Path Guiding in Production
Courses

JIŘÍ VORBA
WETA DIGITAL

Motivation

- TODO: split this into multiple slides, should be almost only pictures and illustrations
- Show average ray numbers and times for rendering typical movie scenes
- Show an example of such a scene
- Not only heavy on data, light transport is the problem
- Infamous MC convergence rate - one over sqrt(N)
- Hardware progress is great, enables many more samples in the given time and will probably get even better in the future (hardware ray-tracing support) but…
- It is advantageous to identify inefficiencies per scene and adjust our sampling method so that we focus our effort where it matters in the scene
- Show examples of indirect / caustics

Presenters

Jiří Vorba
(Weta Digital)

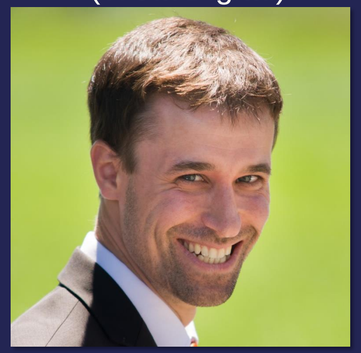

Johannes Hanika
(KIT / Weta Digital)

Sebastian Herholz
(University of Tübingen)

Thomas Müller
(NVIDIA)

Jaroslav Křivánek
(Charles University, Prague / Render Legion)

Alexander Keller
(NVIDIA)

Syllabus

- 14:00 – Opening Statements and Introduction [Jiří Vorba]
  - Overview
  - Introduction
- 14:15 – Theoretical Background [Jaroslav Křivánek]

Syllabus

- 14:30 – Bayesian Inference in Many-Light Sampling [Jaroslav Křivánek]
- 14:45 – Guiding and Shadow Rays [Alexander Keller]

Syllabus

- 15:15 – “Practical Path Guiding” in Production [Thomas Müller]
- 15:45 – Break (15 minutes)

Syllabus

- 16:00 – Volumetric Path Guiding [Sebastian Herholz]
- 16:30 – Guiding in Path Space [Johannes Hanika]
- 17:00 – Open Problems and Future Work [Jiří Vorba]

Goals

- Overview of existing methods
- Sharing practical experience
- Cover theoretical background
- Share open problems with researchers

Introduction

What is path guiding

Path guiding

- What is path guiding?
  - Set of **adaptive** path sampling techniques aware of the scene content
- Applicable in various transport algorithms (unidirection path tracing, bi-directional methods)

Path tracing

- Averaging of many sampled paths
- Efficiency depends on a few sampling decisions

Path tracing – sampling decisions

- Scattering (BRDF sampling)
- Light sampling (Next-event estimation)
- Absorption (Path length)
- Free flight (ray distance sampling)

Scattering (BRDF sampling)

- Challenge: Indirect illumination, visibility

Direct illumination

- Next-Event estimation
- Challenge: Many-light sampling, visibility

Path length

- Ideally short paths, but not shorter
- Russian roulette: albedo based

Key to efficiency

- Standard sampling decisions/schemes are local
- We need global knowledge (radiance)
- Example: BRDF * Radiance
- Zero-variance sampling theory
- Is it useful?

Learning

- Radiance not known a-priory
- Learning approximation from samples
- Improved importance sampling
- Path guiding = guiding the sampling decisions (based on the learned approximation)

Path guiding “How to”

- How to learn from samples?
  - Machine learning
- How to represent the knowledge?
  - Parametic / Non-parametic models
- How to exploit it in the simulation?
  - Depends on the model and the type of the sampling decision

Scattering

Guided directional sampling

History

Learning from photons

- Jensen [1995]

Learning from forward samples

- Lafortune and Willems [1995]

History

Learning from photons

- Jensen [1995]
- Hey and Purgathofer [2002]
- Vorba et al. [2014]
- Vorba et Křivánek [2016]
- Herholz et al. [2016, 2019]

Learning from forward samples

- Lafortune and Willems [1995]
- Pegoraro et al. [2008]
- Bashford-Rogers et al. [2012]
- Müller et al. [2017]
- Dahm and Keller [2018]

Directional path guiding

- TODO: Bit more detail on Vorba et al. 2014 (will illustrate guiding by a concrete method, this method is/can be used in practice)

- TODO: Explain pre-training, used representation, how it is used in the rendering

Path length

Guided Russian roulette and splitting

Guided Russian roulette and splitting

- Importance sampling of path length
- Splitting when expected contribution is high
- Vorba et Křivánek [2016]

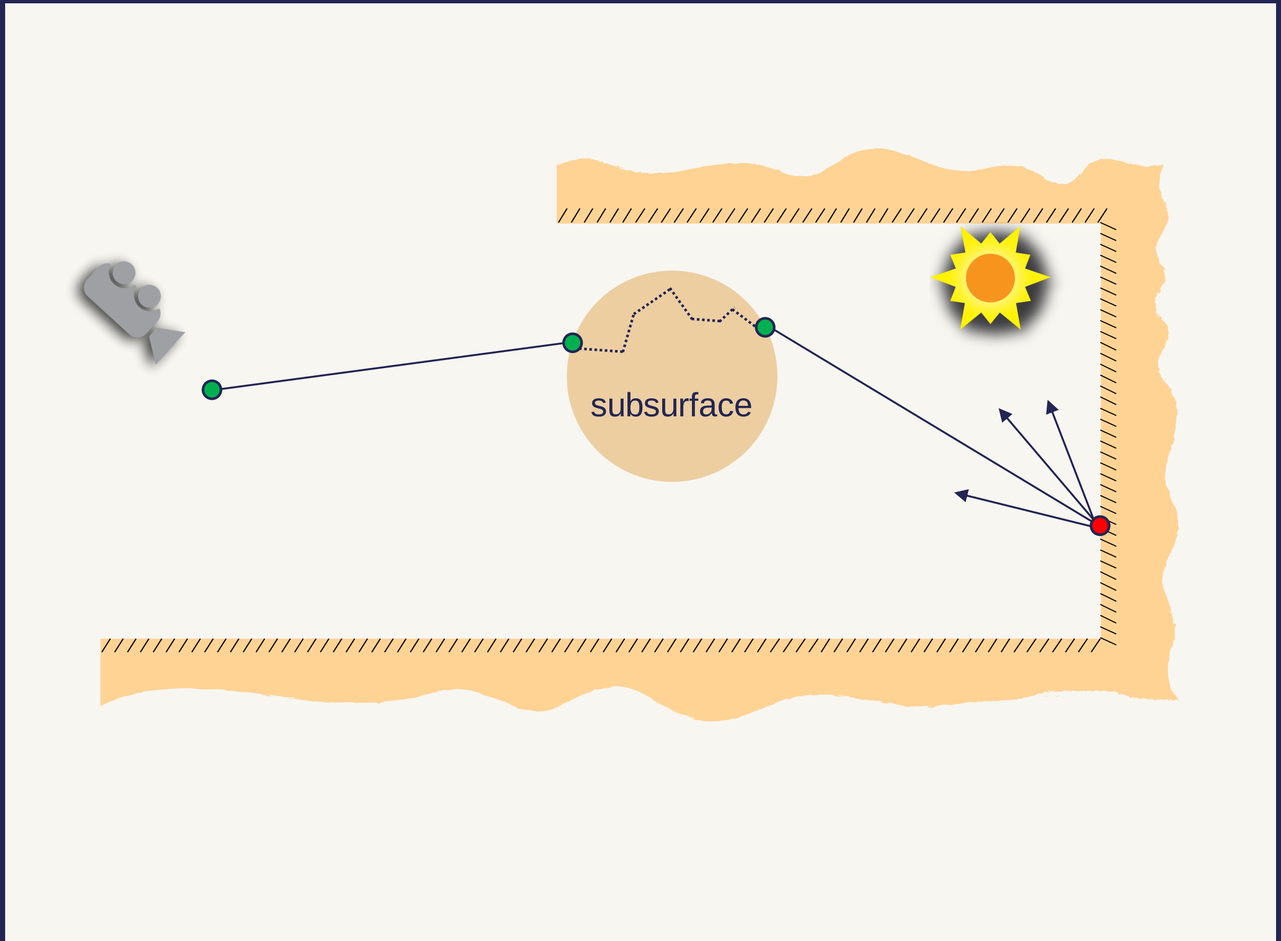

Albedo based Russian roulette

- Termination probability

Current path weight

$$q = \frac{\nu(x_1, x_2, \ldots, x_k)}{t}$$

User given threshold

[Arvo & Kirk 1990]

Albedo based Russian roulette

- Termination probability

Current path weight

$$q = \frac{\nu(x_1, x_2, \ldots, x_k)}{t}$$

User given threshold

[Arvo & Kirk 1990]

$$q = \rho$$

[Jensen 2001]

Albedo

Albedo based Russian roulette

- Termination probability

- Problem: it's local
- Kill paths too early
- Waste time on long paths

Current path weight

$$q = \frac{\nu(x_1, x_2, \ldots, x_k)}{t}$$

User given threshold

[Arvo & Kirk 1990]

$$q = \rho$$

[Jensen 2001]

Albedo


Path tracing (1h)


Path tracing (1h)
Guided RR and splitting (1h)


Path tracing (1h)
Guided RR and splitting (1h)


Path tracing (1h)
Guided RR and splitting (1h)

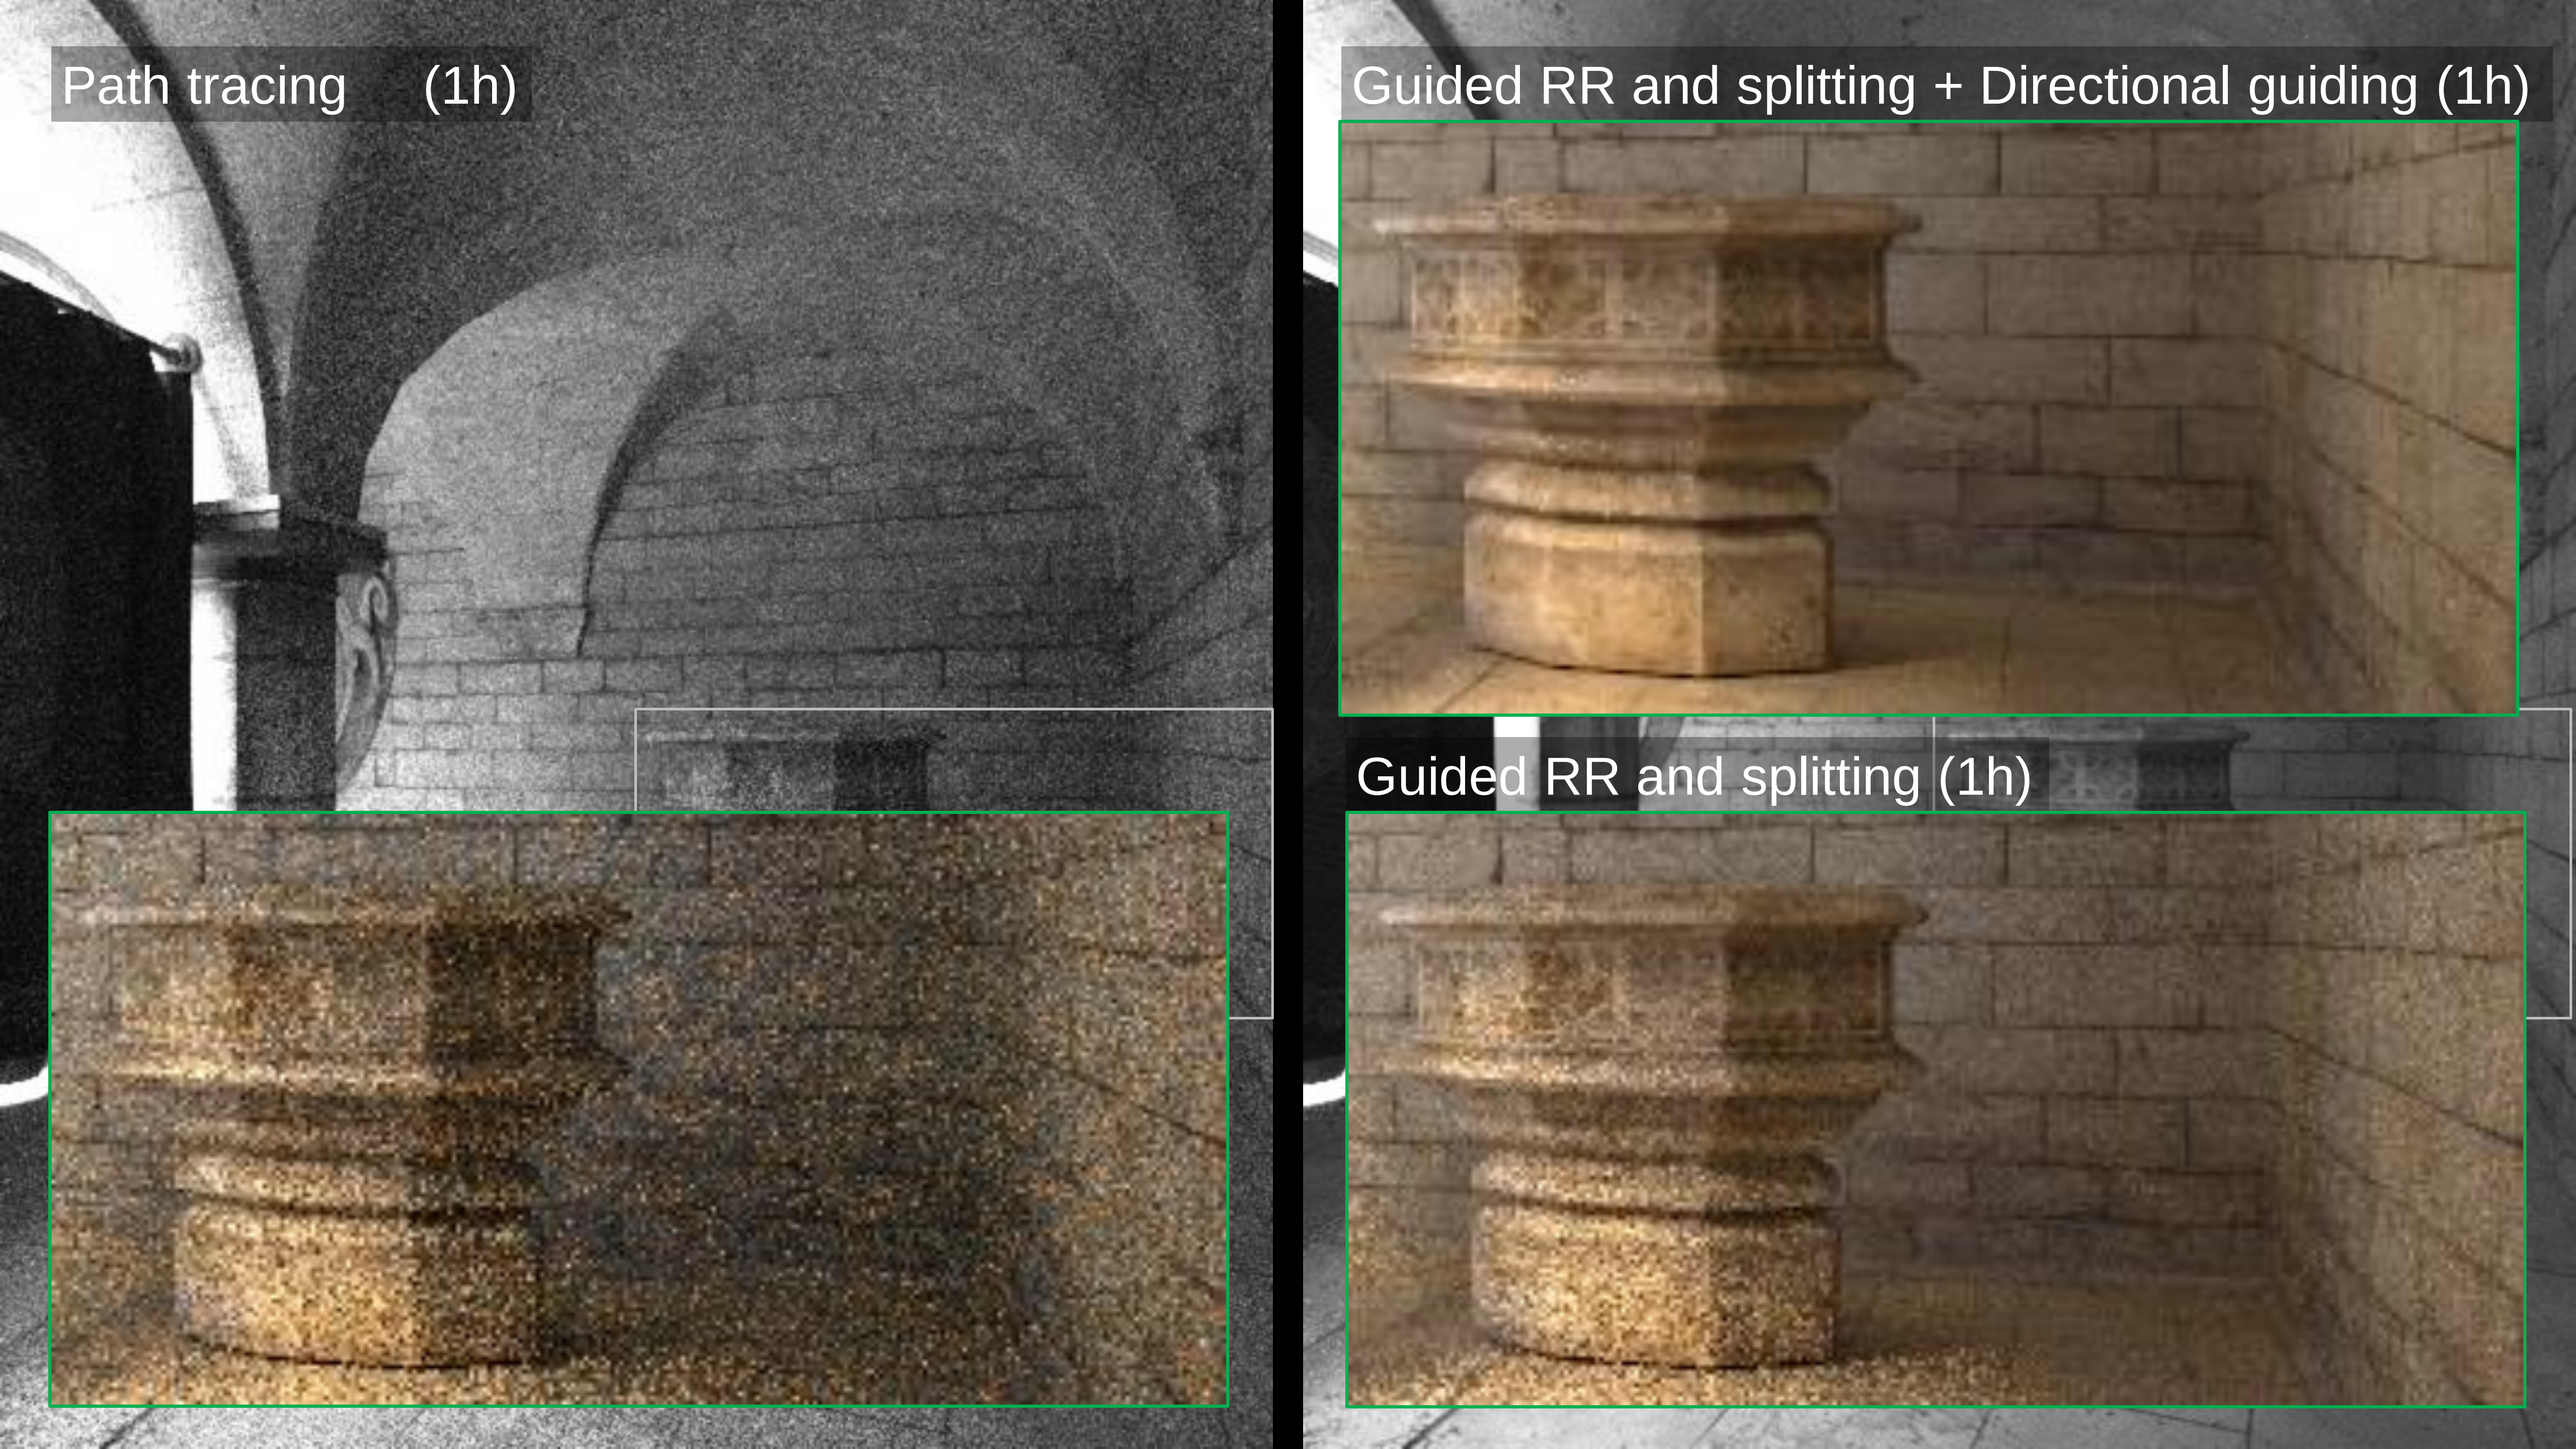
Path tracing (1h)
Guided RR and splitting + Directional guiding (1h)
Guided RR and splitting (1h)

Guided Russian roulette and splitting

- Input
  - approximation of radiance field
  - estimate of pixel values
- Output
  - Termination probability / path split ratio

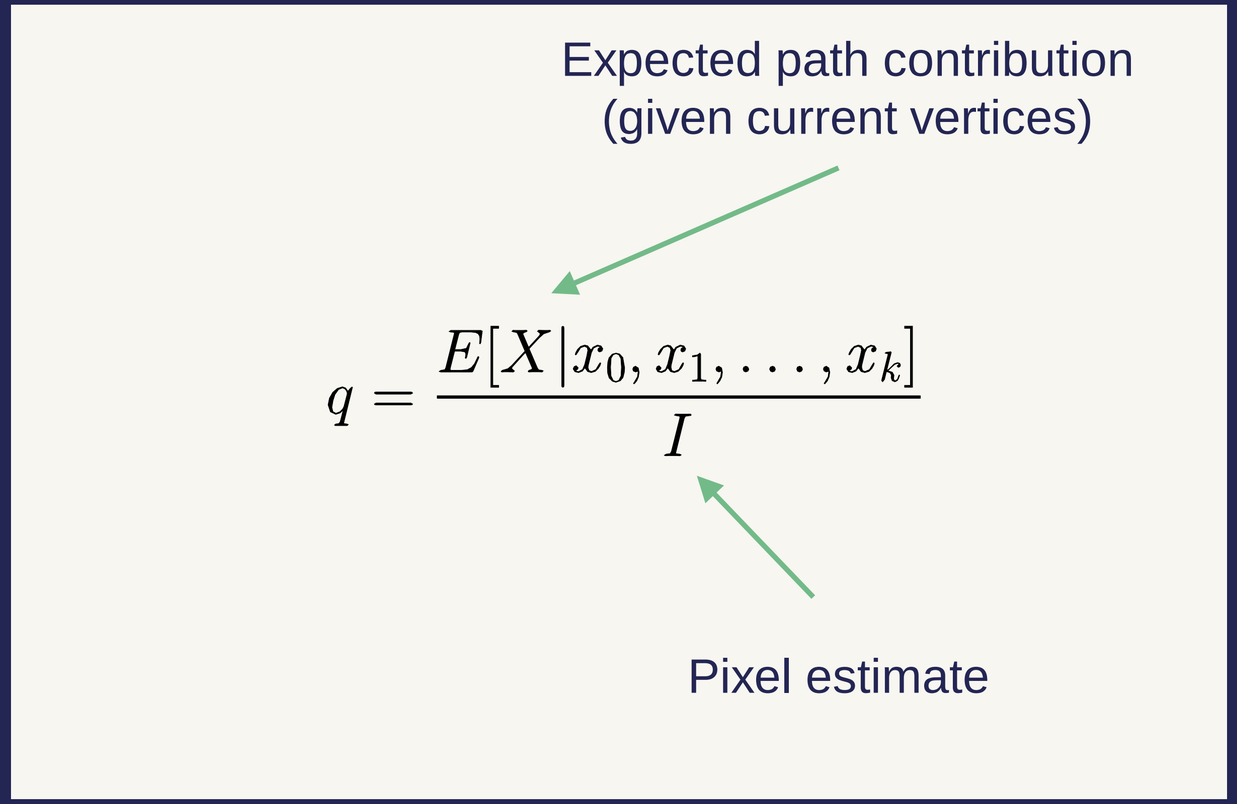

Pixel value estimates (Vorba et Křivánek [2016])

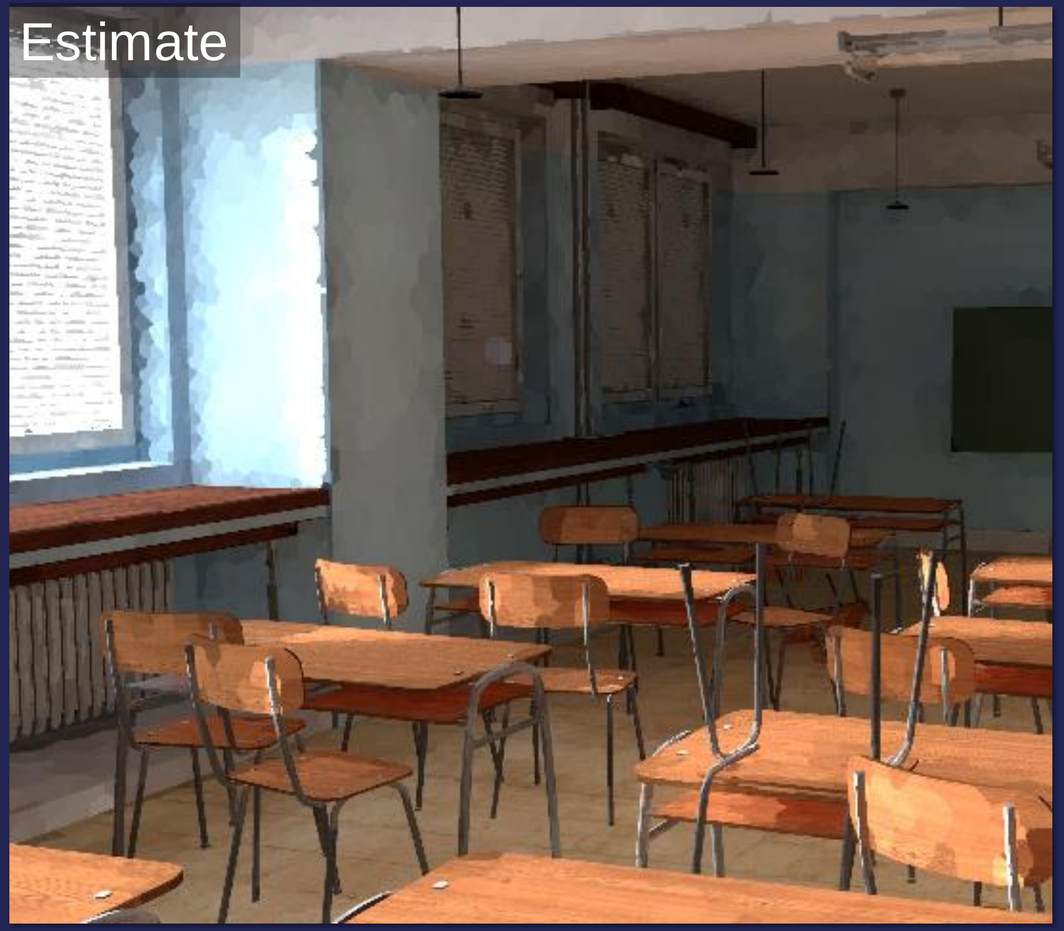

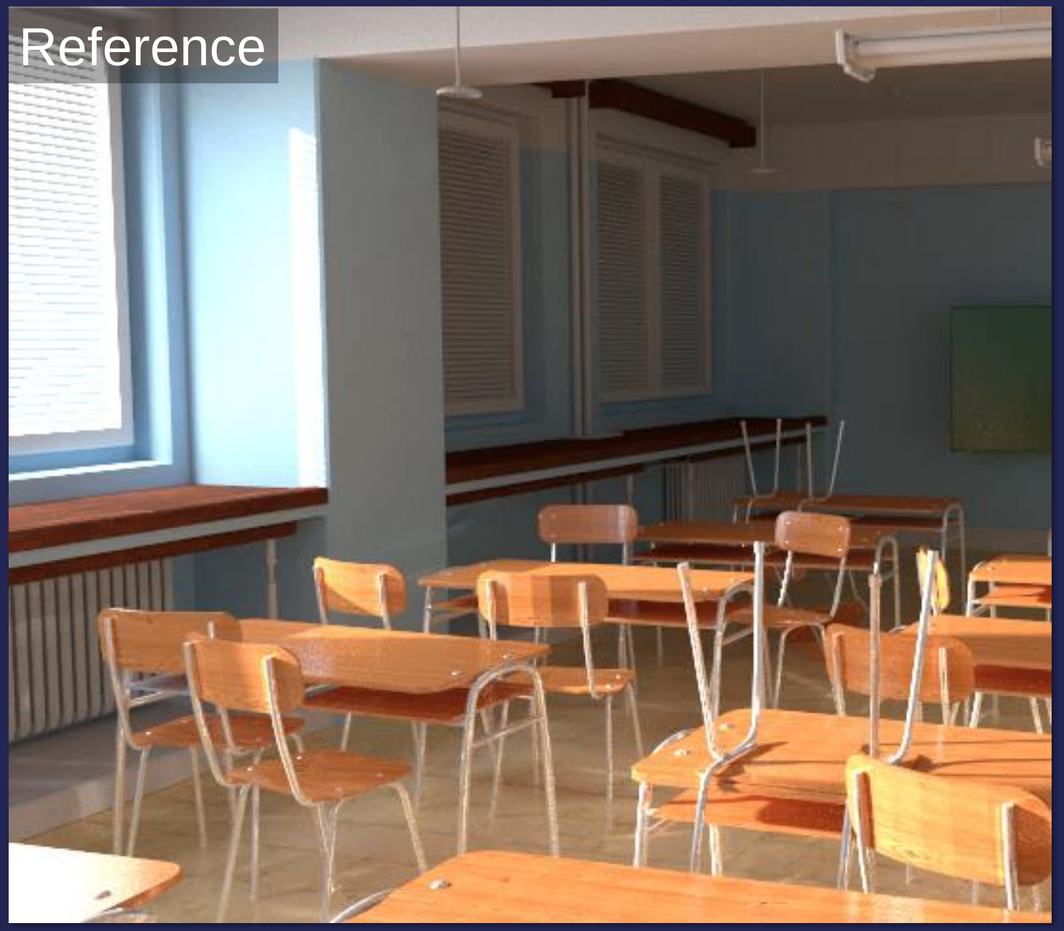

Pixel value estimates (progressive rendering)

- Can be simplified in practice
- Many possible approaches (low sample count -> denoising)
- MIP mapping of beauty image (at Weta)
- TODO: pics

Guided Russian roulette and splitting

- Minimal overhead on top of directional guiding
- Synergic effect
- Makes guiding cheap (even on simple scenes)

Practical method

- Photons – longer time to first pixel
- Forward – fits in progressive rendering
- But forward can learn slowly
  - E.g. caustics
- Ideal method low overhead, is progressive, fast learning

Guiding (photon) emission

- TODO: would be nice to describe what we have done for ABA, caustics and god-rays if we have time

Guiding in Bi-directional algorithms

- Possible to guide
- Say that Alita is PLT (path tracing and light tracing)
- Guided PT is not efficient enough on caustics
- Show/say why
- Photons do not allow for bending physics (for example point-of-entry)
- We do not have light tracing on specular transmission
- We don’t use it in hair, don’t use it on skin
- Together more robust algorithm
- Ideally we wish for forward guiding only method that would cope even with ocean rendering

Wrapping up

- Defined path guiding = Family of adaptive path sampling methods
- Importance sample all the decisions along the path -> Superior convergence rate
  - Scattering (directional guiding)
  - Path length (guided Russian roulette and splitting)
  - Direct light [Jaroslav, Alexander]
  - Free-flight [Sebastian]
- Guided photon emission (caustics in production)
- What would be the ideal production method [Jirka – last session]

THANK YOU