OPEN PROBLEMS AND RESEARCH DIRECTIONS

Jaroslav Křivánek
Charles University | Render Legion | ChaosGroup

ACADEMIA – INDUSTRY

Computer
Graphics
Charles
University

corona

THE LIGHT TRANSPORT CHALLENGE

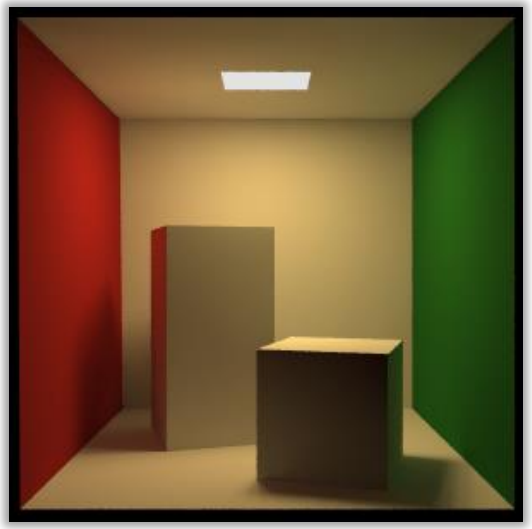

Algorithm that can renders this at least as fast as a path tracer…

… and it can also render this.

TODAYS' RENDERING IS OLD NEWS

- From **Matt Pharr**'s editorial to ACM TOG special issue on production rendering [Pharr 2018]:

*"Today … renderers are … based on … path tracing. Introduced … by Jim Kajiya (**1986**)."*

"Many advancements were made … including
*- more effective light sampling algorithms (Shirley et al. **1996**),*
*- high-quality sampling patterns (Kollig and Keller **2002**), and*
*- multiple importance sampling (Veach and Guibas **1995**),"*

*"… the core ray tracing [got] more efficient (Wald et al. **2001**)."*

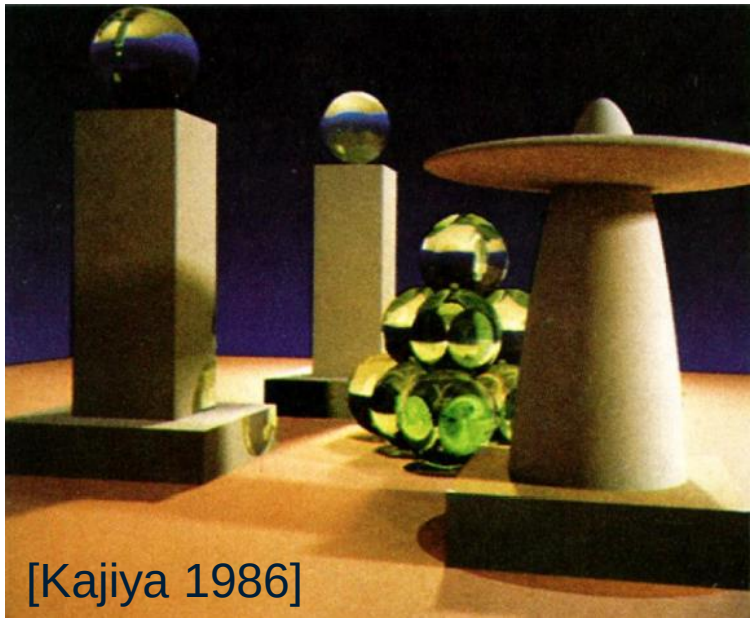

[Kajiya 1986]

TODAYS' RENDERING IS OLD NEWS

- Fundamental blocks all pre-2000's

- Why is that? What happened after 2000?
  - Denoising [Rouselle et al. 2011, 2012, 2013, …]
  - Path guiding [Vorba et al 2014, Muller et al. 2017]
  - "*Collection of special purpose sampling algorithms that work super well in very specific scenarios… Each on their own may be incremental but the toolbox allows us to render many scenes well*"
    (Johannes Hanika, Weta)

ADVANCED LIGHT TRANSPORT

- Why are advanced light transport algorithms not used in practice?

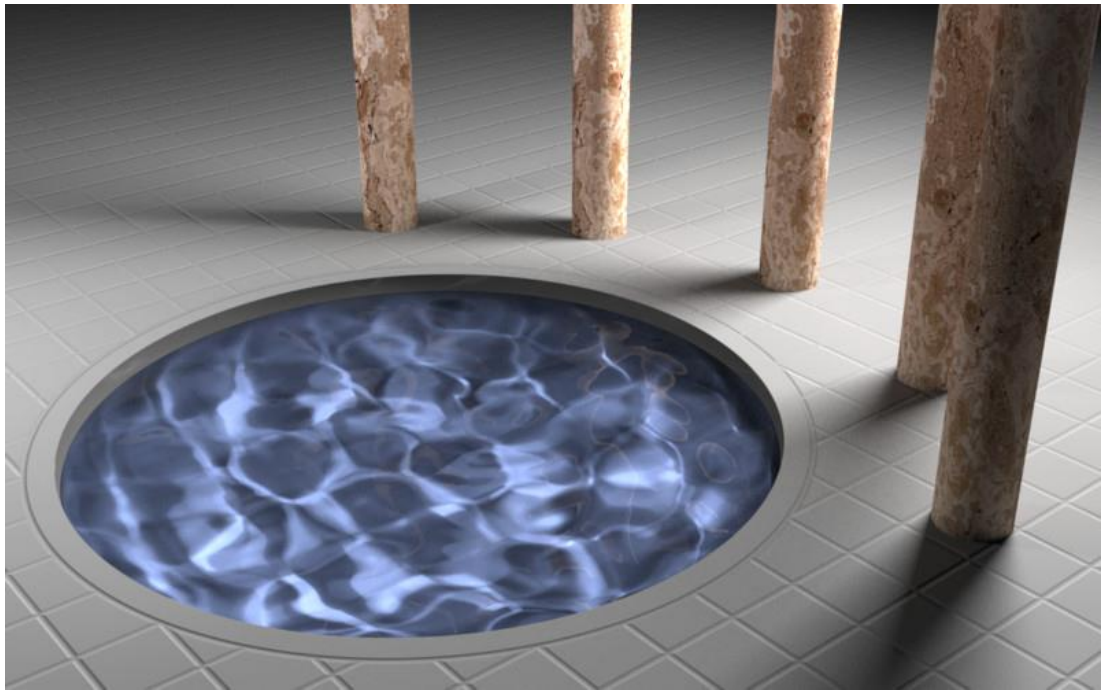

Metropolis Light Transport [Veach and Guibas 1997]

WHY PATH TRACING…

… when we have so many more advanced light transport algorithms?

WHY PATH TRACING?

- Let's see…
  - A "good" light transport algorithm
  - Existing light transport algorithms

A GOOD LIGHT TRANSPORT ALGORITHM ...

- **… does not need to be**
  - Physically-correct
  - Unbiased
- **… but it has to be**
  - Easy-to-use (no parameters)
  - Interactive & progressive
  - Fast in common scenes
  - Robust & reliable
  - Capable of failing gracefully (no conspicuous artifacts)
  - Simple to implement and maintain
  - Flexible (compatible with non-physical tricks)

THE GOOD ALGORITHM CHECKLIST

- Easy-to-use
- Interactive & progressive
- Fast in common scenes
- Robust & reliable
- Graceful fail (no artifacts)
- Simple
- Compatible with unphysical tricks

PATH TRACING

- [Kajiya 1986, Veach and Guibas 1995, Shirley 1996,…]

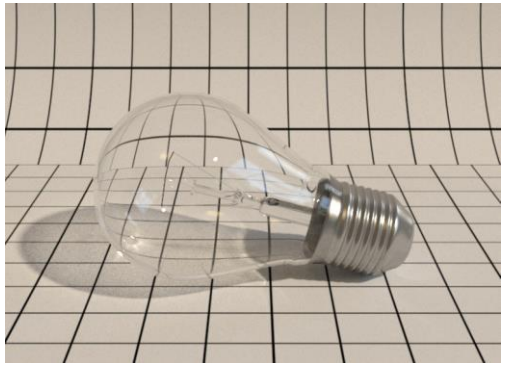

Reference rendering
(VCM)

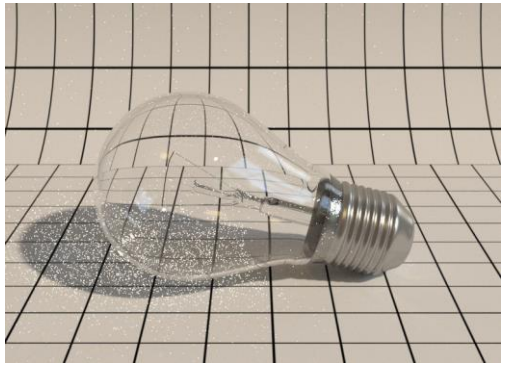

Path tracing
(no clamping)

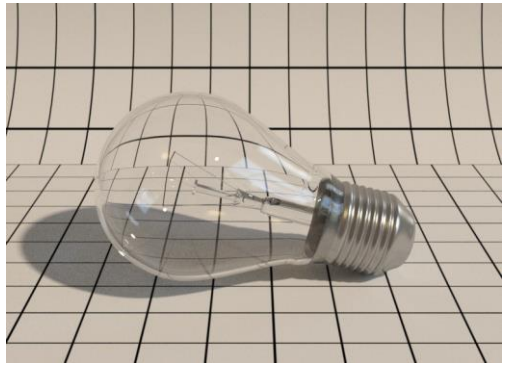

Path tracing
(with clamping)

- Easy-to-use
- Interactive & progressive
- Fast in common scenes
- ~~**Robust & reliable**~~
- Graceful fail (no artifacts)
- Simple
- Compatible with unphysical tricks

PHOTON MAPPING

- [Jensen 1996]

11 minutes on a 733 MHz Pentium III

- ~~Easy-to-use~~
- ~~Interactive & progressive~~
- Fast in common scenes
- ~~Robust & reliable~~
- ~~Graceful fail (no artifacts)~~
- Simple
- Compatible with unphysical tricks

BIDIRECTIONAL PATH TRACING

- [Lafortune and Willems 1993, Veach and Guibas 1995]
- **“Brute-force robustness”** – combine many sampling techniques

- Easy-to-use
- Interactive & progressive
- ~~Fast in common scenes~~
- Robust & reliable
- Graceful fail (no artifacts)
- ~~Simple~~
- ~~Compatible with unphysical tricks~~

Bidirectional path tracing (30 min)


Vertex connection and merging (30 min)

VCM/UPS

- [Georgiev et al. 2012, Hachisuka et al. 2012]
- Photon mapping + bidir
- **“Super-brute-force robustness”**
  - combine even more sampling techniques than bidir

- Addresses robustness
- **Makes overhead even worse**

- Easy-to-use
- ~~**Interactive**~~ & progressive
- ~~**Fast in common scenes**~~
- **Robust & reliable**
- Graceful fail (no artifacts)
- ~~Simple~~
- ~~Compatible with unphysical tricks~~

DIRECTION: “LIGHTWEIGHT ROBUSTNESS”

- Selective use of advanced features - fully automatic
- Initial attempt: “Lightweight photon mapping” [Grittmann et al. 2018]

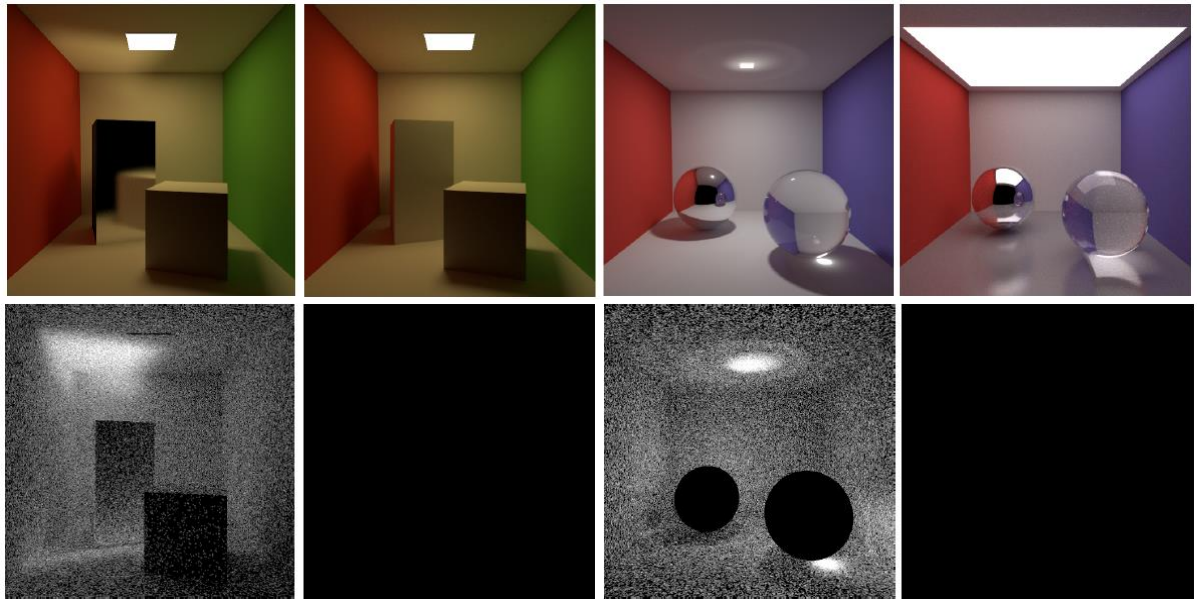

- Easy-to-use
- Interactive & progressive
- **Fast in common scenes**
- **Robust & reliable**
- Graceful fail (no artifacts)
- Simple
- Compatible with unphysical tricks

PATH GUIDING

- [Jensen 1995, Vorba et al. 2014, Mueller et al. 2017]
- **Online adaptive sampling** in path space

- Easy-to-use
- ~~Interactive &~~ progressive
- ~~Fast in common scenes~~
- ~~Robust & reliable~~
- ~~Graceful fail (no artifacts)~~
- ~~Simple~~
- Compatible with unphysical tricks

PATH GUIDING

- [Jensen 1995, Vorba et al. 2014, Muller et al. 2017]
- **Online adaptive sampling** in path space

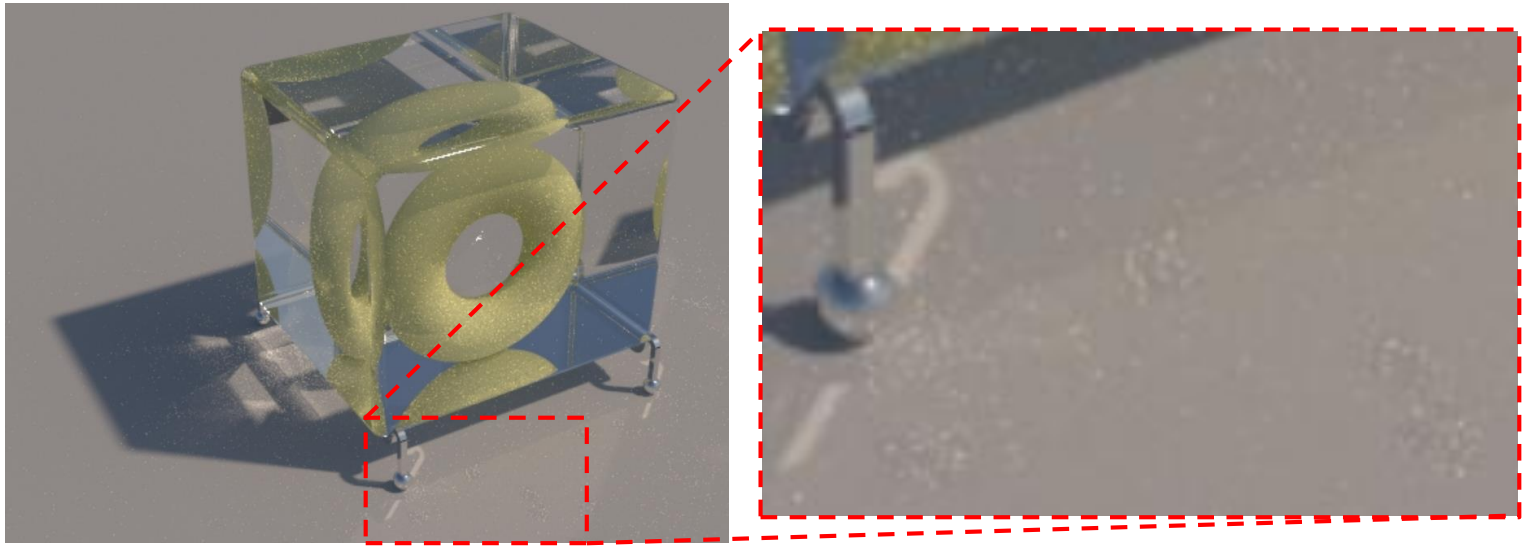

~~Graceful fail~~

- Easy-to-use
- ~~Interactive &~~ progressive
- ~~Fast in common scenes~~
- ~~Robust & reliable~~
- ~~**Graceful fail (no artifacts)**~~
- Simple
- Compatible with unphysical tricks

DIRECTION: “ROBUST ONLINE ADAPTIVITY”

- Adapt to a given scene, but never fail
- Our attempt [Vévoda et al. 2018]
  - specific problem of direct illumination sampling
  - **Bayesian framework for principled online adaptive sampling**

- Easy-to-use
- Interactive & progressive
- Fast in common scenes
- **Robust & reliable**
- Graceful fail (no artifacts)
- Simple
- Compatible with unphysical tricks

Direct illumination only

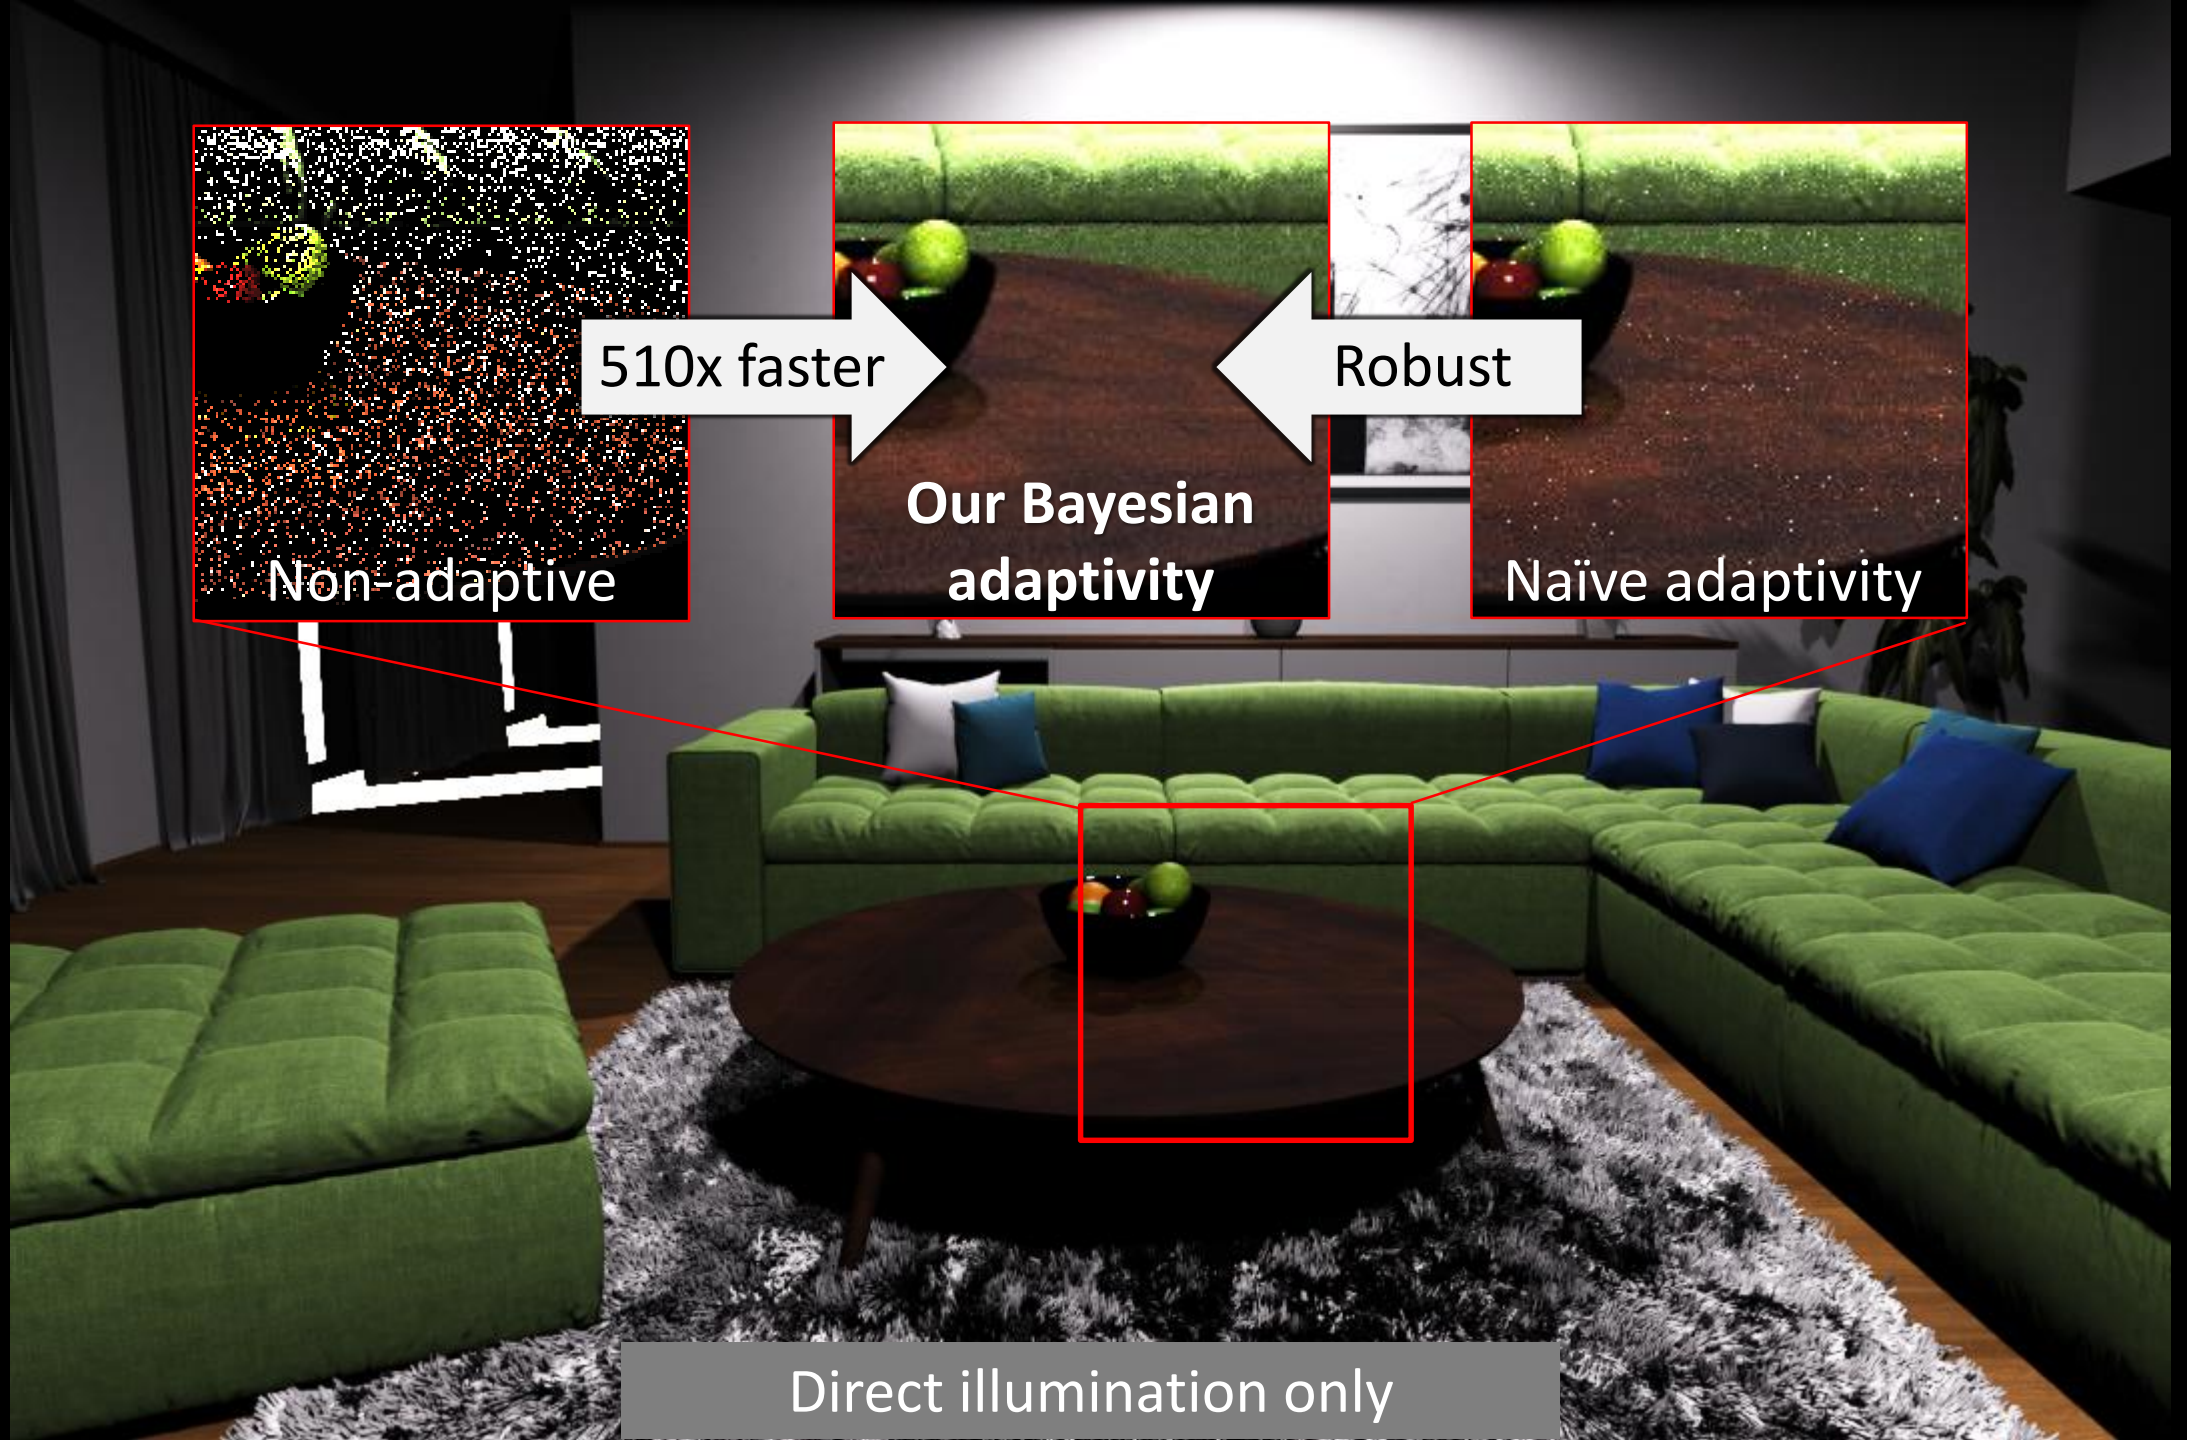
510x faster
Robust
Non-adaptive
Our Bayesian adaptivity
Naïve adaptivity
Direct illumination only

METROPOLIS LIGHT TRANSPORT

- [Veach and Guibas 1997]
- "*… the Metropolis class of algorithms … in my experience these tend to be the most efficient of the unbiased algorithms.*"
  (anonymous reviewer, 2012)
- Problems: uneven convergence, temporal instability

- Easy-to-use
- Interactive & progressive
- ~~Fast in common scenes~~
- ~~Robust & reliable~~
- ~~**Graceful fail (no artifacts)**~~
- ~~Simple~~
- ~~Compatible with unphysical tricks~~

MLT WITH MANIFOLD EXPLORATION

MLT + Manifold exploration [Jakob and Marschner 2012]

Reference

METROPOLIS LIGHT TRANSPORT

- The real source of MLT artifacts unaddressed by papers that build on local differentials
  - [Jakob and Marschner 2012, Kaplanyan et al. 2014, Li et al. 2015]

- Research direction:
  "Global exploration in Metropolis sampling"
  - Some work is there: [Kelemen et al. 2002, Hachisuka et al. 2014, Kaplanyan et al. 2013, Šik and Křivánek 2016, Šik et al. 2016]
  - We need more

- Easy-to-use
- Interactive & progressive
- ~~Fast in common scenes~~
- ~~Robust & reliable~~
- ~~**Graceful fail (no artifacts)**~~
- ~~Simple~~
- ~~Compatible with unphysical tricks~~

PRIMARY SAMPLE SPACE MLT

PSS-MLT [Kelemen et al. 2002]

Reference

METROPOLIZED VCM

Metropolised VCM [Šik et al. 2016]

Reference

LIGHT TRANSPORT RESEARCH

What will it take to make more meaningful progress?

LIGHT TRANSPORT RESEARCH

- Rendering is a mature field

- Practice shows there are numerous unresolved problems
- Underlying algorithm still (the fragile) unidirectional path tracing

- Academic research in rendering hindered by several factors
  - Getting up to speed takes a long time – theory, code
  - Low-hanging fruit taken, the rest labelled “incremental” and often rejected
  - Rendering practice does not communicate its needs to the academia clearly enough

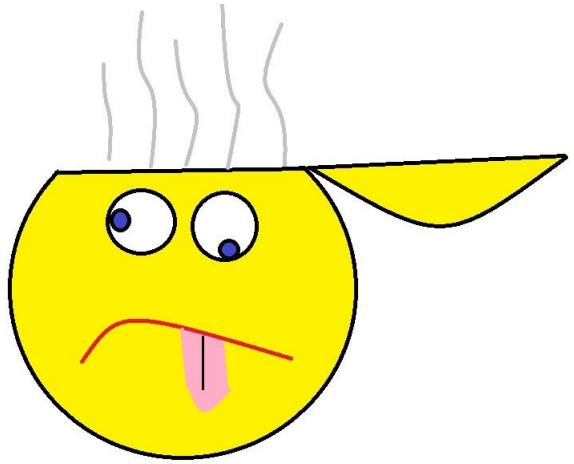

A Msc/PhD student entering the light transport field (official portrait)

SUITABILITY CHECKLIST

- Does the method solve a real problem? Does it address the real root of the problem?
- Does it have reasonable assumptions? (e.g. BRDFs are almost always black-boxes in practice)
- Does one need to set up any parameters?
- Does the method slow down cases where it is not needed?
- Does it fit interactive workflow (no/minimal precomputation)?
- Does the method fail gracefully? (because users always find a way to break it).
- Does the method complicate other code?
- ….

LIGHT TRANSPORT BENCHMARK

- Objective benchmarking of algorithms
- Counteract some current flaws of the publication machine
  - Cherry-picking of results (to meet the 10x speedup expectation)
  - Paper production overhead (extensive comparisons often requested)
- Challenges
  - Calibrate the currently unrealistic **reviewer expectations** (not all papers will bring 10x speedup – that's reality)
  - **Lack of standards** in graphics (even a BRDF of the same name is different in different renderers)
  - Different ways of dealing with **non-physical issues**, such as shading normal/geometry normal inconsistency
  - Lack of commonly accepted objective **image quality metrics** (What is an *'image artifact'*? What is a *'firefly'*?)

DON'T BE ASHAMED TO **GO 'INCREMENTAL'**

- "*If I have seen further it is by standing on the shoulders of Giants.*" (Newton)
- Rendering is a mature field – making breakthroughs is difficult, does not mean it's a solved problem
- Challenges
  - Counteract **reviewer's preference for "fresh ideas"** – reviewers need to be more ready to accept work that address shortcomings of current methods
  - "*I feel that the contribution of this paper is mainly in being a well-written account on how to combine these two algorithms, but that alone doesn't make it SIGGRAPH paper.*" (VCM, Georgiev et al. 2012)
  - **"***The paper represents excellent work … However, It does not represent a completely new direction in GI research.*" (Path guiding, Vorba et al. 2014)

INTRODUCE BIAS! (IN A CONTROLLED WAY)

- Unbiased not useful by itself
- We need more work on trading (lots of) variance reduction for (little) bias
- Denoising already does that (and it does it well). Can we do better?

HOMEWORK FOR THE INDUSTRY

- Communicate needs and open issues more openly and clearly
- Adhere to standards
  - Publish reference implementations
- Share data/resources (kudos to Disney Animation for the Moana dataset)

CONCLUSION

CONCLUSION (MY PART)

- Commercial renderers rely on some fairly old light transport technology
  - Vanilla path tracing ticks many check-boxes
  - Fails for caustics, complex indirect illumination, complex occlusion
- Advanced light transport algorithms rarely used in practice
  - Overhead in simple cases
  - Do not fail gracefully
- Future work
  - Light transport benchmark & calibration of reviewers' expectations
  - Understand real issues of existing solutions, then solve
  - Shift focus to realistic content

THANK YOU!

- Acknowledgments
  - Johannes Hanika
  - Funding: Czech Science Foundation (16-18964S)
  - Colleagues and students from Charles university & Corona developers
  - Course presents
  - You

EXTRA SLIDES

REVIEW QUOTES FOR “FRESH IDEAS”

- “*I feel that the contribution of this paper is mainly in being a well-written account on how to combine these two algorithms, but that alone doesn't make it SIGGRAPH paper.*” (VCM, Georgiev et al. 2012)
- **“***The paper represents excellent work … However, It does not represent a completely new direction in GI research.*” (Path guiding, Vorba et al. 2014)
- “*… little algorithmic novelty is present: the idea of building adaptive distributions is almost 20 years old. This algorithm is a nicer implementation of the same basic idea.*” (Path guiding, Vorba et al. 2014)
- *“There wasn’t … something surprising that I learned from this paper, which I expect from a SIGGRAPH-level submission.*” (Metropolised VCM, Šik et al. 2016)

REVIEWERS JUDGING PRACTICAL MERIT

- "*The authors argument that the presented approach, however, can also handle spatially varying BRDFs and is algorithmically simpler. These are correct, yet not very strong arguments, since* ***spatially … varying BRDFs are a very special scenario****.*" (Filtered importance sampling submission to EGSR 2007)
- "*I'm not convinced long-term, with the use of ultra-fast raytracers, that an approach like this will have small enough overheads to be useful.*" (Path guiding, Vorba et al. 2014)
- "*While many path tracers implement some form of Russian roulette (since we need to terminate path generation at some point), nearly none of these path tracers implement splitting even though Arvo and Kirk introduced them since 1990. There seems to be some big obstacles preventing people from implementing splitting in their path tracers.*" (ADDRS, Vorba and Křivánek 2016)
  - Corona, V-ray, Arnold, etc. all use splitting
- "*… in my opinion the work does not demonstrate enough practical value even for further qualitative theoretical analysis*." (about an algorithm running in Corona for two years, improving performance 1.5x – 2x on average)