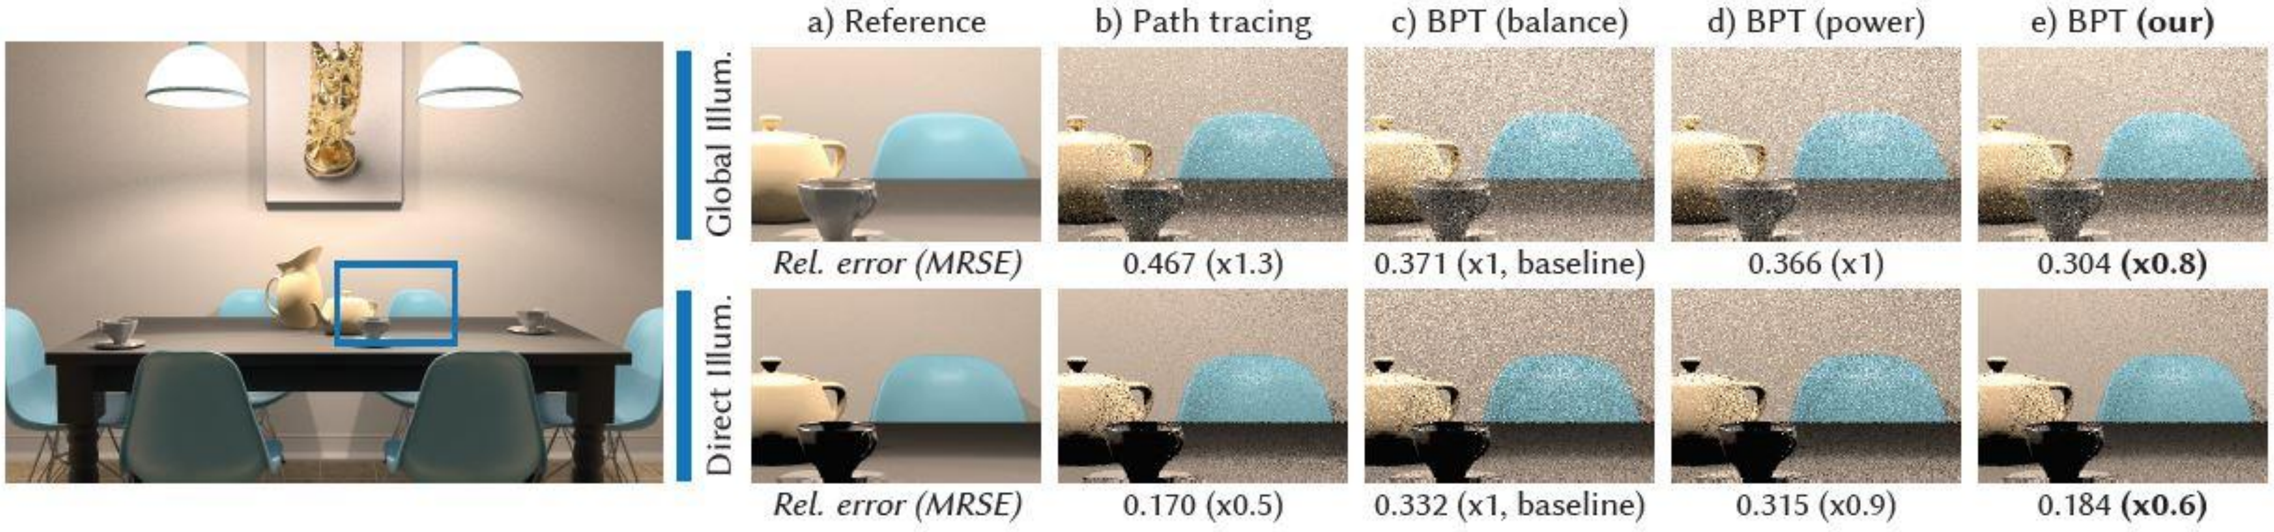

Grittmann et al. – Variance-aware multiple importance sampling
SIGGRAPH Asia 2019

Fixing MIS for Bidirectional Methods – Part II

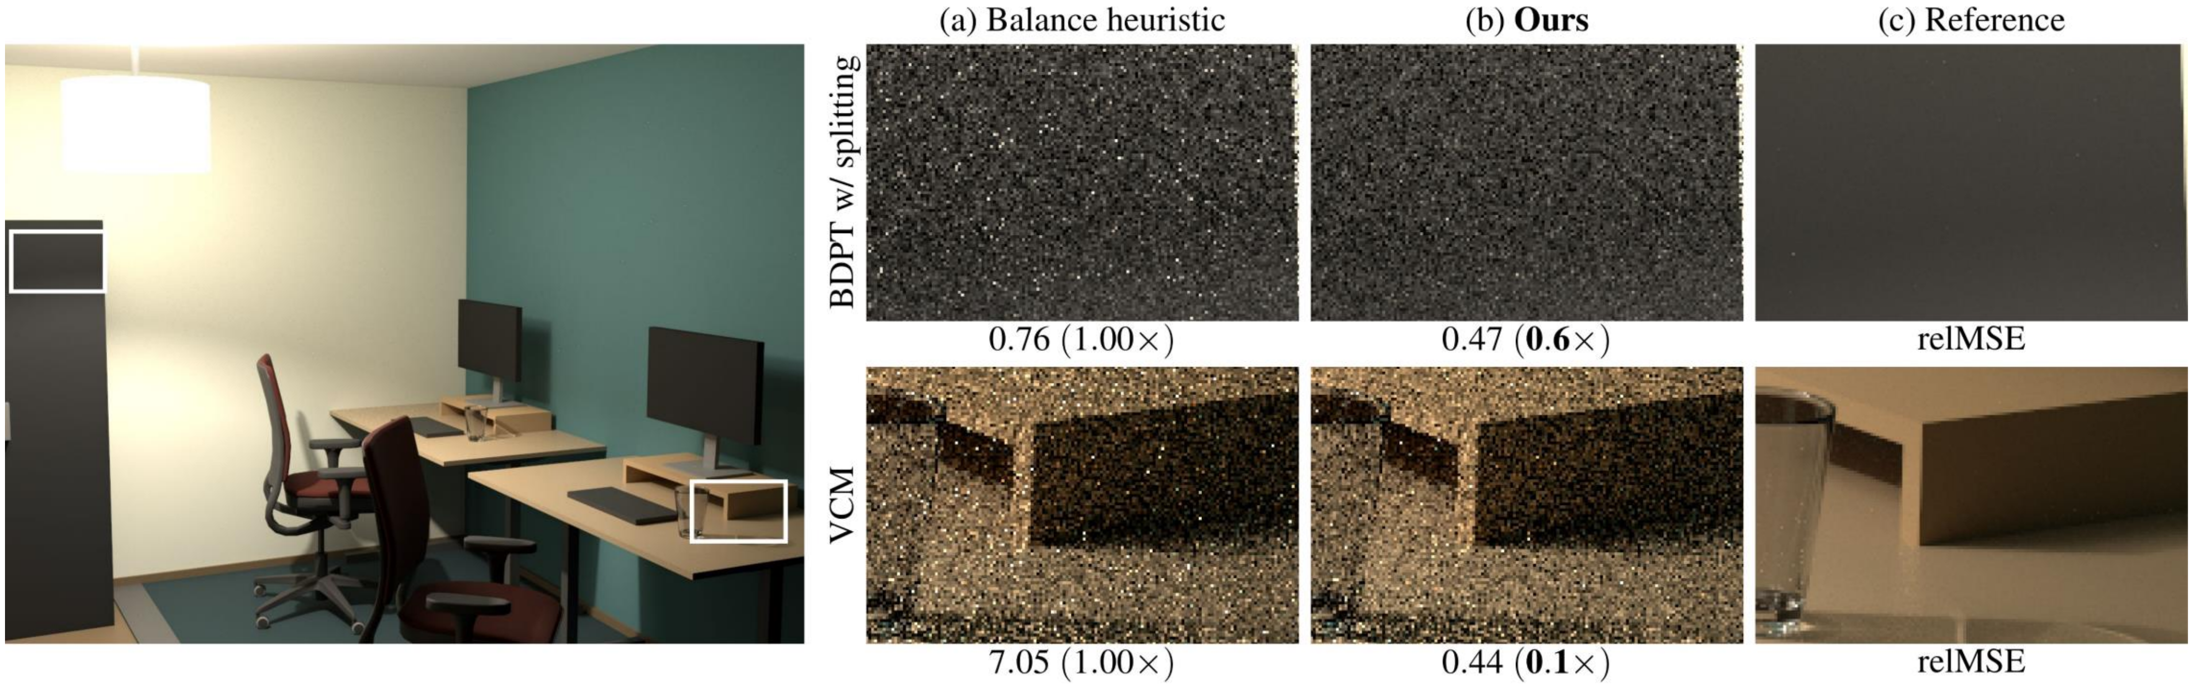

Grittmann et al. – Correlation-aware multiple importance sampling
Eurographics 2021

Identifying Guiding Targets not on Surfaces

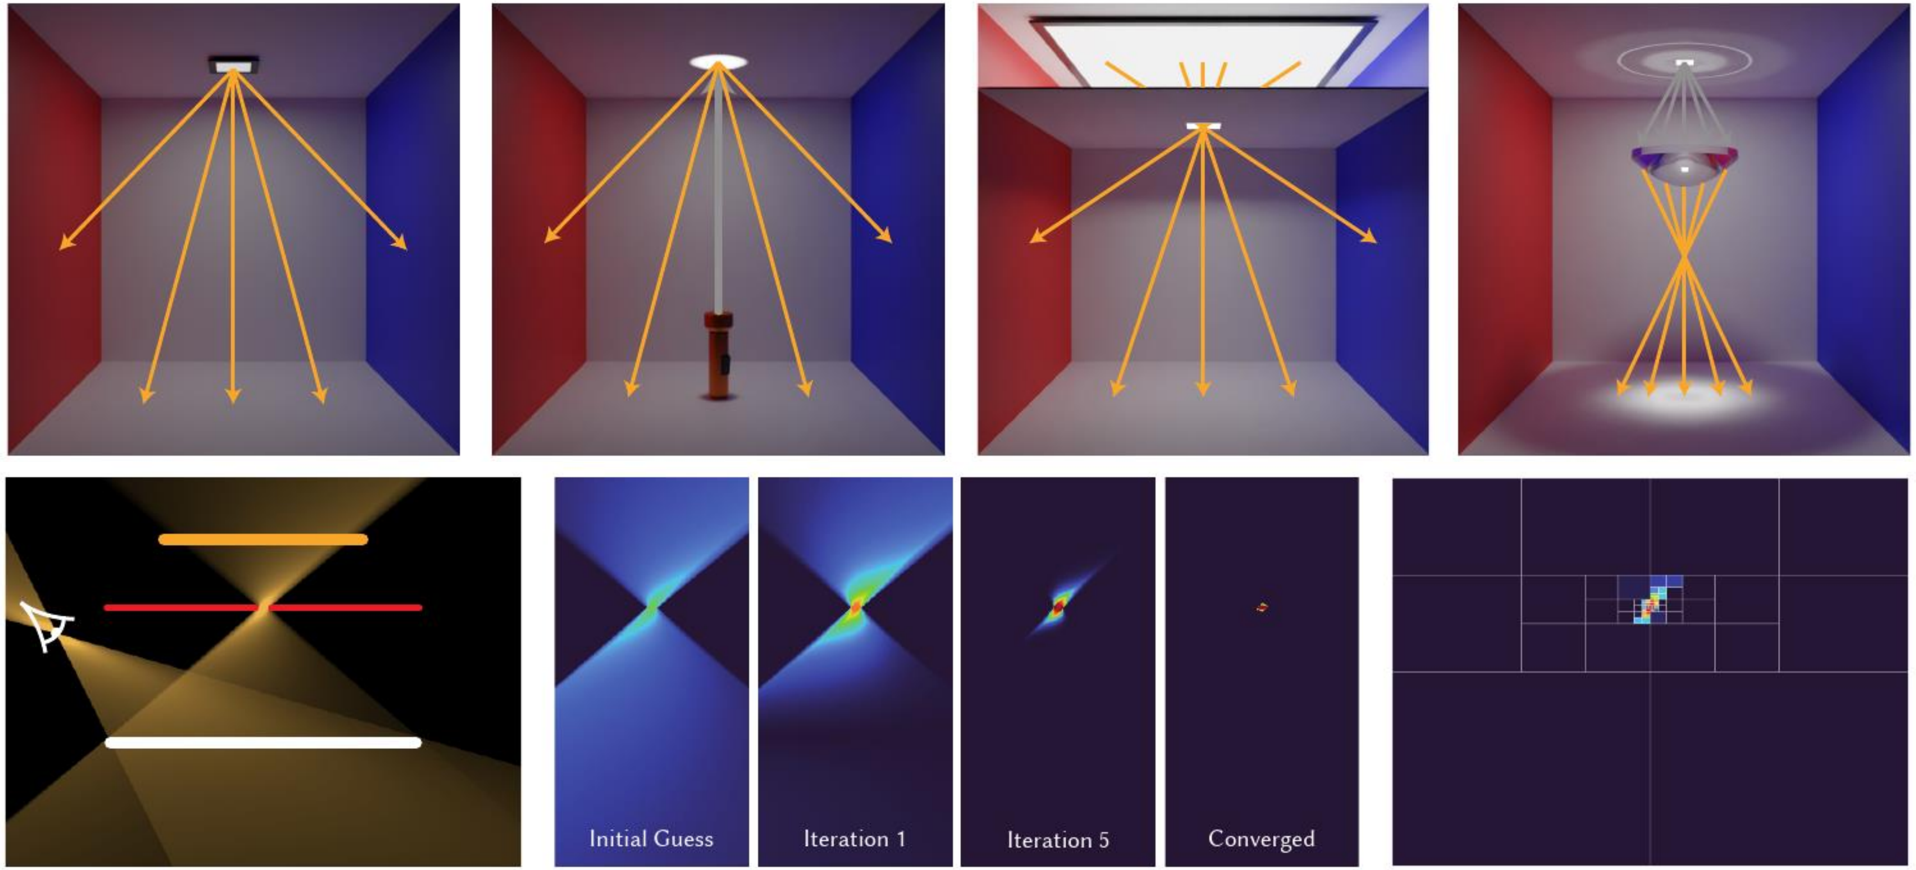

Rath et al. – Focal Path Guiding
Siggraph 2023

Learning Compact Scene Representations

Weier, et al. – Rendering with mixed geometric and neural representations.
Siggraph 2023 + 2024

Other Research From Saarbrücken

- Some more examples from my research group

DFKI-ASR: Agents and Simulated Reality

How to design AI systems that can provide guarantees and that humans can understand and trust?

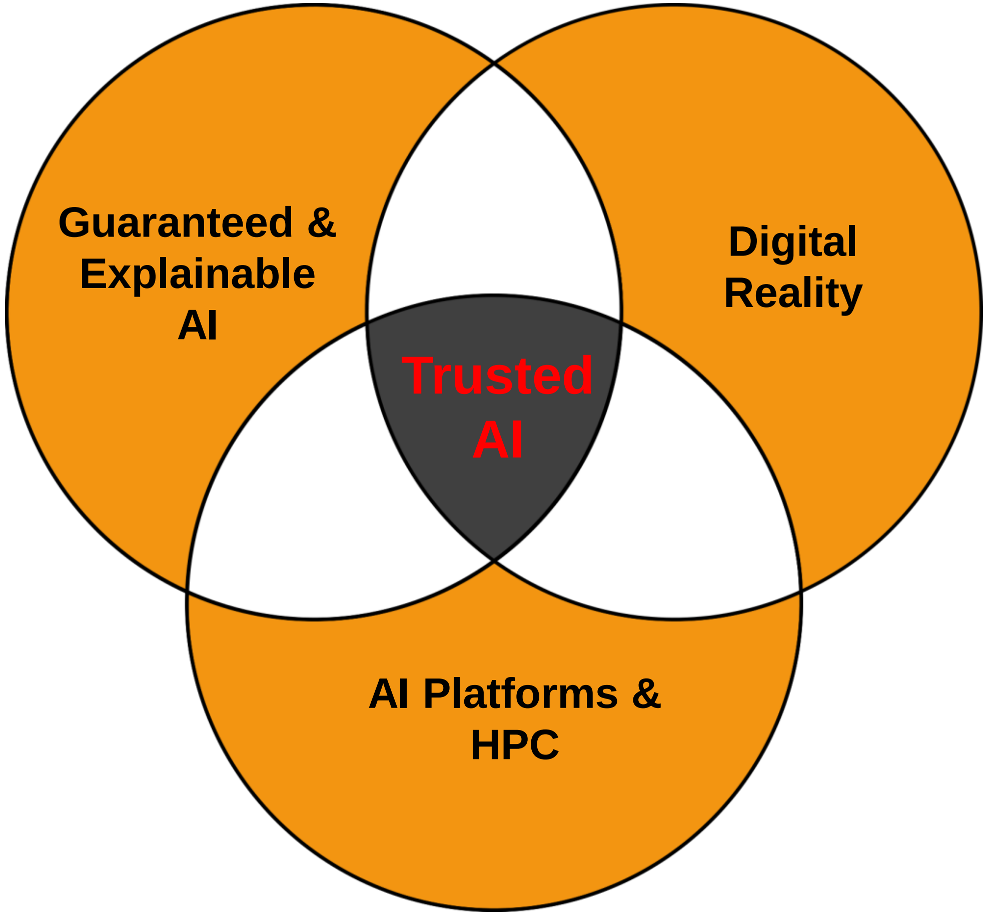

How can synthetic data from parametric models and simulations be used for *training, validating, and certifying AI systems?*

How can AI-systems be realized technically in a *reliable and efficient way?*

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)

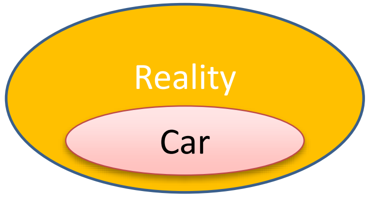

Digital Reality

- 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

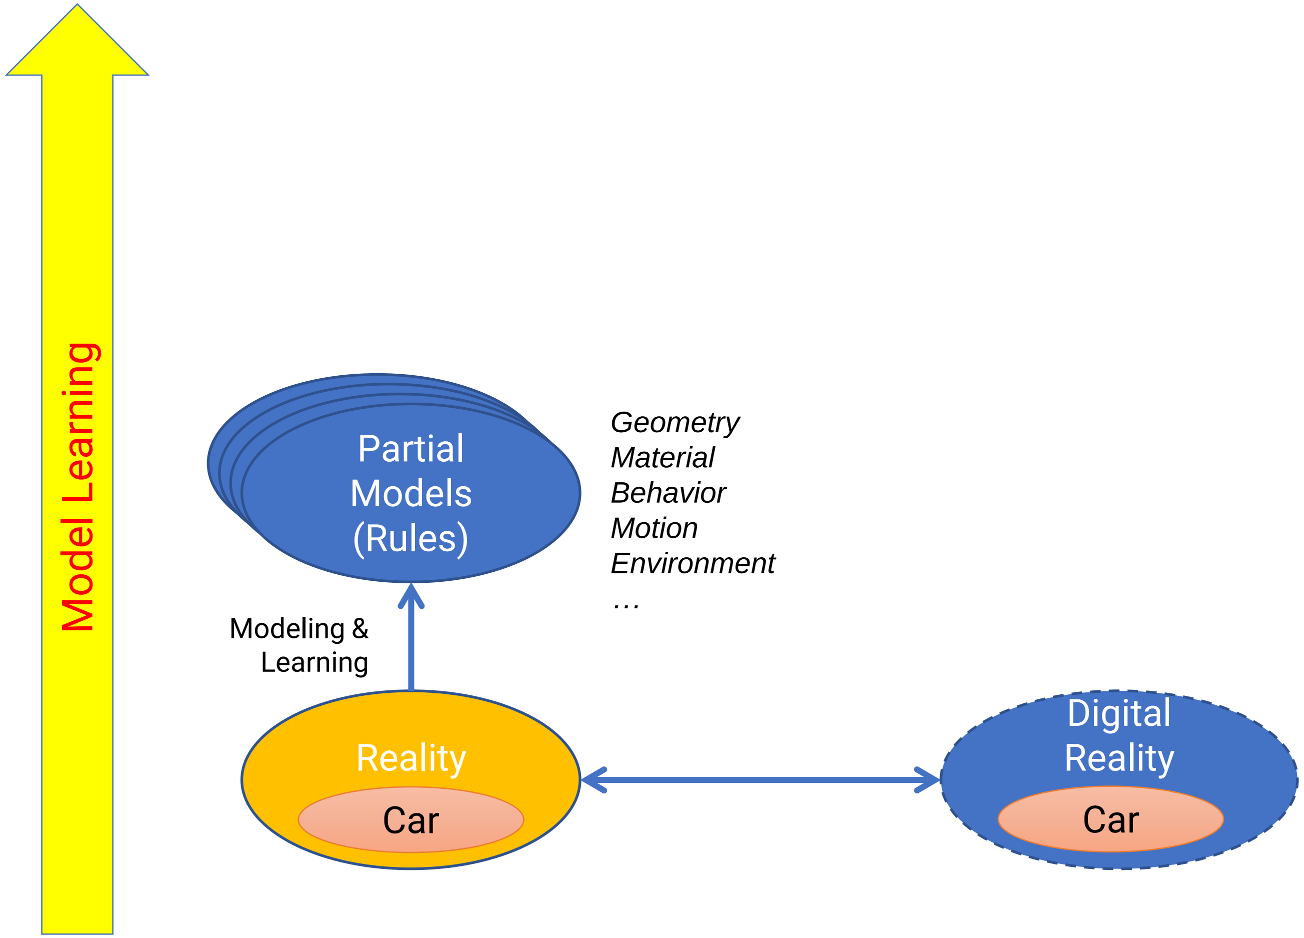

Digital Reality: AI to Certify AI

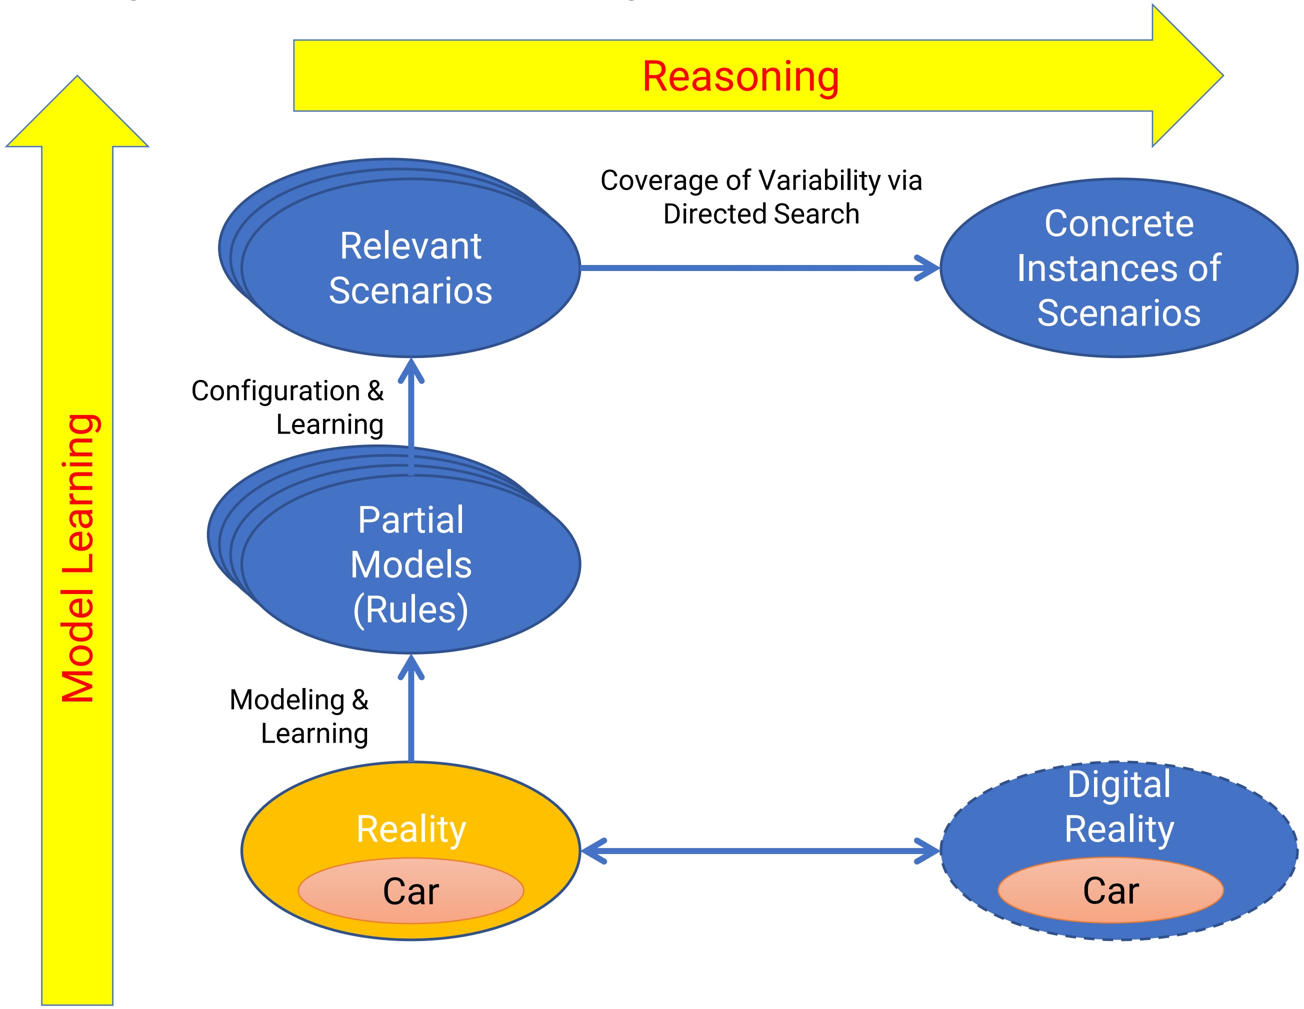

Digital Reality: AI to Certify AI

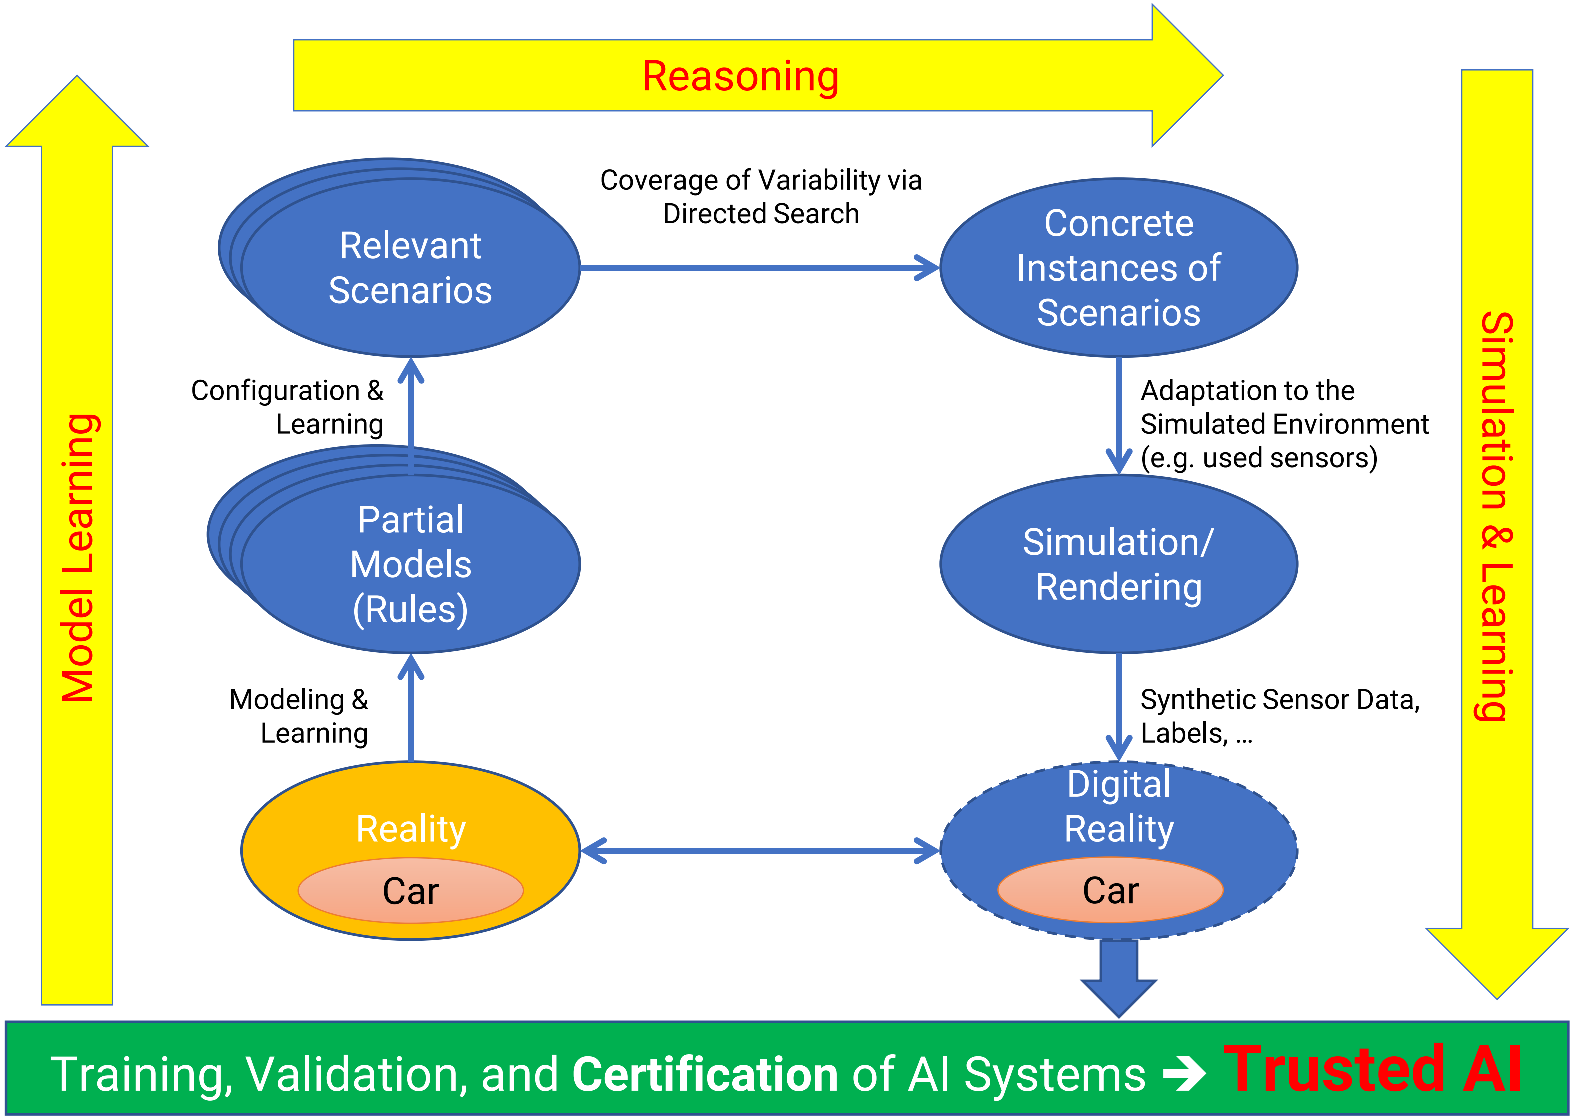

Digital Reality: AI to Certify AI

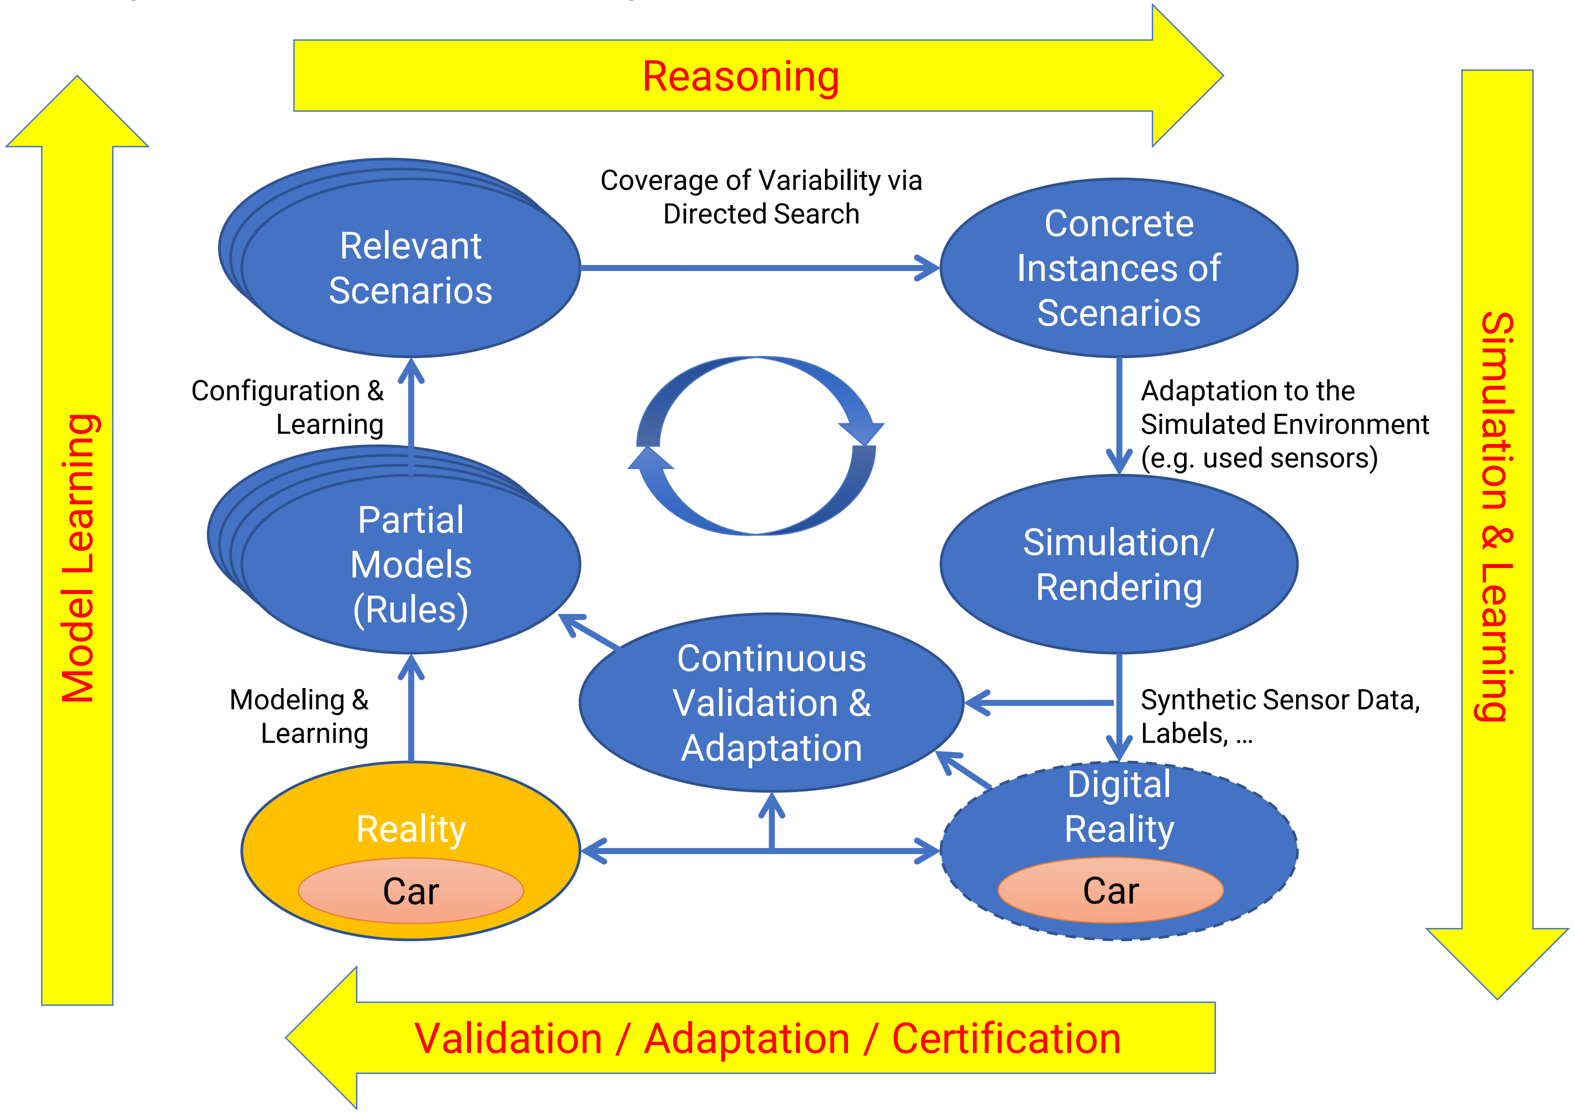

Digital Reality: AI to Certify AI

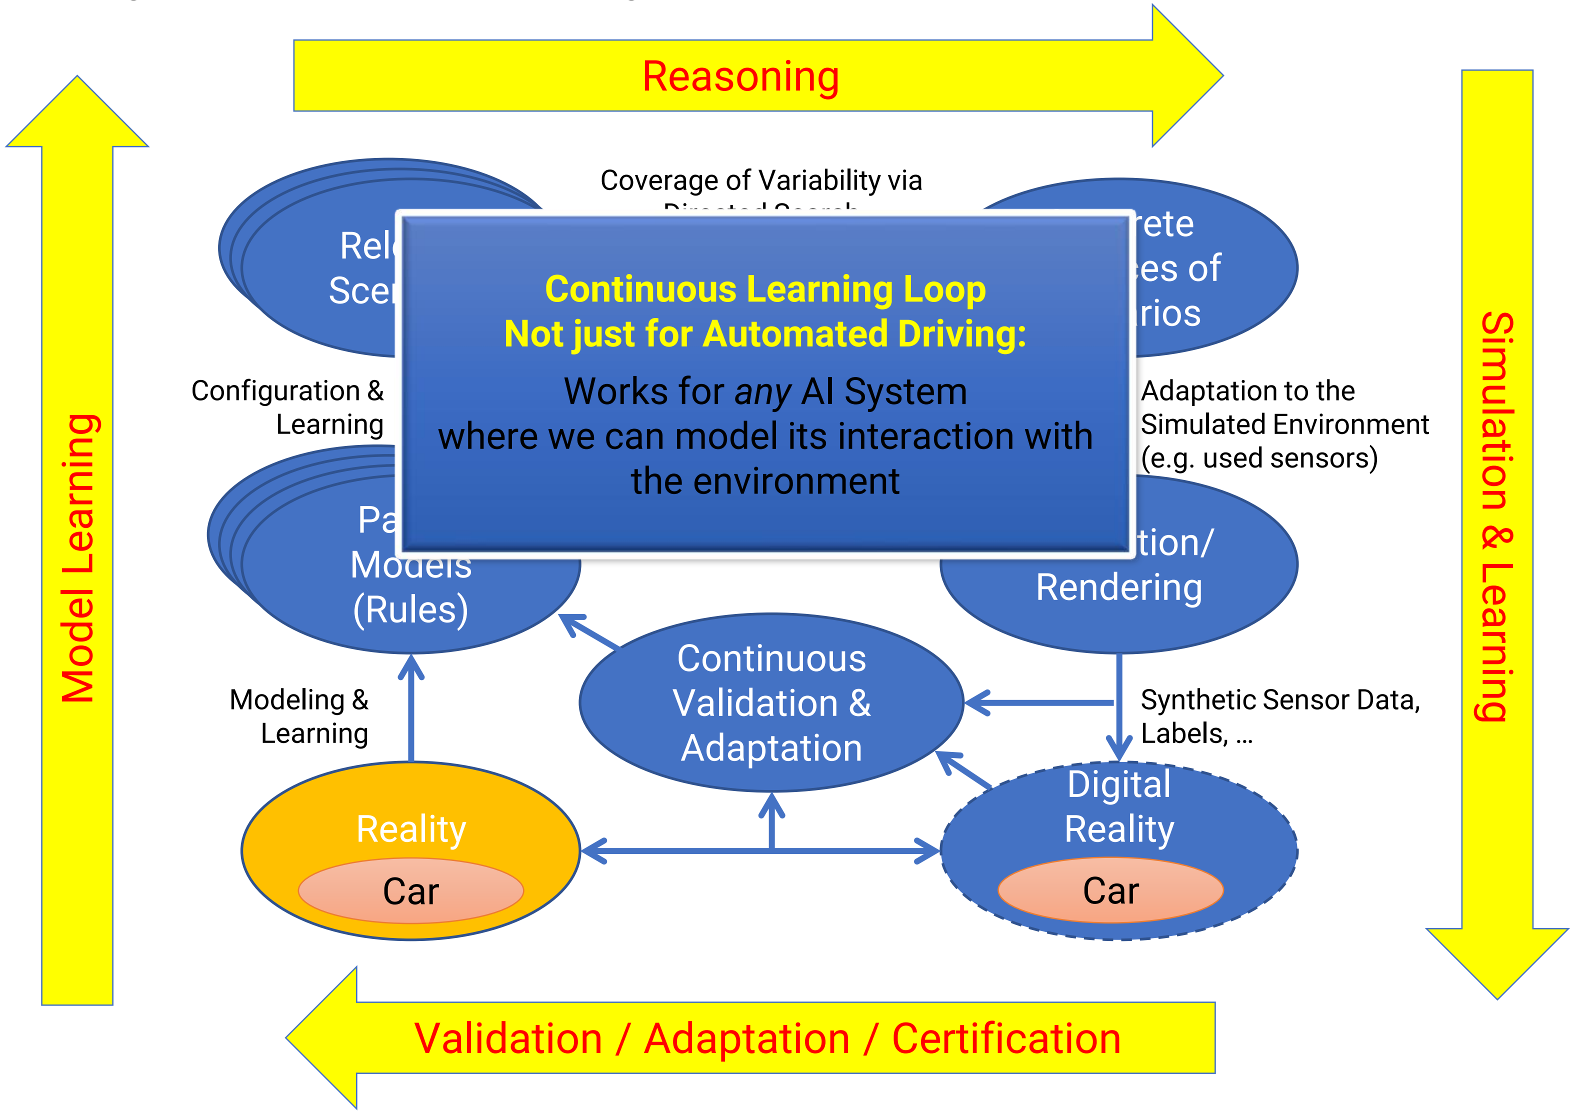

Radar Simulation

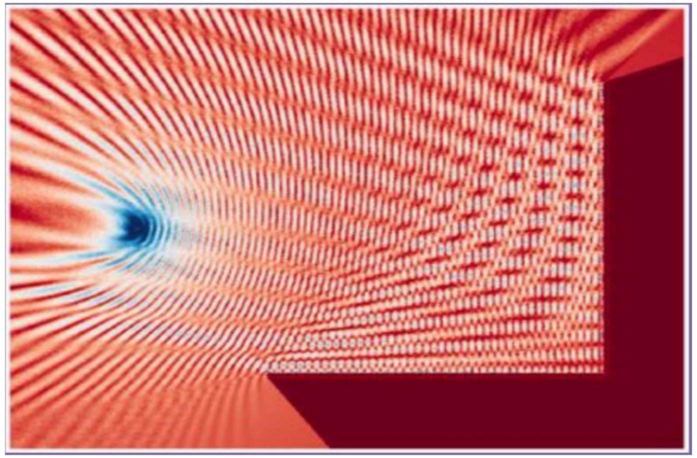

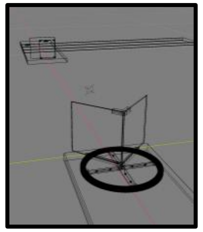

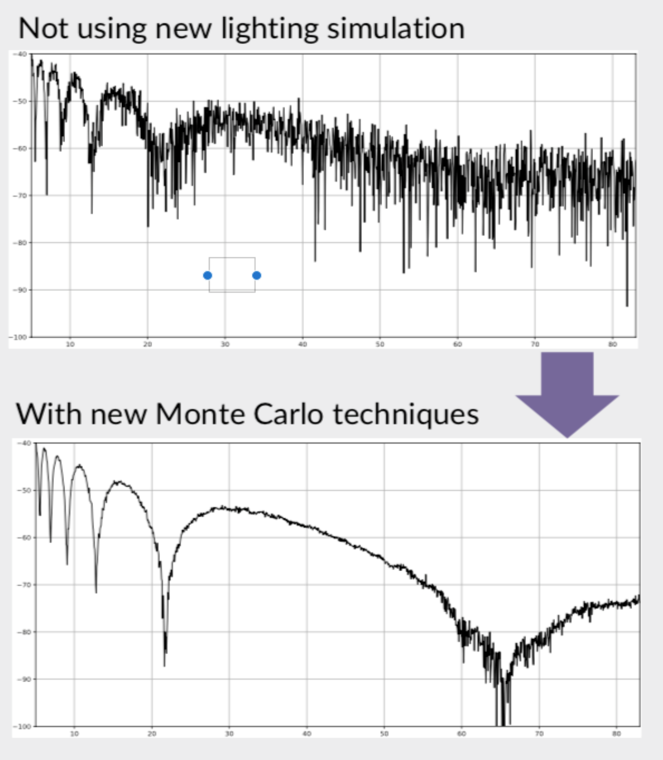

Ours Method: 33 seconds

(Physical Optics + Monte Carlo)

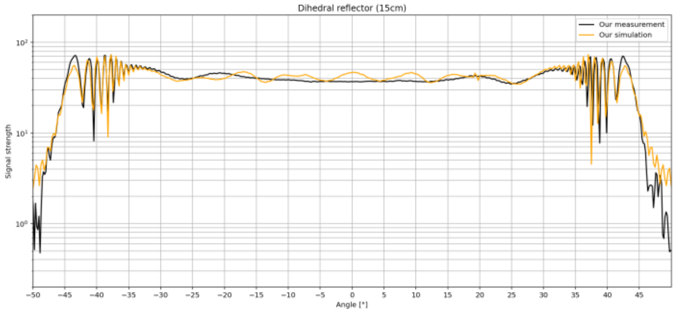

EM.Illumina: 13.4 hours

(Physical Optics + Finite Elements)

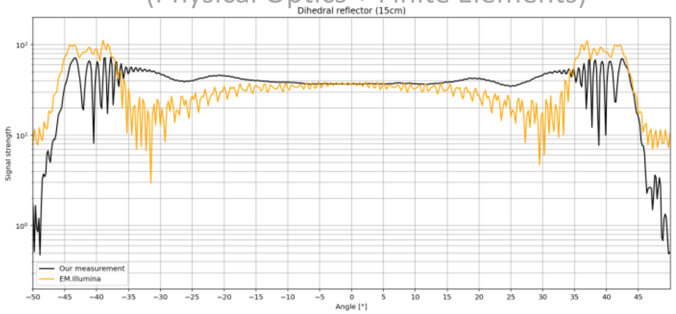

Bridging the gap between radar simulation & modern computer graphics

Our resulting method is over 1,000x faster than existing commercial software, while still achieving better accuracy

Autonomous Driving: Training using Synthetic Sensor Data and Realistic Models (TÜV, VDA, ZF, Conti, …)

Motion Modeling and Synthesis

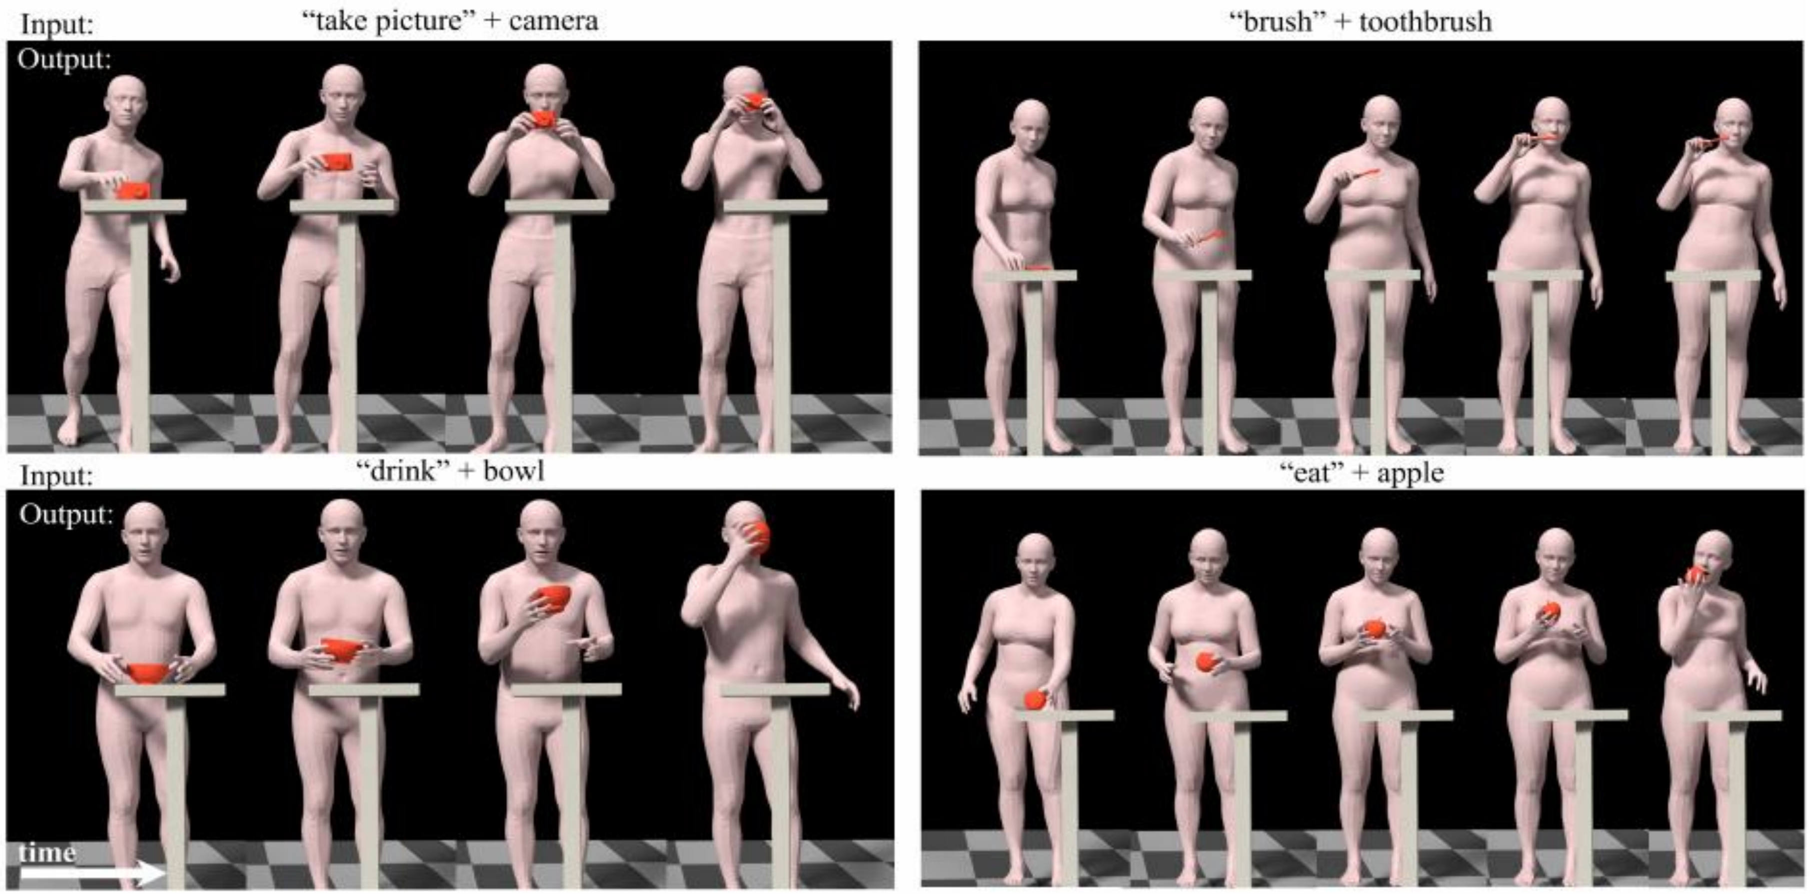

Gosh et al – Using action descriptions to drive motion synthesis via learned models
Eurographics 2023

Collaborative Robotics and Simulated Reality (VW, Airbus, …)

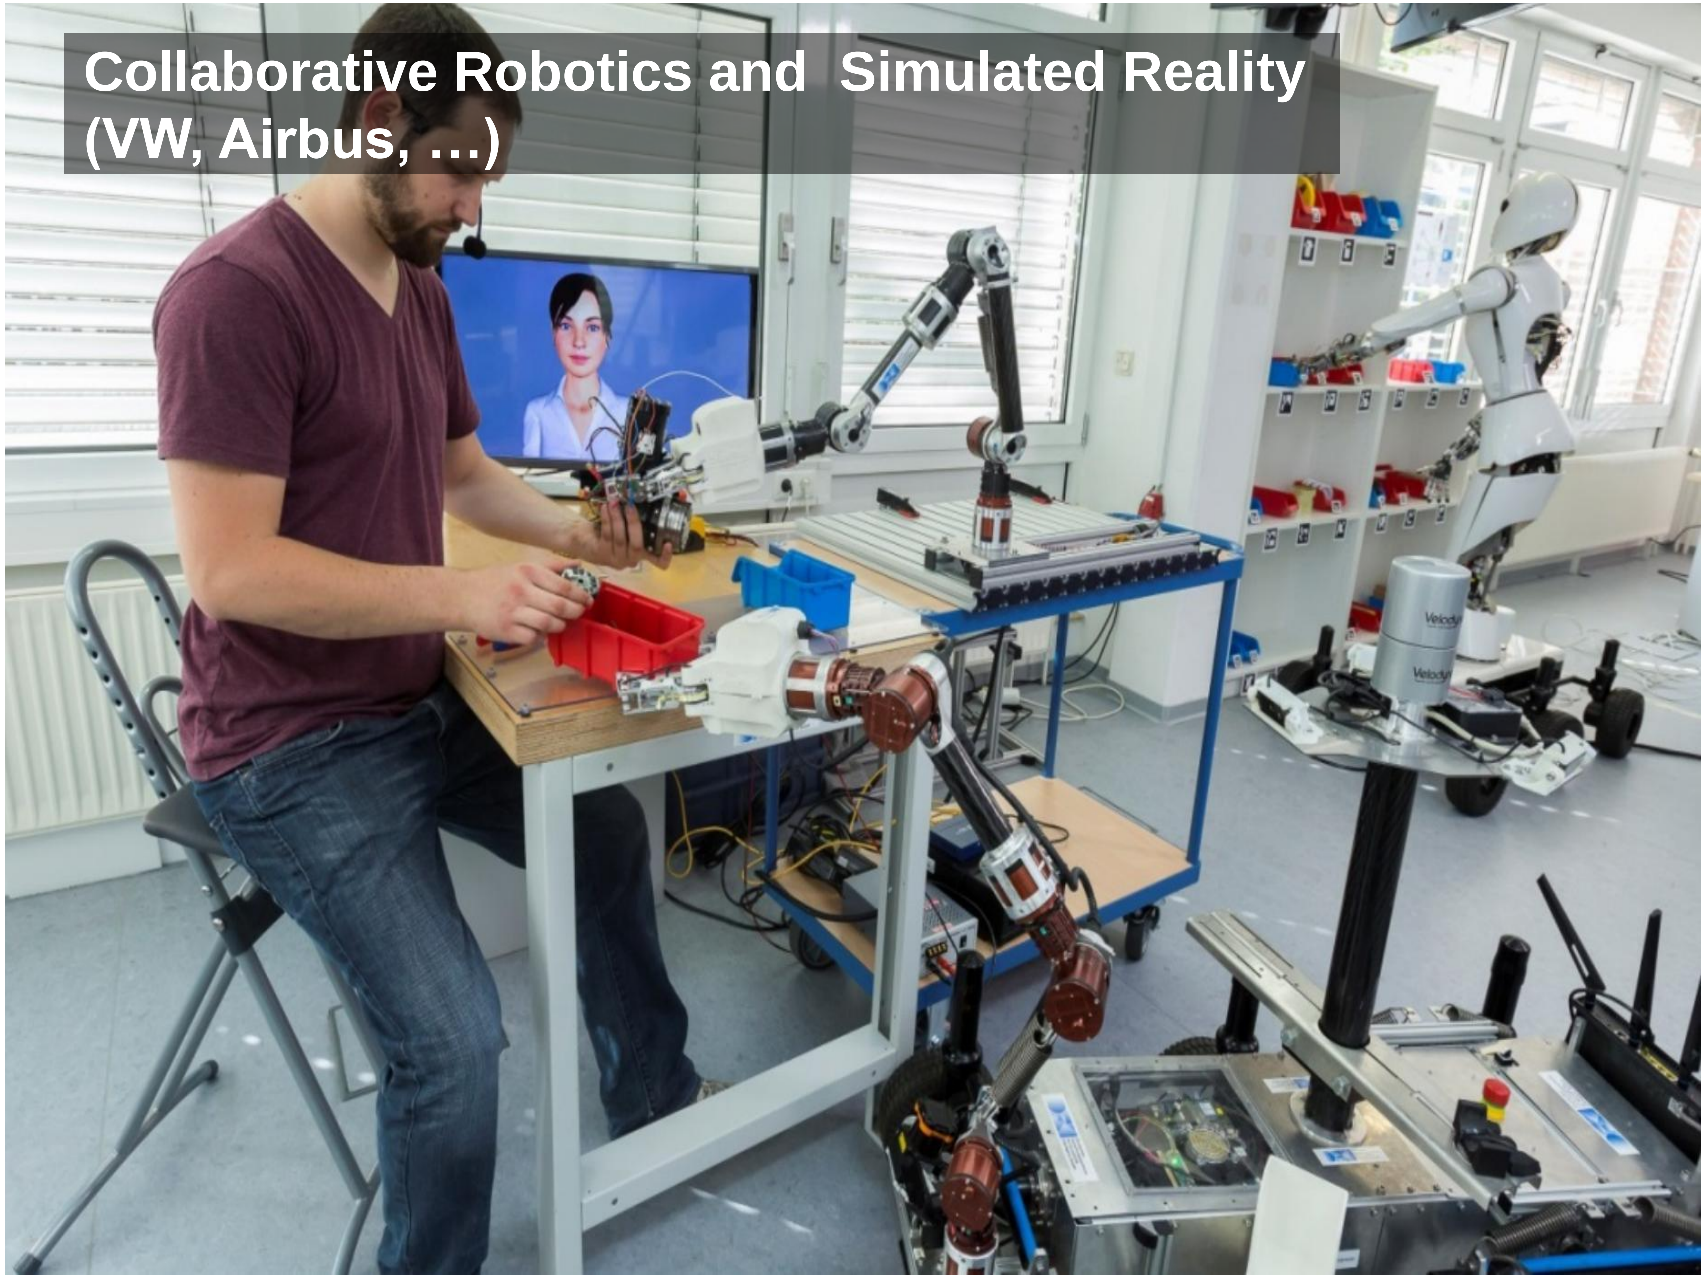

Models of the World

- Long history in motion research (>50 years)
  - E.g. Gunnar Johansson's Point Light Walkers (1974)
- Humans can easily identify more than what we see
  - Identify the person with high probability
  - Perceive properties like gender, age, weight, mood, ...
  - Based on minimal information
- Can we teach machines the same?
  - Currently, only bottom-up analysis
  - Neuroscience: Humans strongly perceive also top-down

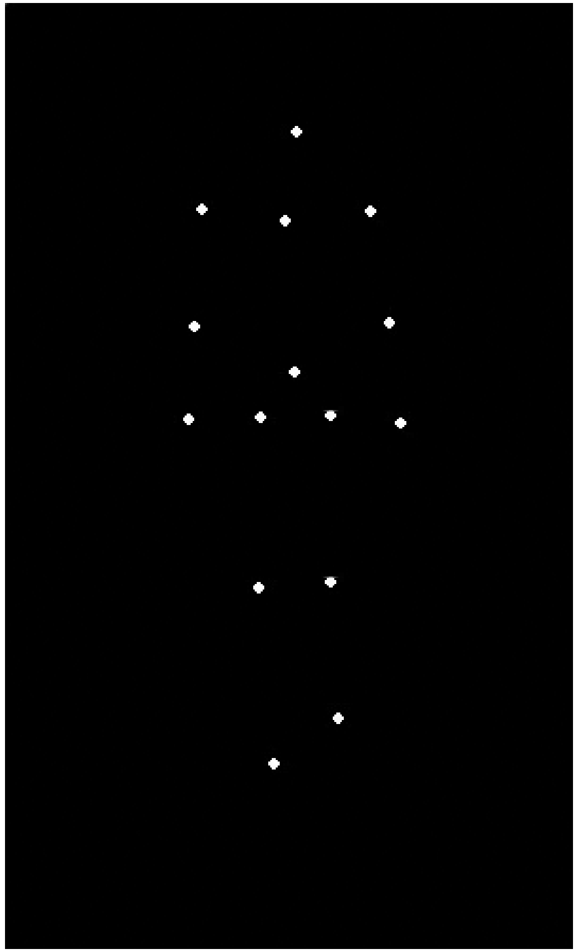

Models & AI:
Relationship Between Humans and AI

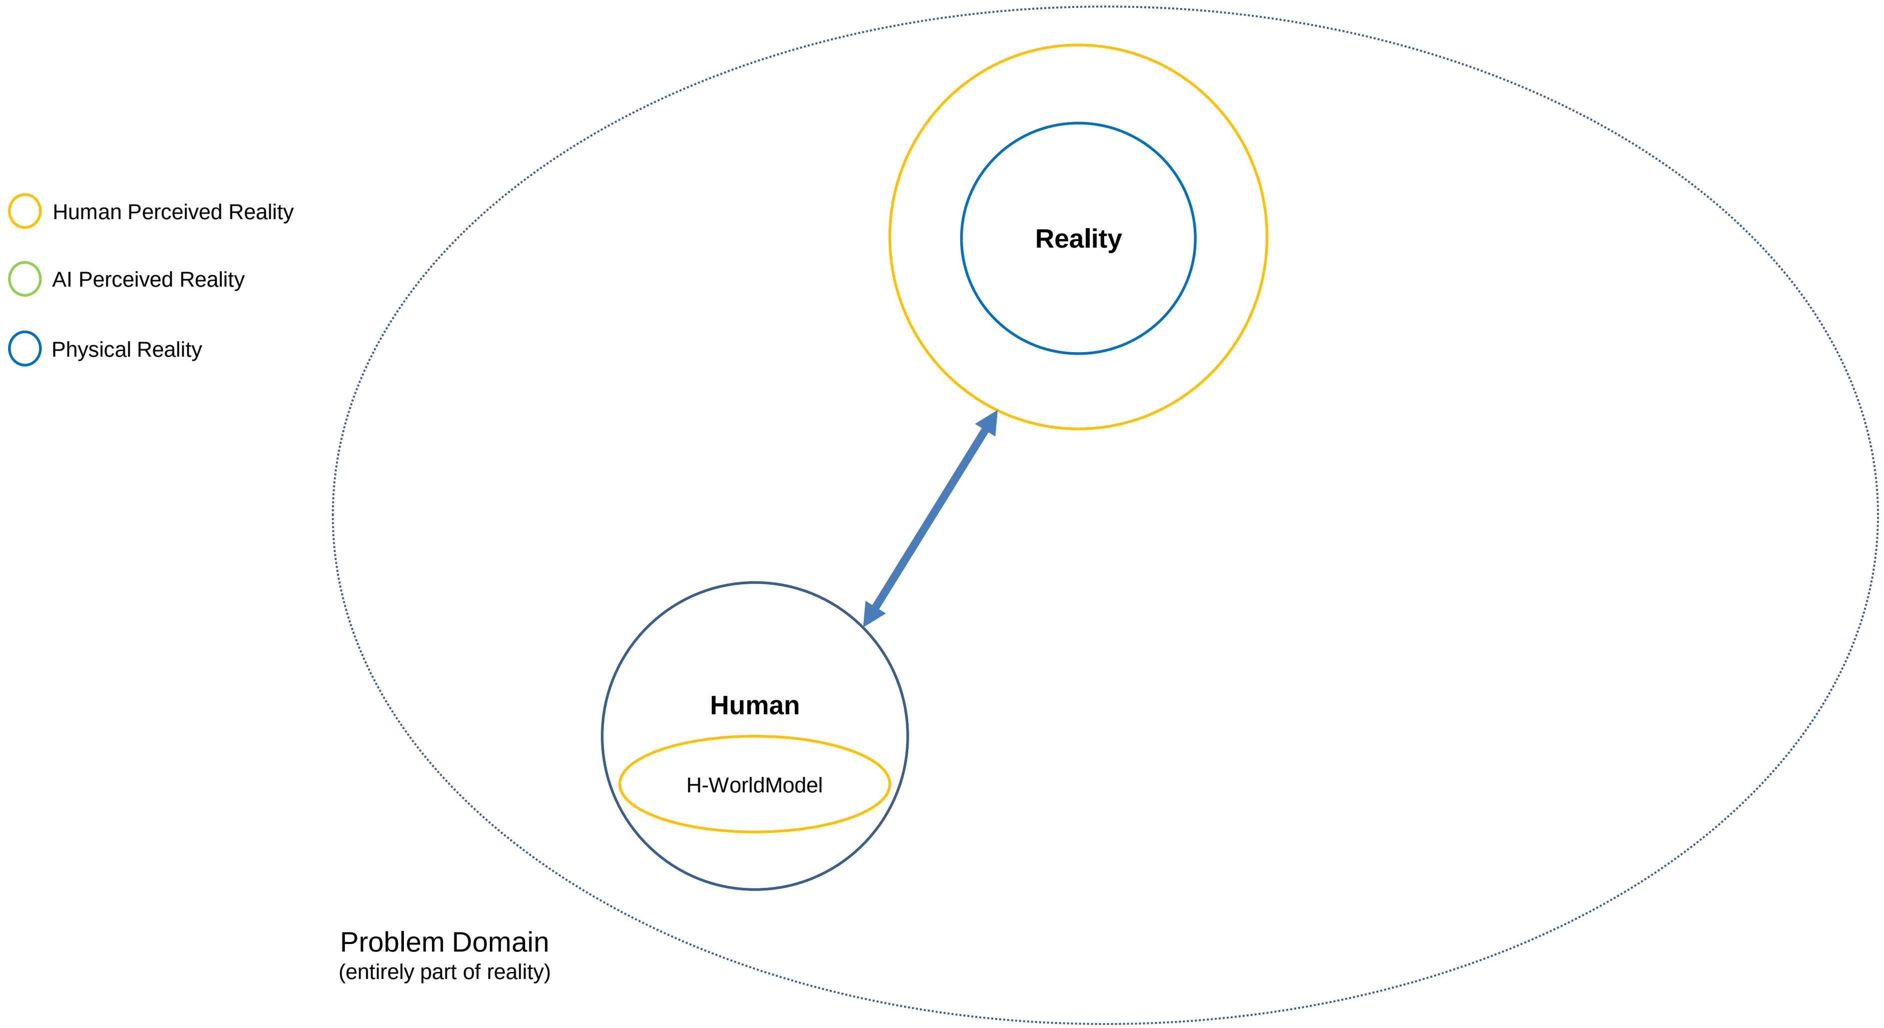

Models & AI:
Relationship Between Humans and AI

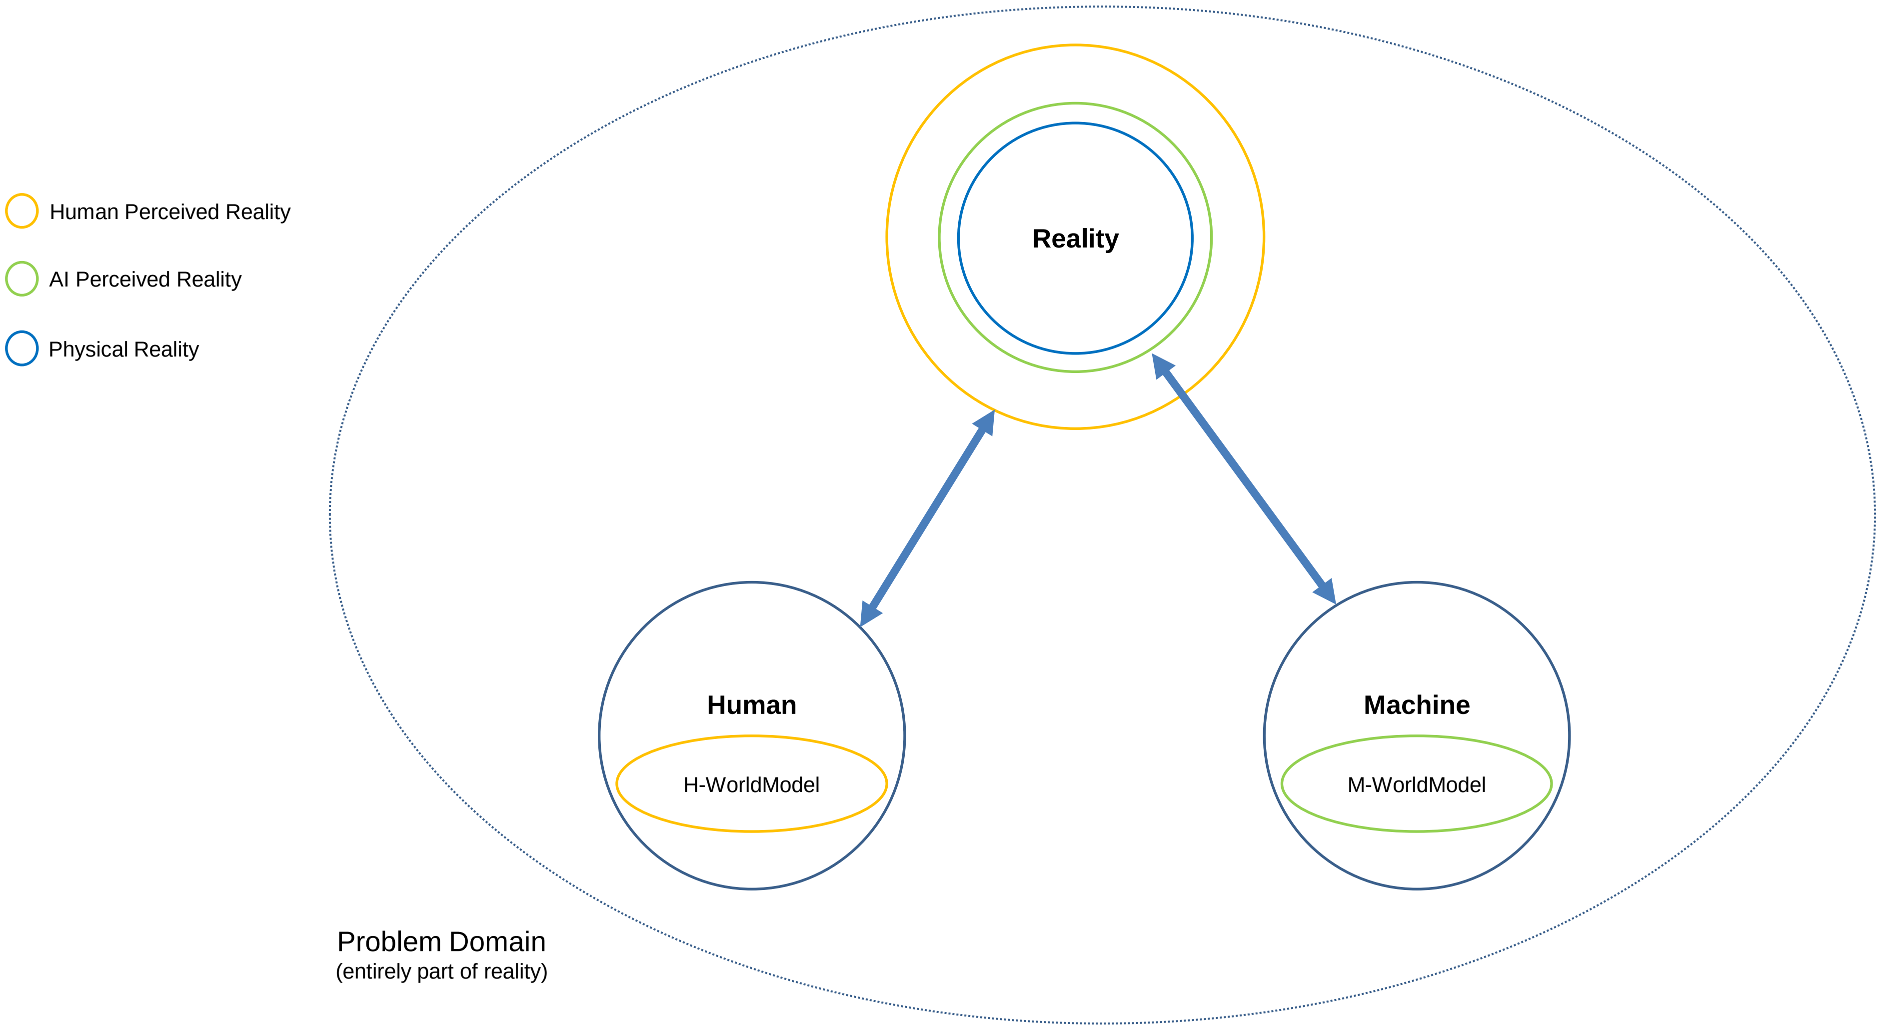

Models & AI:
Relationship Between Humans and AI

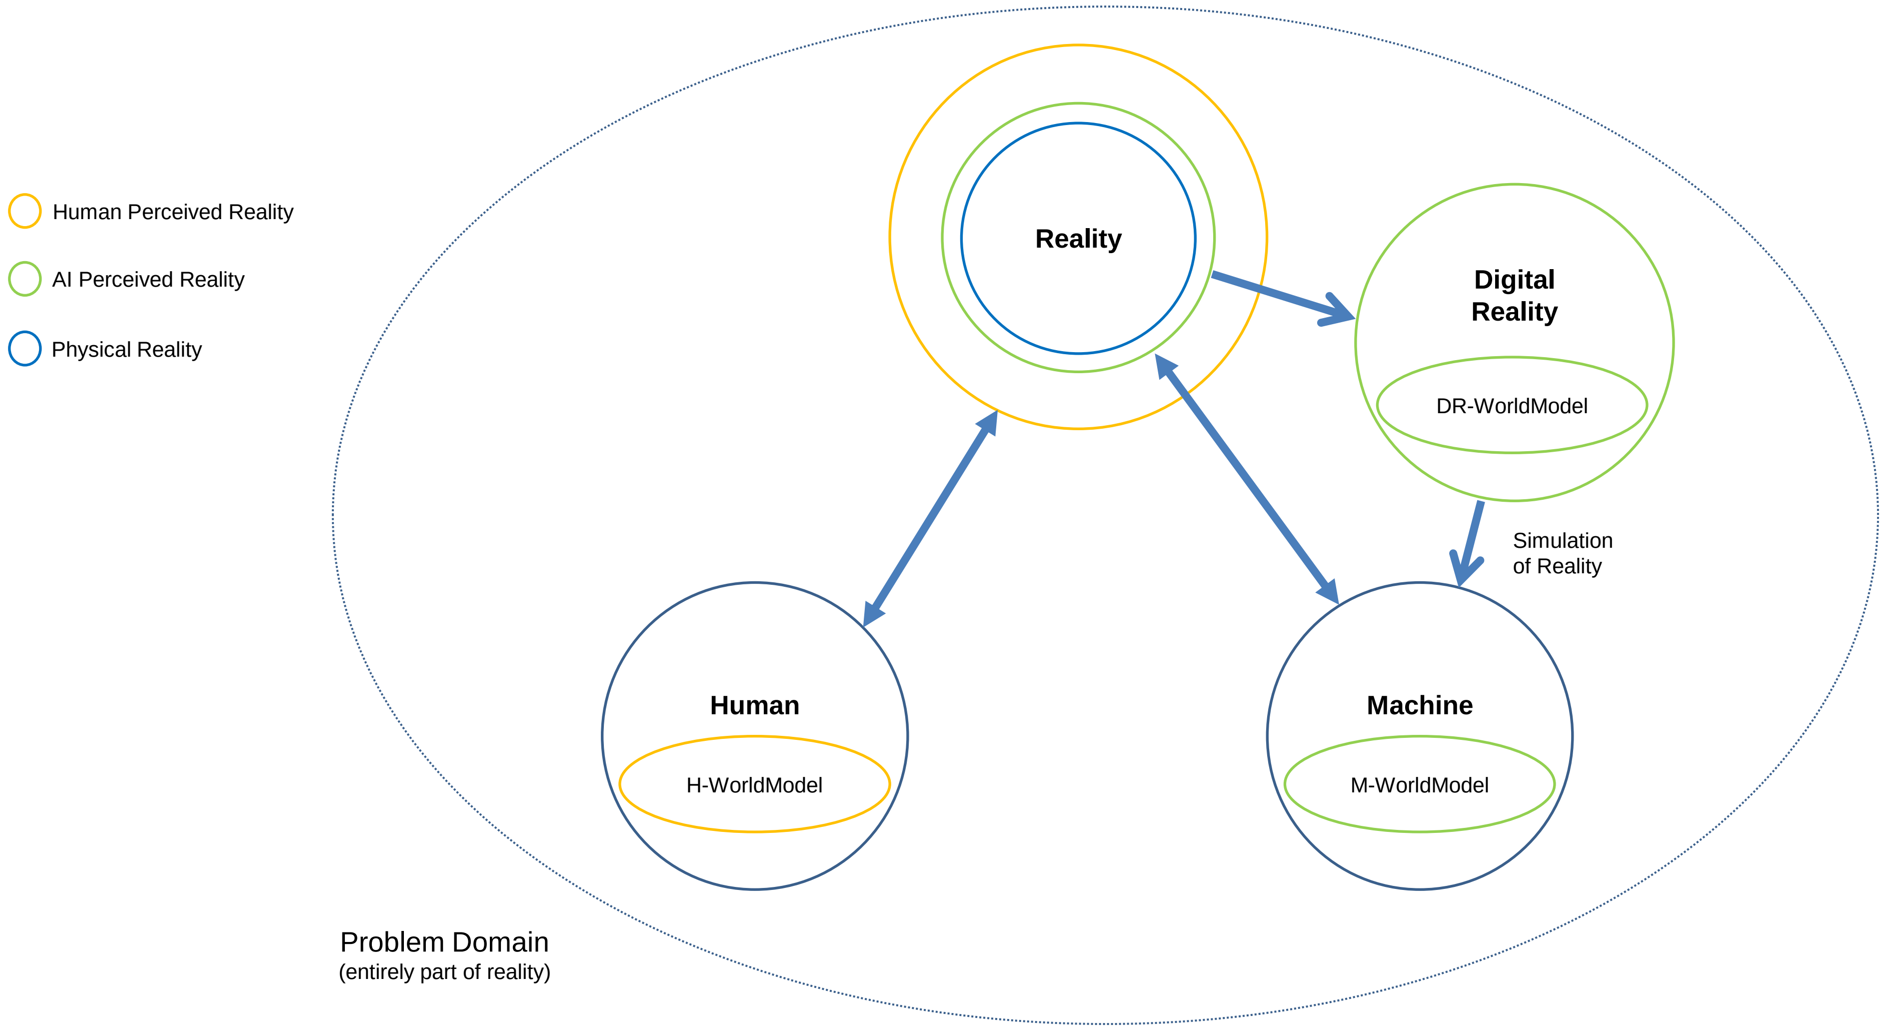

Models & AI:
Relationship Between Humans and AI

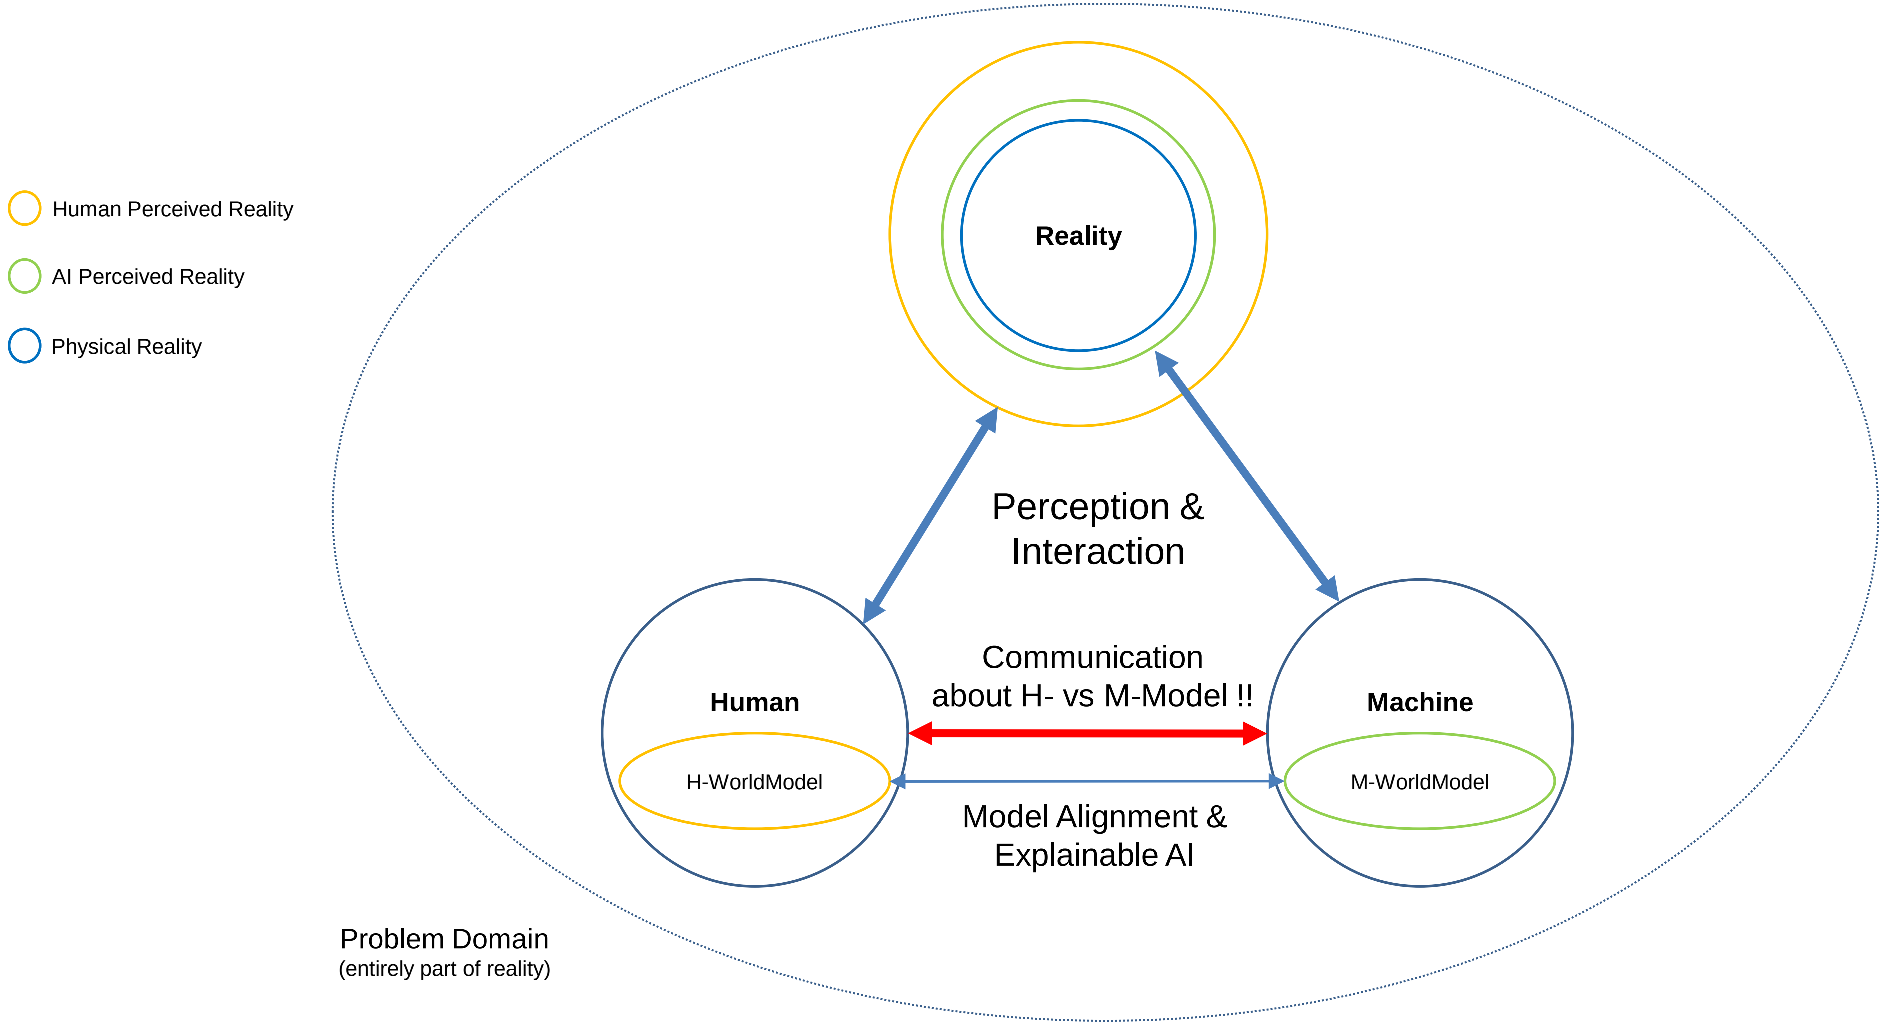

Neuro-Explicit AI Models (1)

- Need to move from ChatGPT to ActGPT
  - Not just words (or pixels) but modeling the physical world
    - 3D structures, motion, masses & forces, illumination, surface properties, …
  - Need vectors, representing these properties and their relationships/context
- Neuro-explicit AI models
  - We already know how the world works (physics, chemistry, …) – no need to re-learn
  - Use explicit models as the core (differential equations, simulations, logic models, …)
  - Use neural models to learn and model the difference to the real world
- Key Role for Trusted AI
  - Need for guarantees about the behavior of physical/embedded AI systems
  - ChatGT hallucinating text is already really bad
  - But a hallucinating robot can (literally) wreak havoc

CERTAIN: Trusted AI should give Guarantees for…

Guarantees for Trusted AI

By Design	By Tools	By Insight	By Interaction
• Intrinsic Correctness • Deductive Arguments & Proofs • (Physical) Laws, Rules & Constraints	• Modelling and Simulating the Real World • Systematic Testing with Synthetic Data • Monitoring, Auditing	• Explanations, Reasons • Causality • Transparency, Accountability, Visualization	• Human Experience, Influence, Control • Reinforcement Learning from Human Feedback (RLHF) • Useable Trust, Trust Calibration

Neuro-explicit AI Models

Ethics

Standards

Data

Ultimate Goal: Can we Teach Computers to "Understand" and Simulate the World Around Us?