



How to Present a Rendering Paper

Morgan McGuire
CS 888 Fall'19
University of Waterloo

[These are lecture notes from a graduate course, which were prepared with feedback from co-teachers Stephen Mann and Craig Kaplan and colleagues Peter Shirley and Ward Lopes]

Recommended Presentation Structure

1. Impact of the paper (intended + hindsight for classic papers)
2. Contributions of the paper ← *add your own emphasis here*
3. Details of the problem
4. Previous work
5. Limitations & drawbacks
6. Key ideas
7. Selected equations, listings, diagrams
8. Show results ← *critique results + evaluation methodology*
9. [*Maybe critique exposition*]
10. Set up and facilitate discussion, leading with the authors' own conclusions

2

Target 20 minutes of you speaking for a ~45 minute presentation, including interruptions and discussion

“I’m not sure that I understood this part” is fine

Periodically stop and invite questions

Use screen-grabs directly from the paper where possible for equations and images (to save your time as well as quote the work)

Choose what to focus on and what to mention only in passing. Do not try to present every aspect or every detail.

Respectfully critique the exposition (What made it easy to understand and should be advocated? What could have been improved, in retrospect?)

Be sure to explain the notation

Some good discussion topics:

- What you didn't understand
- Clever, reusable tricks
- Relationship to other papers
- Potential for followup work

I'll now give a brief example

Prefer Minimal Slides

- ▶ Avoid clip art and irrelevant images
- ▶ Simple template and font
- ▶ Target fewer than 25 words/page
- ▶ Target one important image or equation per slide

3

This slide template isn't bad, although I prefer a completely white or black background...I used one with this simple pattern on the side because there were four slide decks for the first lecture and I needed to be able to distinguish them at a glance! White backgrounds are good because it is easy to put screenshots from paper PDFs against them.

While you don't want to numb your audience into sleep with 20 plain slides of bullets, usually the right solution is to find more content to focus the presentation around and not to decorate the presentation. Sometimes that can be the strategic use of color for emphasis; primarily it will be diagrams, equations, and images. You'll see how I do that when presenting Kajiya's work in a moment.

Feel free to insert a blank slide where you need to talk about something and don't have a visual aid. Don't talk over a slide about something else or feel compelled to insert images.

Avoid transitions and animations, except for builds, which are a great way to pace material within a slide. Remember that the audience may look away and miss an animation—and that they don't always hold up well on video conferencing.

You don't need all of the content on the slide as text. You can use presenter notes to write a transcript if you wish. The slide is a visual aid to the structure of your presentation, not the presentation itself—you actually *presenting* verbally and through gesture, intonation, and pace is the main content of the presentation.

Use References

- ▶ Cite equations and figures
- ▶ Put page numbers on your slides
- ▶ Distinguish the paper from your editorial content
- ▶ Include a bibliography

4

It is important for the audience to understand the origin of material that you are presenting. If it is really obvious that an image is a screen capture from the paper that the entire presentation is about, then you can skip a few citations. But if you are bringing in material from multiple sources or adding your own, then you need to make that clear.

Page numbers on your slides will help the audience to refer back to specific slides during questions and discussion.

You don't *show* the bibliography slide, but it is essential for others reading your presentation later to understand what work your citations are referencing. This can be a screenshot from the paper plus text for references that you added.

Define Key Terms

- ▶ Vocabulary
- ▶ Variables
- ▶ Axes
- ▶ Metrics

5

For technical material, it is important to make clear the assumptions of jargon and notation.

If a word is used in a specific way, such as “specular”, you must define it for the audience. Ideally both verbally and on the slide, so that it is very clear.

All variables that are shown must be defined. I like to do this with colored callouts around an equation from the paper if it does not provide its own notation table. If you feel like defining some of these variables is not a good use of time, then that is a sign that the equation probably should not explicitly appear in the presentation as it is too detailed for the venue.

If there is a result graph or table, tell us how to interpret it before diving into analysis of the results. “The horizontal axis is time...”, “a good result will increase up and to the left...”, “the vertical axis is in log joules, so one tick mark indicates 10x more energy”.

If you show an image, tell us what we are looking at (e.g., “this is a 1m^3 box. The left wall is painted red and the right wall is painted green...”) before analyzing the quality of the result.

Questions before I give an example?

Example



The Rendering Equation

J. T. Kajiya, SIGGRAPH 1986

Presented by A. Student
September 6, 2019

The impact of this paper was enormous. It is probably the most famous paper in physically-based rendering, and underlies that whole field today. It builds on the shoulders of giants but crosses two critical thresholds to make the important breakthroughs.

The paper is written with an understanding of the significance of its core result: framing rendering as a general case for Monte Carlo integral equation solvers, rather than tackling individual phenomena separately.

It downplays the significance of the particular solution of path tracing, so the impact of that aspect comes to us from history and not the paper itself.

Contributions

Cast all rendering as a general case	<i>Conceptual breakthrough</i>
Rendering as Markov Chain Monte Carlo	<i>Conceptual breakthrough</i>
Surface area Rendering Equation	<i>Rarely used</i>
Hierarchical sampling algorithm	<i>Rarely used</i>
Path tracing algorithm	<i>Algorithmic breakthrough</i>
+ Notes on importance sampling	<i>Very important</i>

8

Here are the contributions as presented in the paper.

I consider the generalization of rendering and presenting it in the context of Maxwell's equations to be a conceptual breakthrough for the field. This is "the rendering equation".

The actual form of the equation as given turns out to be less useful than the concept because it has an awkward parameterization over areas.

The hierarchical sampling algorithm is a good idea, but for this application it was supplanted by quasi-Monte Carlo methods and so I won't focus on that today.

The path tracing algorithm generalizes ideas from distribution ray tracing and radiosity to create an algorithmic breakthrough, and the (unimplemented) notes about importance sampling are also very relevant.

Part of the path tracing derivation is showing that rendering is an MCMC problem. The Radiosity algorithm previously leveraged the Markov aspect, but didn't connect it to other work in monte carlo approximations.

Problem

- ▶ Realistic rendering!
- ▶ Arbitrary materials: glossy, matte, transmissive, specular
- ▶ Arbitrary light sources—any surface
- ▶ Efficient simulation
- ▶ Capture all illumination phenomena
- ▶ No “ambient” hacks

Previous Work

Heat transfer [Siegel and Howell '81]

- ▶ *Energy transport as an integral equation for a different field*

Radiosity [Goral et al. '84, Cohen & Greenburg '85, Nishita & Nakame '85]

- ▶ *Limited to perfect Lambertian*
- ▶ *Integral equation formulation; solved by linear algebra*

Distributed ray tracing [Cook et al. 1984]

- ▶ *Efficient sampling of high-dimensional space*
- ▶ *Magic ambient & too much time spent deep in the path tree*

Monte Carlo integration [von Neumann and Ulam via Rubenstein '81]

- ▶ *Core technique not previously applied to rendering*

Questions?

Limitations

Monochrome ray optics

- ▶ No wavelengths (despite RGB implementation)
- ▶ No polarization
- ▶ No phase or diffraction
- ▶ *No time*

Poor convergence for:

- ▶ *Sharp caustics (specular-to-diffuse paths)*
- ▶ *Small light apertures*
- ▶ *Large lights over very shiny surfaces*

11

The paper mentions the limitation to monochrome ray optics of the theory and resulting algorithm.

It doesn't mention (and Kajiya likely was not aware at the time of) the poor convergence properties in scenes where the paths with the most energy transported are hard to discover by searching backwards from the camera or where importance sampling is difficult. [We'll see later that Veach and Guibas address these limitations a decade later with an alternative algorithm.]

Key Ideas

- ▶ Rendering is simulation of Maxwell's integral equations [which can be simplified for this purpose]
- ▶ Light transport is a Markov process with the light surfaces as sources
- ▶ Light is transported along paths
- ▶ [Heuristic:] Equal computation should be expended on direct and indirect illumination at each path node
- ▶ About equal computation should be spent at each depth in the path tree (no branching)

The Rendering Equation

The rendering equation is

$$I(x, x') = g(x, x') \left[\epsilon(x, x') + \int_S \rho(x, x', x'') I(x', x'') dx'' \right]. \quad (1)$$

where:

- $I(x, x')$ is related to the intensity of light passing from point x' to point x
- $g(x, x')$ is a “geometry” term
- $\epsilon(x, x')$ is related to the intensity of emitted light from x' to x
- $\rho(x, x', x'')$ is related to the intensity of light scattered from x'' to x by a patch of surface at x'

13

Kajiya’s own words and notation explain this for you quite clearly. The key point is that this is an integral equation: I is written in terms of the integral of I .

We’re computing the light “intensity” transmitted from point x to x' , which arrived at x from the points x'' in S .

As a guide for the modern reader,

G is the cosine of the angle of the incidence with an extra area term in the denominator, which is typically separated from emittance and put into the integrand today; he had it outside and the unusual units as an artifact of that point parameterization of the integral.

S is the set of all surface points that are visible from x [which will be determined by ray tracing]

Rho is usually written f today—it is the BSDF

Epsilon is the emittance function, with units of W/m^2

I has units of W/m^4 .

Questions?

What is really neat is that...this is it! A little ray tracing and good choice for rho and you have a full renderer. And your code even really follows this structure. Let's see how he turns it into code...

Markov Chain Monte Carlo

If we wish to calculate the value of one coordinate of x , say x_1 , then we calculate the quantity

$$\hat{x}_1 = \left(\prod_{i=0}^{l(\omega)} m_{n_{i-1}n_i} \right) a_{n_l(\omega)} \frac{1}{p(\omega)}$$

averaged over all paths $\omega \in \Omega$. Simply taking expected values verifies that this quantity gives the desired quantity.

14

This is a weighted average of the product of the transport factors, where the weight is the inverse of the probability with which a given path was selected. Note that the denominator here is probability...this is NOT the rho variable from the previous slide.

“ x ” is going to be the intensity (not the point; an awkward notation choice), m is the reflectance and geometry, and a is the light emittance. This formulation assumes that light is only emitted at the end of a path, which doesn’t match the code.

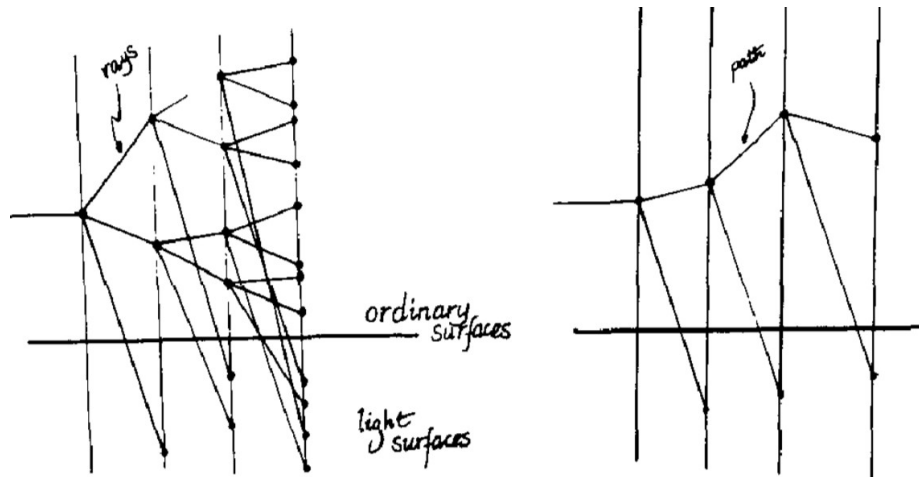
Path Tracing

1. Choose a point \mathbf{x}' in the scene visible through the imaging aperture to a selected pixel \mathbf{x} on the virtual screen.
2. Add in the radiated intensity.
3. For the length of a Markov path do
 - 3.1 Select the point \mathbf{x}'' and calculate the geometrical factor $g(\mathbf{x}, \mathbf{x}')$.
 - 3.2 Calculate the reflectance function $\rho(\mathbf{x}, \mathbf{x}', \mathbf{x}'')$ and multiply by $\epsilon(\mathbf{x}', \mathbf{x}'')$.
 - 3.3 Add this contribution to the pixel intensity.

15

Here's the code. What is implicit is that the nodes on the markov path are found by ray tracing in uniformly random directions.

Path Tracing



16

This is a diagram of a search path through the scene, where the scene is the vertical axis and the iteration depth is the horizontal axis. It is a really wonderful diagram once you understand it.

You can see that previous ray tracing methods spent most of their computational time deep in the tree...which is where the incremental addition to the light intensity is *lowest* [because it is Markov]. Path tracing instead spends about equal time at all levels, and gives equal computation to direct and indirect light (as a heuristic). This is actually more in the spirit of Cook '84 than Cook '84 itself was: Kajiya is treating iteration depth as one of the dimensions of the sampling space, and ensuring that we sample whole transport paths, rather than oversampling deep and undersampling primary rays.

Results: Comparison




Figure 5. A comparison of ray tracing vs. integral equation technique. Note the presence of light on the base polygon scattered by the sphere from the light source.

17

Figure 5 shows a model rendered via two techniques. On the left side is the model rendered via the standard ray tracing technique (albeit with ambient coefficients set to 0 and the single branching ratio speedup mentioned above). The right image shows the result of rendering via the integral equation.

Results: New Phenomena




Figure 6. A sample image. All objects are neutral grey. Color on the objects is due to caustics from the green glass balls and color bleeding from the base polygon.

18

In figure 6 we show an image illustrating the power of the integral equation technique. All objects in the scene are a neutral grey except for the green glass balls and the base polygon (which is slightly reddish). Any color on the grey objects would be missing from a ray tracing image. Note that the green glass balls cast caustics on objects in the scene. There is color bleeding from the lightly colored base polygon onto the bottom of the oblate spheroid in the upper right. For simplicity and comparison purposes, the opaque surfaces in this scene are lambertian, but there is no restriction on the lighting models that can be used.

Results: Performance

Both images are 256 by 256 pixels with a fixed 40 paths per pixel. The images were computed on an IBM-4341. The first image took **401 minutes** of CPU time, the second image took **533 minutes [10 hours]**.

Figure 6 is a 512 by 512 pixel image with 40 paths per pixel. It was computed on an IBM 3081 and consumed **1221 minutes [20 hours]** of CPU time

Exposition

- ▶ Awkward notation
 - ▶ Kajya uses $I(\mathbf{x}, \mathbf{x}')$: $J/(\mathbf{m}^4 \mathbf{s}) = \mathbf{W} \mathbf{m}^{-4}$ because he's working with a point parameterization. An angular parameterization gives the more common radiance $L(\mathbf{x}, \omega)$: $\mathbf{W}/(\mathbf{m}^2 \text{sr})$.
 - ▶ Most of section 2 would be eliminated by the angular parameterization
 - ▶ “Three point reflectance” is today called the bidirectional scattering distribution function and was already in use at the time as the BRDF
- ▶ Section 3 is an exemplary demonstration of generality and related work survey.
- ▶ Brilliant economy: Eq. 1, the path tracing listing, the MCMC derivation, and the path tracing figure could fit on one page and carry the key results

Discussion

- ▶ Why does it *still* take 5-20 hours to path trace an image for an offline film renderer today, if computers are now orders of magnitude faster than the ones Kajiya used?
- ▶ Is extension to wavelengths, polarization, phase, and diffraction really as easy as implied?
- ▶ Integrating over time is fundamentally different than the other quantities, because it affects the ray cast [set S] and not “just” $\rho(x, x', x'')$. How can we approach this?
- ▶ Other applications in rendering for [stratified] hierarchical adaptive importance sampling?

Bibliography

- J.L. BENTLEY, J.H. FRIEDMAN "Data structures for range searching", *ACM Comp. Surv.*, 11,4, pp.397-409, 1979.
- S. CHANDRASEKAR *Radiative Transfer*, Oxford University Press, 1950.
- M.F. COHEN, D.P. GREENBURG "The Hemi-cube: a Radiosity solution for complex environments", *Computer Graphics* 19,3, pp.31-40, 1985.
- R.L. COOK, T. PORTER, L CARPENTER "Distributed ray tracing", *Computer Graphics* 18,3, pp.137-146, 1984.
- R.L. COOK, "Stochastic sampling in computer graphics", to appear in *ACM Transactions of Graphics*
- R. COURANT AND D. HILBERT, *Methods of mathematical physics* 2 vols., Interscience, New York 1953, 1962.
- F.C. CROW "Summed area tables for texture mapping", *Computer Graphics* 18,3, pp.207-212, 1984.
- R.P. FEYNMAN AND A.P. HIBBS *Quantum Mechanics and Path Integrals*, McGraw-Hill, New York 1965.
- H. GOLDSTEIN *Classical Mechanics* Addison-Wesley, Reading, Mass. 1950.
- C. M. GORAL, K.E. TORRANCE, D.P. GREENBURG "Modeling the interaction of light between diffuse surfaces", *Computer Graphics* 18,3, pp.213-222, 1984.
- I.H. HALTON "A retrospective and prospective survey of the monte carlo method", *SIAM Rev.* 12, pp.1-63, 1970.
- J.T. KAJIYA, B. VON HERZEN "Ray tracing volume densities", *Computer Graphics* 18,3, pp.165-174, 1984.
- M.E. LEE, R.A. REDNER, S.P. USELTON "Statistically Optimised Sampling for distributed ray tracing" *Computer graphics* v.19,3 pp.61-67.
- T. NISHITA, E. NAKAMAE "Continuous tone representation of three dimensional objects taking account of shadows and interreflection", *Computer Graphics* 19,3, pp.23-30, 1985.
- R.Y. RUBENSTEIN *Simulation and the Monte Carlo Method*, J.Wiley, New York, 1981.
- R. SIEGEL, J.R. HOWELL *Thermal Radiation Heat Transfer*, McGraw Hill, New York, 1981.
- T. WHITED "An improved illumination model for shaded display", *Comm. ACM*, 23,6, pp.343-349, June 1980.
- S. YAKOWITZ, et. al. "Weighted monte carlo integration", *SIAM J. Num. An.* 15,6, pp.1289-1300, 1978.